

An empirical model of imperfect dynamic competition and application to hydroelectricity storage

Olli Kauppi and Matti Liski*

September 2008

Abstract

The Nordic power market presents a unique opportunity for testing the nature and degree of market power in storage behavior due to preciseness of data on market fundamentals determining hydro resource use. We develop an explicit model of dynamic imperfect competition mapping the primitive distributions to market outcomes as a function of the market structure. We estimate the market structure that best explains the main behavioral patterns in pricing, storage, and production in years 2000-05. Exceptional events in the data allow us to identify a pattern for market power. We simulate the expected efficiency loss from the pattern and find limited scope for social losses. Market power however increases expected reservoir and price levels, and also implies an increase in price risk.

JEL Classification: D43; L1; L9; Q2; Q4

Keywords: storage; hydroelectricity; resources; market power; the Nordic power market

*Kauppi <kauppi@hse.fi> and Liski <liski@hse.fi> are at the economics department of the Helsinki School of Economics. Liski is a faculty associate at MIT-CEEPR. Both authors are grateful for funding from the Nordic Energy Research (NEMIEC and NEECI), and Yrjö Jahnsson Foundation. We are grateful for comments and very constructive help from Pat Bajari, Karl-Axel Edin, Reyer Gerlagh, Mitri Kitti, Maria Kopsakangas-Savolainen, Michel Moreaux, J.-P. Montero, Pauli Murto, Thomas Tangerås, Otto Toivanen, Frank Wolak, Juuso Välimäki, and seminar participants at Catholic University of Chile, Catholic University of Leuven, EAERE-2008, EEA-ESEM 2008, HECER-Helsinki, IIOC-2008, Research Institute of Industrial Economics in Stockholm, Toulouse School of Economics and University of Manchester.

1 Introduction

The Nordic wholesale market for electricity covers the four continental Nordic countries – Finland, Denmark, Norway, and Sweden – which, through their national transmission system operators, own and run a common power exchange, the Nord Pool. Private parties can procure and sell electricity in the Nord Pool, allowing a division of labor for a diverse set of generation technologies including hydro, nuclear, and various forms of thermal power. Hydroelectricity is the key technology in this market. On average, one half of the annual Nordic consumption is met by hydroelectricity but its availability varies widely within and across the years – the annual deviation of water availability can deviate from a typical year by an amount that translates into ca. 1.3 bn € (2 bn \$) using average historical prices. Under multiple uncertainties regarding future inflows, temperature-driven demands, and fuel prices of alternative production, the Nordic hydro power stations face a nontrivial problem of allocating the water stocks between the current and future uses.

We find that the Nordic market presents a unique opportunity for an empirical application of an explicit model of dynamic imperfect competition. In this market, the institutional, technological, and economic framework naturally shape the model structure, thereby leaving relatively little scope for speculations regarding the main ingredients of the model. It is hard to think of other inherently dynamic markets where the 'state' of the market can be measured with similar preciseness. As opposed to many other dynamic markets such as those for aircrafts (Benkard 2004) or cement (Ryan 2006),¹ the producers' dynamic decision is in principle simple: how much water to release and save today? The reduction in the complexity of the economic problem allows us to take steps in empirical matching of the market structure with a quite detailed data.²

We develop a model that is computationally tractable and can map the multiple distributions of market fundamentals into price, output and reservoir distributions as a

¹While much of the literature on dynamic competition aims to capture the evolution of an industry by focusing on entry and exit, the hydro industry is pronouncedly static with respect to its capacity but extremely dynamic and potentially also strategic with respect to usage of the existing capacity. In this market, the static-dynamic breakdown (see Doraszelski and Pakes, 2006) is thus reversed such that the economic problem is further simplified without loss of realism, allowing a more detailed study of dynamic pricing and production for a very sharply defined commodity.

²We note there are not many applications of explicit dynamic models of imperfect competition, despite the considerable recent interest in developing a framework for such testing (e.g., Bajari, Benkard, and Levin 2007, Pesendorfer and Schmidt-Dengler 2008; see Doraszelski and Pakes 2007 for a review).

function of the market structure, which allows us to choose the structure that best fits with the historical data. The approach is not specific to the Nordic market and therefore applies to market power issues in storable-good markets, and electricity markets with hydro technologies more generally. This paper is also the first explicit attempt to evaluate the best-fitting market structure in hydro use in the Nordic market and among the first in general.³ We find a pattern for market power and evidence that it systematically distorted the reservoir levels during the years 2000-2005. The model can explain the main behavioral patterns in pricing, storage, and production, and 90 per cent of the estimated welfare loss. We also simulate the expected long-run social loss from such behavior and find an extremely low number: the best-fitting market structure increases the expected average price of electricity by less than 1 €/MWh. This leads us to conclude that the scope for social losses from imperfect competition is not large in the Nordic market for hydroelectricity, but at extraordinary events, such as the shortage of water availability in 2002, large sporadic deviations from the first-best outcome can occur.

Market power in storable-good markets has traditionally been notoriously difficult to detect because price-cost margins depend on expected future market conditions that cannot be observed *ex post* – in a pure storage decision such as the hydro release, the marginal cost is only the opportunity cost from not being able to sell the same unit in the future. Thus, to evaluate the price-cost margins, one needs to evaluate the expected future values at the state of the market where the decision is made. Perhaps for this reason, while there is a well-developed theory on competitive storage,⁴ there is little work on market structure and storage and, in particular, empirical applications or tests are practically nonexistent.⁵ The Nordic market gives us an advantage in testing the effectiveness of storage behavior. As an electricity market, the Nordic market is, and has been, subject to a regulatory oversight, providing a wealth of data that we can use to estimate relatively objectively how market participants should view the market fundamentals such as inflows, demands, and thermoelectric supply.

³Evaluating market power in the Nordic market requires a framework for imperfect competition in hydro use. We are the first to provide such a framework and its empirical application. Amundsen and Bergman (2002) and (2006), and von der Fehr, Amundsen and Bergman (2005) provide valuable analysis of the issues relevant in the Nordic market.

⁴The work by Williams and Wright is summarized in their book (1991); see also Deaton and Laroque (1992) and (1996).

⁵McLaren (1999) builds on Newbery (1984) to describe a Markov perfect equilibrium in an oligopolistic storage market. Market power leads to reduced storage levels and increased price risk. Rotemberg and Saloner (1989) include storage as a strategic device supporting collusive oligopoly equilibria.

We depart from other studies of market power in electricity markets in that the focus is on the long-run usage of capacity rather than short-run market power in the spot market. In the Nordic market, the hydro stocks are long-lived and the main market fundamental determining how the division of labor between capacity types evolves within and between the years. The stocks create a firm link between the current spot prices and expected future prices, thereby stipulating efficiency analysis of the long-run price levels. Studies of other early deregulated electricity markets focus on the short-run market power, for which electricity markets provide an interesting case: there is relatively precise engineering (expert) data on marginal costs, allowing a direct evaluation of price-cost margins from price-quantity data. This approach has been used by Wolfram (1999) in the British electricity market, and by Borenstein, Bushnell, and Wolak (2002) in the California's market; in later work, for example, Hortaçsu and Puller (2008), Puller (2007), and Bushnell, Mansur, and Saravia (2008) put more focus on the market structure.

A hydro-dominated market requires a very different methodological approach from that used in the previous work on electricity markets. Any attempt to include hydro as a part of the aggregate marginal cost curve will require a behavioral element in the analysis because one must solve the equilibrium valuation of water; this value does not exist as a primitive input in an expert data set. A realistic computation of the socially optimal water values is in most cases a large scale numerical problem. To obtain a realistic benchmark for our market power analysis, we first develop an aggregative model of competitive storage, where data inputs include 52 weekly distributions for both inflows and consumer demand estimated from historical data as well as weekly supply curves for other technologies. The hydro demand is then constructed as the residual using the consumer demand and nonhydro supply curve. In this procedure, we must estimate how the nonhydro capacity is supplied in each potential future state of the market; otherwise one cannot form expectations determining the value of the current storage. This is an important difference to the past studies based on expert data sets on marginal cost curves.⁶

A model set up this way can be used to map the primitive distributions of market

⁶Our paper is a natural extension to the literature on hydroelectricity markets in that we present the first explicit empirical model of imperfect competition that is fitted to the market outcomes. Analytically, Crampes and Moreaux (2001) show that market power can be exercised by exploiting differences in demand elasticity of different periods, without spilling water. Bushnell (2003) finds potential for such behavior in a numerical multi-period short-term Cournot game calibrated to the western United States electricity markets. Scott and Read (1996), Garcia, Reitzes and Stacchetti (2001), and Thille and Genc (2008) also study a mixed hydro-thermal system with market power.

fundamentals and nonhydro supply curves to socially optimal weekly price, output and reservoir distributions. The moment properties of the price distributions reveal that the Nordic market has features of an exhaustible-resource market. About 50 per cent of the annual inflow is concentrated to Spring weeks, leading to a market arbitrage that seeks to use this endowment to equalize expected discounted prices until the next Spring. Indeed, the socially optimal expected market price increases at a rate very close to the interest rate throughout the hydrological year, while in the end of the year the price is expected to drop at the arrival of the new allocation. The market has also features of a traditional storage market: favorable demand-inflow realizations lead to storage demand and savings to the next year. Towards the end of the hydrological year weekly price distributions have moment properties familiar to those observed in other storable-commodity markets.

Using the socially optimal policy we can evaluate the historical market experience in 2000-2005, a period over which the economic environment was relatively stable. We find a 7.3 per cent welfare loss, or that the cost of meeting the same demand could have been 636 mill. € lower. We also find a systematic deviation between the socially optimal policy and the market usage of water: the reservoir target levels are systematically different leading to a market shortage of water in late 2002 and to a considerable price spike.

When developing the model of dynamic imperfect competition, we keep the primitives of the socially optimal framework but change the behavioral assumption: some fraction α of the total reservoir and turbine capacity is assumed to be strategically managed, and the remainder of the hydroelectricity generation is competitive. We do not have data detailed enough to map actual firm level capacities into the model, and given the dimensionality of the problem, this approach would render the model intractable. Our dominant firm (or cartel) approach is pushing the computational limits while still being an explicit model of dynamic competition.

The existing techniques for finding the underlying structural parameter (the capacity share α) do not directly apply in our dynamic game.⁷ The computational problem is caused by the need to evaluate the market expectations of the behavior of the large firm in each possible state. We develop an algorithm for solving this fixed-point problem, and then solve the game through a large backward-induction exercise. By repeatedly solving the game for varying α -values, we find a mapping from primitive distributions

⁷We considered using the Bajari, Benkard and Levin (2007) approach, where the market policy is first estimated and then used to recover structural parameters. However, the hydro policies are conditional on a rich set of variables, and it is difficult to estimate these relationships with preciseness that can lead to an outcome competing with our direct method.

plus market structure to weekly price, output and reservoir distributions. We then use the Generalized Method of Moments for the three moment restrictions to find the best-fitting market share parameter.

We find that the market structure where 30 per cent of the storage capacity is strategically managed provides the best match with the historical data. The result is robust to various forms of data aggregation (weekly, monthly, quarterly, or semi-annual aggregation). To evaluate if some unobserved or mismeasured factors can produce a similar match with the data, we force the competitive behavioral assumption and estimate structurally the best-fitting constraints in the hydro system, the discount rate, and out-of-sample expectations for demand and inflows. Sufficient adjustment of both lower and upper limits on available hydro capacity can almost match the fit provided by our behavioral assumption, but with gross deviation from what the data indicates for the available capacity.

How is the market power then exercised? We rule out (or penalize) spilling of water since such behavior is easily detected. Under this constraint, the current availability can be reduced by shifting supply to the future, thereby increasing the expected reservoir levels as well as prices and price risk. However, in expected terms the social loss from such behavior is extremely low: the best-fitting market structure increases the expected average price of electricity by merely 1 €/MWh. The reason for the relatively large loss estimated from the historical data is that the market experienced an inflow shortage in late 2002 that occurs on average once in every 200 years. Such extraordinary events provide a unique opportunity for exercising market power, and this is what our model predicts: the model can replicate the price shock experienced and explain 90 per cent of the welfare loss.

The paper is structured as follows. In Section 2, we provide an overview of institutional framework and the market fundamentals that are the main ingredients of the model. In Section 3, we describe the formal model used in the socially optimal hydro allocation problem. While complicated due to multidimensional state and uncertainties, it is a standard stochastic dynamic programming problem. The model is general enough to give traditional storage and exhaustible resource models as special cases, but when specified to match the power market framework, the implications become specific to this market. We explain how this model is calibrated and discuss the properties of the socially optimal path in detail. In Section 4, we formally develop the alternative market structure that is then, for a given α , calibrated similarly to the socially optimal model (with some increase in computational complexity). We develop a test statistic and search for the best

matching α , and also explain the implications of market power in this storage market. In Section 5, we test for the robustness of our results by studying whether the observed behavior could be explained by socially optimal hydro use under alternative parameterizations of the model. The final section concludes and discusses the shortcomings of the approach.

2 Institutions and market fundamentals

2.1 System price

In this section we give an overview of the institutional and market environment; the data and its sources will be discussed in detail in Section 3.4. The Nordic wholesale power market developed to its current form through a series of steps, as the four continental Nordic countries (Finland, Denmark, Norway, Sweden) underwent electricity market liberalization at different times in the 1990's. Full integration was achieved in October 2000, when East Denmark was integrated into the market. Wholesale electricity trade is organized through a common pool, Nord Pool, a power exchange owned by the national transmission system operators.⁸ Market participants submit quantity-price schedules to the day-ahead hourly market (Elspot market).⁹ The demand and supply bids are aggregated, and the hourly clearing price is called the system price. The Nordic market uses a zonal pricing system, in which the market is divided into separate price areas. If the delivery commitments at the system price lead to transmission congestion, separate price areas are established. However, we do not focus on the hourly electricity market but define the relevant market at the weekly level. Our objective is to analyze hydro storage for which extraordinary events may have ramifications over several years and, given this objective, we define prices as well as other economic variables as weekly averages. Decisions in an hourly market do not lead to significant changes in hydro stocks and, therefore, one is forced to aggregate over hours to make the dimensions of stocks and flows relevant for the analysis. At this level of aggregation, there are good reasons to argue that the Nordic area is a relatively well integrated electricity market. The Nordic market forms a single price area for a significant fraction of time, as indicated by Table 1 which shows deviations from the system price for the main price areas as percentage

⁸For more information, see www.nordpool.com. For a concise description of the Nordic market, see Amundsen and Bergman (2006).

⁹The day-ahead Elspot market is the relevant spot market. While there is a real-time market (Elbas market) closing an hour before delivery, volumes in the Elbas market are small relative to the Elspot.

departures in weekly averages. About 94 per cent of the hydro resource stocks are located in the Norwegian and Swedish price areas, in which deviations from the system price are on average the smallest. It would be difficult to choose any other price than the system price as the reference price for hydro storage decisions.¹⁰

Figure 3 shows the weekly system price over the six years 2000-05, and an estimated price that we discuss later. We focus on this period when matching the model with the data because the institutional and economic environment was relatively stable; that is, the market was not yet affected by the European emissions trading scheme and further integration to the continental Europe. To present a snapshot of the market development, we may call the years 2000-01 as years of abundant availability of hydroelectricity which is reflected in the prices of Figure 3. The year 2002 in turn was exceptional: the Fall rainfall and thus inflow was scant and the stocks were drawn down to approach historical minimums by the turn of the year. The price spike resulted, and it took almost two years for the stocks to recover.

2.2 Capacities

The attraction of a joint Nordic power market is due to the favorable mix of generation technologies resulting from the integration of the national markets. Roughly one half of annual Nordic generation is produced by hydro plants. In 2000-05, 61 per cent of hydroelectricity was generated in Norway and 33 per cent in Sweden.¹¹ Sweden is the largest producer of thermoelectricity with a share of 46 per cent of annual mean production, followed by Finland and Denmark, with shares of 35 and 19 per cent, respectively. The direction of trade between the countries varies from year to year, depending mainly on the availability of hydroelectricity. In years of high precipitation, the hydro power is exported from the hydro dominated regions to Denmark and Finland. In these years, a sizeable fraction of total thermal capacity is idle through much of the year. When inflow is scarce, the flow of trade is reversed, and power is exported from the thermally intensive

¹⁰The direction of congestion in the transmission links varies from year to year depending on the division of labor between hydro-intensive and thermal-intensive regions in the market. Thus, the frequency with which hydro producers receive a price deviating up or down depends on the state of the market and, in principle, one could estimate the expected departure in the price and then use this information when evaluating the hydro producers' behavior. In the current paper, we do not model the hydro resource stocks in different price areas separately and, therefore, cannot incorporate the information about area price differentials in a meaningful way.

¹¹The capacities cited here are reported by the Organisation for the Nordic Transmission System Operators (www.nordel.org) unless otherwise noted.

regions to Norway.

Hydro availability therefore is the one single market fundamental that would alone cause considerable price volatility within and across the years even without other sources of uncertainty. Figure 1 depicts the mean and the empirical support for aggregate weekly inflow over the years 1980-1999. The mean annual inflow in the market area was 201 TWh of energy, and the maximum deviation from this -49 TWh in 1996. This difference translates into a value of ca. 1.3 billion € using the average system price in 2000-05.

Within-the-year seasonal inflows follow a certain well-known pattern, as illustrated by Figure 1. The hydrological year can be seen to start in Spring when expected inflows are large due to the melting of snow; on average 50 per cent of annual inflow arrives in the three months following week 18. The aggregate reservoir capacity in the market is 121 TWh, or 60 per cent of average annual inflow. There are several hundred hydro power stations in the market area, with a great variety of plant types. At one extreme, the run-of-river power plants have no storage capacity, and usually produce as much electricity as the current river flow permits. At the other extreme, there are power stations connected with one or more large reservoirs, that may take months to fill or empty. In 2005, the total turbine capacity of the hydro plants was 47 445 MW, or 72% of peak demand. Hydro production is also constrained by environmental river flow constraints. These constraints together with the must-run nature of the run-of-river plants bound the hydro output from below.

For our empirical application, it is important to emphasize the following features of the hydro system. First, there is an almost deterministic inflow peak in the Spring: in our historical data, the Spring inflow has never been less than one third of the mean annual inflow. In this sense, at the start of each hydrological year, the market receives a reasonably large recurrent water allocation that must be depleted gradually. The annual consumption of this exhaustible resource has marked implications for the equilibrium price expectations, as we will explicate. Second, the remaining annual inflow, on average 50 per cent, is learned gradually over the course of the Fall and Winter. This uncertainty is important for the storage dynamics over the years: abundant Fall inflow, for example, can lead to storage demand and savings to the next year; in case of shortage, a drawdown of stocks can take place. The Nordic market for water can be seen, on one hand, as an *exhaustible-resource market* and, on the other, as a *storage market* for a reproducible good. For understanding the potential for market power, it is important to understand these two interpretations, as we will see. Third, the reservoir, turbine, and various flow constraints for production affect the degree of flexibility in using the overall hydro

resource. We take an estimate for these constraints from the data and previous studies, but we also structurally estimate a set of constraints best fitting the data. The purpose of this procedure is to distinguish the effect of potentially mismeasured constraints on the equilibrium from the effect of potential market power.

2.3 Demand for hydro

Like hydro inflow, the overall electricity demand follows a seasonal pattern, which is closely temperature related. Figure 2 depicts the mean demand and empirical support over the weeks of years 2000-2005. The relevant concept of demand for the purposes of this paper is the residual demand for the hydro: when consumer demand is given, the supply from non-hydro technologies determines the residual demand for hydro. In the Nordic area, the non-hydro production capacity consists of nuclear, thermal (coal-, gas-, biofuel-, waste- and oil-fired plants), and wind power. An important part of thermal capacity is combined heat and power (CHP) plants which primarily serve local demand for heating but also generate power for industrial processes and very cost-efficient electricity as a side product. An implication of CHP capacity is that the non-hydro market supply experiences temperature-related seasonal shifts, which we seek to capture in our estimation procedure detailed later. Table 2 provides a breakdown of capacity, number of plants, and the utilization rates of the capacity forms over the period 2000-2005. At the market level, there is thus a rich portfolio of capacities with large number of plants in each category determining a relatively smooth supply function or, alternatively put, a smooth residual demand function for hydro.

The elasticity of this residual demand is almost exclusively determined by the slope of the non-hydro supply curve because the consumer demand is insensitive to short-run price changes. For this reason, in the analysis we will take the consumer demand as a given draw from a week-specific distribution that we estimate from the data. The industrial consumers have more flexibility in responding to short-run price changes, but their own generation capacity is included as part of the overall market supply curve and, therefore, their price responsiveness is accounted for.

3 Socially efficient allocation

3.1 The model

We describe now the socially optimal resource allocation problem. This way we introduce the basic elements of the model which, for the most part, remain the same throughout the rest of the paper.

Time is discrete and extends to infinity, $t = 0, 1, 2, \dots$. One year consists of 52 discrete time periods. It will be important to keep track of the periods within a year, and therefore we introduce another time index for the week, ω . Let S_t denote the aggregate hydro stock (measured in energy) in the reservoir, x_t is the demand for energy, and ω_t is the week at t . State, denoted by s_t at t , is the vector

$$s_t = (S_t, x_t, \omega_t).$$

The timing of decisions within period t is the following:

1. state s_t is observed;
2. water usage from the stock, denoted by u_t , is chosen;
3. residual demand $z_t = x_t - u_t$ is met by non-hydro production;
4. inflow available at $t + 1$ is realized.

In the empirical application the key variables are discrete and defined on a finite grid, and this is what we assume also for the theory model. In particular, the action set $u_t \in U(s_t)$ is finite as well as the possible physical state space for S_t . Choices are constrained, e.g., by the availability of water, reservoir and turbine capacity, and river flow restrictions.

Demand realization is drawn separately for each week from a week-specific distribution:

$$\begin{aligned} x_t &\sim G_\omega(x), \\ \omega &= \omega_t \in \{1, \dots, 52\}, \end{aligned} \tag{1}$$

where G_ω is a cumulative distribution function (CDF) on some finite set of outcomes X_ω (each element bounded). An alternative to this formulation would be to assume week-by-week realizations of demand schedules depending on price, incorporating demand

elasticity in a more realistic manner. However, the analytical loss is small since for our purposes the interesting elasticity is given by the residual demand for hydro. This elasticity is to a large degree determined by the slope of the non-hydro supply curve. Yet another formulation would be to include persistence in seasonal shocks, as high demand in some week due to a cold spell may have implications for the next week's demand. Since we are uncertain about the relevance of this phenomenon in the Nordic area, we do not want to expand the state space by assuming correlated shocks in demand.

Production by other than hydro capacity has a week-specific aggregate cost curve

$$C : \omega \times z \longrightarrow R_+^1$$

which is increasing in z each week ω . We denote the weekly cost by $C_\omega(z)$. As explained, the seasonal variation comes from the availability of CHP capacity and from the maintenance pattern for nuclear and large coal plants. The definition of $C_\omega(z)$ incorporates the level of fuel prices and we could also include changing fuel prices explicitly. Indeed, we solve the planner's model under a stochastic fuel-price process when we evaluate the robustness of the results in Section 5. However, fuel prices are not structural variables of the Nordic market in the same sense as inflow and demand are because we cannot estimate fuel price distributions with the same accuracy. We find it important not to mix fuel prices with the market fundamentals because, as will be demonstrated, excluding the fuel price uncertainty has little effect on the predicting power of the model. Thus, we set up the benchmark model with a cost function depending on supply z and period ω only.

The final stochastic element of the model is the water inflow which we denote by r_t . The inflow at t is observed only after the hydro usage u_t is chosen but it is observed before the choice of the next period water use u_{t+1} . The inflow realization is, like demand, drawn separately for each week from a week-specific distribution:

$$\begin{aligned} r_t &\sim F_\omega(r), \\ \omega &= \omega_t \in \{1, \dots, 52\}, \end{aligned} \tag{2}$$

where F_ω is a CDF on some finite set of outcomes R_ω (bounded elements).

Finally, the physical state, i.e. the hydro stock, develops according to

$$S_{t+1} = \min\{\bar{S}, S_t - u_t + r_t\} \tag{3}$$

where we include the reservoir capacity \bar{S} . Any inflow leading to a stock exceeding \bar{S} is spilled over and left unused. The next period stock cannot go below a nonnegative lower

bound \underline{S} ; this constraint will be implemented through the choice set $u_t \in U(s_t)$. Now, if we fix a policy rule $u_t = g(s_t)$ and start from a given state s_0 , the development of the state vector s_t is fully determined by the stochastic processes for x and r , and by the law of motion for S_{t+1} . To determine the optimal policy, we define next the per-period payoff for the decision maker at each t as

$$\pi(s_t, u_t) \equiv -C_\omega(x_t - u_t).$$

Maximizing π is equivalent to minimizing the cost of non-hydro production. If we let β be the discount factor per period, the optimal policy $u_t = g(s_t)$ maximizes the discounted sum of the expected per period payoffs, or alternatively put, minimizes the social cost of meeting the current and future demand requirements generated by (1). Let $v(s_t)$ denote the maximum social value at state s_t . This value satisfies the Bellman equation

$$v(s_t) = \max_{u_t \in U(s_t)} \{\pi(s_t, u_t) + \beta E_{s_{t+1}|s_t} v(s_{t+1})\}.$$

Note that the existence of the optimal policy follows directly from the Blackwell's Theorem because the rewards are bounded and the state space is finite (see Stokey et al. 1989).

In the empirical application, all production is dispatched by market clearing in a spot market, where the residual demand $x_t - u_t$ is left for non-hydro producers. If the market is competitive,¹² it is cleared through bidding such that the spot price satisfies

$$p_t = C'_\omega(x_t - u_t).$$

We express the socially optimal hydro dispatch policy immediately in terms of the (socially optimal) market price p_t because the price will give (or approximate due to discrete action space) the shadow cost of not using a unit of water in the current period. Using the optimal policy $u_t = g(s_t)$, we see that the state s_t follows a stationary Markov process, and therefore it generates a stationary weekly price distribution. Let $p_t = p_g(s_t)$ denote the socially optimal price following when optimal policy g is applied at state s_t . As $t \rightarrow \infty$, we obtain a limiting week-by-week distribution for the state vector by the stationarity of the underlying Markov process, and thereby also a limiting week-by-week distribution for the prices:

¹²In the empirical part, we estimate the non-hydro supply from data without invoking competitive behavior. Thus, $C'_\omega(z)$ is interpreted as the inverse supply curve rather than the true marginal cost curve. See Section 3.4 for detailed discussion.

$$\begin{aligned}
p_t &\sim P_\omega(p), \\
\omega &= \omega_t \in \{1, \dots, 52\},
\end{aligned}
\tag{4}$$

where $P_\omega(p)$ is the discrete CDF on some finite set of possible prices.

Denoting the first moments of the long-run weekly price distribution by μ_ω , from (4), we can describe the basic economic logic of the equilibrium using the long-run price distribution. The model allows various interpretations, depending how the market fundamentals are specified.

3.2 Interpretations

Exhaustible-resource interpretation. Suppose the long-run price moments satisfy

$$\mu_1 = \beta\mu_2 = \dots = \beta^{51}\mu_{52} > \beta^{52}\mu_1,$$

a situation that can arise, e.g., when the annual inflow is concentrated to the first week (or to some other week initiating the hydrological year). Then, the allocation problem is effectively an exhaustible-resource problem within the weeks of the year, equalizing the expected present-value prices across the weeks but not across the years: the new inflow at the beginning of the year makes the resource reproducible. Assuming that the decision maker indeed has enough flexibility to equalize expected prices within the year (to be discussed in detail below), the drop in the expected price must arise at the turn of the year as long as there is expected annual scarcity.

Storable-good interpretation. The long-run price moments can satisfy

$$\mu_{\omega_t} > \beta\mu_{\omega_{t+1}},$$

for all weeks when the weeks are relatively similar in terms of inflow and demand for hydro. In this situation, the equilibrium progresses as in standard competitive commodity storage models (Williams and Wright, 1991): inventories are held to the next period after relatively favorable inflow-demand conditions, implying storage demand up to the point where the current price equals the expected next period price, $p_t = \beta E p_{t+1}$; when the current inflow-demand conditions are relatively unfavorable, stockout may take place, and $p_t > \beta E p_{t+1}$. However, when periods are ex ante similar in terms of inflow and demand, the expected long-run storage cannot be positive and the price means satisfy $\mu_{\omega_t} > \beta\mu_{\omega_{t+1}}$. Consistent with this reasoning, the long-run price distribution is skewed

as the storage demand eliminates extremely low prices that would arise when storage is not allowed (see also Deaton and Laroque, 1991).

When the market fundamentals are estimated from the Nordic market data, we observe that both of these interpretations are useful. The socially optimal long-run prices support the exhaustible-resource view of the expected year but the storage market view describes well the decisions at the annual level.

3.3 Characterization

The long-run price means are useful in conceptualizing the nature of the market, but the realized price sequences may follow a logic that can be difficult to relate to the long-run price distributions. For ease of interpretation of the empirical results, we explain next how the state-dependent optimal policy, the current price, and the market fundamentals are linked.

Consider the optimal policy $g(s_t)$, and let $d_t = d(s_t)$ be an alternative policy that deviates from $g(s_t)$ only at current t ,

$$d(s_t) = \Delta + g(s_t),$$

where $\Delta \neq 0$ and coincides with $g(s)$ at all other dates and states. We can define

$$\bar{p}_t = \bar{p}(s_t, \Delta) = \frac{\pi(d(s_t)) - \pi(g(s_t))}{\Delta}$$

as the average cost change caused by the one-shot deviation Δ . Recall that the grid for actions determines the smallest feasible Δ ; when Δ is small, then $\bar{p}(s_t, \Delta)$ is approximately equal to the market price, p_t . We can thus interpret \bar{p}_t as the approximate price in the following:

Proposition 1 *Assume there is an alternative policy to $g(s_t)$ at s_t , i.e., $\Delta \neq 0$ and $d_t \in U(s_t)$. Price \bar{p}_t and the alternative have the following relationship:*

$$\Delta > 0 \iff \bar{p}_t \leq \beta^k E_t \bar{p}_{t+k} \text{ for some } k \geq 1. \quad (5)$$

$$\Delta < 0 \iff \bar{p}_t \geq \beta^{k'} E_t \bar{p}_{t+k'} \text{ for some } k' \geq 1 \quad (6)$$

Proof. See Appendix. ■

In the empirical application, feasible choices are constrained, e.g., by storage and turbine capacity, water availability, and river flow restrictions. When these constraints allow a deviation upwards from the optimal policy at state s_t , i.e. $\Delta > 0$, then the

cost saving today, given by \bar{p}_t , is weakly lower than the expected loss from future cost increase implied by increased usage today. That is, the current "price" is lower than some expected future discounted "price". Similar reasoning holds in the other direction.

When inflow and demand distributions for hydro vary widely across weeks, the set of conceivable prices can shift from one period to the next, and there is no general way of achieving the present-value price equalization. Even when the optimal policy is unconstrained in equilibrium, i.e., it is possible to use or save more water at state s_t , the current price can be lower than some expected future price

$$p_t < \beta^k E_t p_{t+k}$$

and higher than some other expected future price

$$p_t > \beta^{k'} E_t p_{t+k'}.$$

This pattern in no way contradicts Proposition 1. The optimal policy seeks to minimize the difference in expected present value prices but no price equalization is guaranteed. For this reason the long-run price moments can satisfy

$$\mu_\omega \leq \beta \mu_{\omega+1}$$

over some weeks when, for example, inflow is high in week ω so that the storage capacity is likely to be binding. Then, in expectations water is frequently dumped to the market in that period. Alternatively, expected demand may be high enough to frequently require maximum production in week ω but even more so in the next week $\omega + 1$. Finally, minimum flow requirements at low demand periods can bias price moments downwards from what would otherwise hold for some particular weeks.

3.4 Calibration of the benchmark model

In this section, we describe the data and the estimations needed for the calibration of the planner's model. Here, we calibrate the model as suggested by the data, but in Section 5 we re-evaluate the data inputs and the distributional assumptions using a structural estimation procedure. The data, estimations, and the program for computing the model are available at the authors' webpage. We use weekly observations from the six years 2000-2005 which is a period over which the institutional and market environment was relatively stable.

For demand, we use weekly demand data for the Nordic market in 2000-05 as published by the Organization for Nordic Transmission System Operators. As explained earlier, in a

given week, the consumer demand is assumed to be inelastically drawn from the demand distribution. We assume that demand is normally distributed with the weekly means and standard deviations computed from the data.¹³ The distribution is then mapped to a finite grid. The step length of the grid was fixed at 200 GWh, leading to an average of 5.4 demand states per week¹⁴. The weekly support of demand in the model follows the empirical support as observed in the data.

Inflow energy is assumed to be log-normally distributed, and the parameters of the distributions are estimated using data from the period 1980-1999. National inflow data is published by Norwegian Water Resources and Energy Directorate (NVE), Swedenergy and the Finnish Environment Institute. As with demand, inflow is mapped to a finite grid, with an average of 27.5 possible inflow levels per week.

Hydroelectric generation is represented by a single reservoir and power plant, and we use the aggregate market reservoir capacity of 120 TWh and the aggregate weekly turbine capacity of 7.9 TWh as the key parameters of the hydro sector. There is no publicly available information about minimum flow constraints but, after presenting the main results, we experiment with different levels of minimum production. For the minimum reservoir level, we use a lower bound of 10 TWh for the whole Nordic system.¹⁵

For the residual demand of hydroelectricity, we can follow two routes. We can use engineering data on the fleet of non-hydro power plants in the Nordic area to build an aggregate marginal cost curve.¹⁶ Using this data we can in principle follow the approach from Wolfram (1999), also used in Borenstein et al. (2002), to construct the theoretical supply curve for nuclear and thermal plants. In this market the theoretical non-hydro

¹³Demand for electricity showed little trend growth over the sample period.

¹⁴For example, demand varies between 8.2 and 9.6 TWh in the first week of January. All variables measured in energy must be discretized using the same step length to keep track of the evolution of the reservoir level. Thus, while a finer grid for demand might seem plausible, decreasing the step length would also increase the reservoir space. The current choice of step length is determined by the computational burden and memory requirements of the market power model.

¹⁵The lower bound of the aggregate reservoir level is based on the importance of the hydro resource as a fast power reserve supporting the electrical system. Bye et al. (2006) refer to a statement by the NVE, according to which the actual minimum level of Norwegian reservoirs was 8 TWh in the spring of 2003. Nordel uses 5% (6 TWh) of total reservoir capacity as the lower bound for aggregate reservoir level in the simulations of its Energy Balances publication (Nordel 2006). Amundsen and Bergman (2006) refer to a total minimum reservoir level of 15 TWh in 2002, and to 12 TWh in 2003.

¹⁶A data set containing all plants of relevant size in Finland, Sweden and Norway has been collected by the firm EME Analys for use with the PoMo market simulation model. We thank Per-Erik Springfeldt and Karl-Axel Edin for sharing this data with us.

supply curve experiences considerable seasonal shifts because of heating demand (making electricity a side product) and planned maintenance outages. Moreover, for the hydro usage decisions we need to know the expected future supply of the non-hydro power; the value of water in a given state can be computed only by evaluating its value in possible future states. At this point, the expert data set becomes dependent on subjective assessments of patterns in capacity availability and maintenance.

For the above reason, we rather estimate the seasonal supply of the non-hydro capacity than use the engineering data. We thus estimate the weekly supply function of the thermal sector from data on the weekly system price and total demand in 2000-05. A conceptual difference to Wolfram (1999) follows: by estimating the thermal (all non-hydro) supply from the data, we include all the strategic distortions that may exist in this part of the market (nevertheless, it is a conceptually valid approach to evaluate the efficiency of hydro use separately, given the behavior of the thermal sector).

The system price data is published by Nord Pool, while electricity production by technology is reported by the Organization for Nordic Transmission System Operators. We used the European Brent spot price for the price of fuel oil as reported by Reuters. We regress the thermal supply on the price of electricity, the prices of fossil fuels and the time of year. A majority of the marginal cost of thermal plants consists of the price of the fuel. As explained, the thermal generation costs vary within the year for reasons related to heating demand and maintenance, both of which follow a seasonal pattern (nuclear plants and other large thermal power plants follow a seasonal maintenance schedule). To capture these effects, we include month dummies d_t in the regression equation,

$$z_t = \beta_0 + \beta_1 \ln p_t^{elec} + \delta q_t + \gamma d_t + \varepsilon_t,$$

where z_t is the thermal supply, and q_t is the vector of fuel prices. The thermal generation is composed of all other production than hydro, including wind power and the net import of electricity. The price depends on thermal generation, and is thus endogenous. There are two natural candidates for instruments, the hydro production and the level of reservoirs, both of which influence the price level but not the cost of thermoelectricity. We report our estimation results in Table 3. The first panel of the table contains the results of the first stage of the two-stage least squares regression. The first column of the table represents the model with fossil fuel (coal and oil) prices as regressors and aggregate reservoir level as the instrument for price. Fossil fuel prices are strongly multicollinear, and the price of coal is dropped from the model depicted in the second column. Finally, the third column reports the results of the same model as in the second column, but using

hydro output instead of reservoir levels as the instrument. As expected, there is a strong negative relationship between reservoir levels and price. The same holds true for total hydro output and price. The second panel of Table 3 presents the second stage results. The parameter values and the model fit are very similar for the two instruments. We take this as an indicator of the strength of the instruments since the correlation between output and reservoir levels is not perfect. Given its slightly better fit in the first stage, we use the model with reservoir levels as instruments in the calibration.

We note here that the purpose of the estimation is to find a stationary supply curve that shifts only because of the seasons within the year. This way we seek to obtain a fair description of how the hydro producers viewed their residual demand ex ante; it would not be difficult to estimate the non-hydro supply more precisely using information that is available ex post. We want to include only supply shifters that we can include into the state vector defined earlier. We set the fuel price equal to observed average from the period 2000-05, but later solve the planner's model with a stochastic fuel price using the above estimated curve. However, we cannot solve the market power model with a stochastic process for the fuel price because of the curse of dimensionality. We find no evidence that the fuel price is important for our results regarding the market structure.

Given x_t , the estimated supply z_t gives the relationship between hydro output and market prices, and this is how the value of hydro is evaluated throughout the remaining of the paper. It is therefore important to illustrate how well this key input to the model describes reality: Fig 3 depicts the historical weekly prices and the prices obtained by using historical values for z_t and the estimated thermal supply. The fit is reasonably accurate for the whole period; in particular, the estimated price equation captures the price spike of 2002-03. However, the predicted prices deviate more from the actual prices after the price spike, which may be due to the fact that thermal plants rescheduled their maintenance patterns in response to the shortage of hydro after the price spike.

The annual discount rate is 8 per cent.

We develop an algorithm for solving the model using a combination of backward induction and modified policy iteration. The algorithm begins with an initial estimate of the value of water at the end of the year. Given this end value, we can solve for the optimal policies and water values for the entire year by backward induction. Then, using modified policy iteration (see Puterman 1994), we iterate over the value of water in the first week of the year. For a given policy estimate we compute its value over a fixed number of years. The value of the evaluated policy then replaces the current estimate of the value of water in the end of the year. We iterate until the week-by-week value

function converges.

3.5 The benchmark results

We first generate the long-run weekly price moments by running the model over 2000 years, using the market fundamentals that we calibrated as explained above. Recall that we are not projecting the market to the future but, rather, studying how the model maps the distributions of the fundamentals, describing the market in 2000-05, to socially optimal price distributions. The first moments of the weekly prices are in the upper panel of Fig. 4 , and the second moments together with skewness of the prices are in the lower panel. The weekly long-run price means reveal the exhaustible-resource nature of the market: the Spring inflow is in expectations depleted over the course of the year, leading to expected prices increasing quite closely at the rate the rate of interest until next inflow peak. The drop in the price expectation from week 18 to week 19 is .063, a number close to the discount rate.¹⁷ In this sense, various constraints in the hydro system, as specified above, do not prevent a relatively close equalization of the present-value expected prices across the weeks. The average price level is 26 € which is almost identical to historical average of 26.3€ from the period 2000-05.

From the lower panel we see that the socially optimal price risk, indicated by the second moment of the weekly prices, increases towards the end of the hydrological year. This makes sense: Summer and early Fall are the periods of relatively abundant storage and predictable demand. Considerable uncertainty regarding the overall annual inflow is revealed gradually during the Fall, and unfavorable sequences of rainfall, or cold spells increasing demand, can lead to drawdown of stocks. Such risks are larger, the longer the period under consideration, which is why the socially optimal price risk must increase with time, until removed by a new inflow at the turn of the season. The skewness of price is positive and also increases towards the end of the hydrological year. This relates to the fact the storage motives across the hydrological years dominate the market dynamics exactly there: the storage demand for the next year tends to eliminate the extremely low price realizations so that there are relatively few downward price spikes to match the upward spikes (see also Deaton and Laroque 1991 for discussion).

Let us now examine a particular sequence of events, i.e., the historical realizations

¹⁷The peak price is on week 17 and the lowest price on week 20. The reduction is .085 which is slightly higher than the discount rate. Regressing the expected price on a constant and weeks, starting from week 18 and ending at the next year's week 17, gives the slope .085 for the price curve.

of demands and inflow over the period 2000-05. Figure 5 shows two panels over the weeks of 2000-2005. The upper panel is for the aggregate storage and the lower one is for hydro output, both measured as gigawatthours (GWh). The socially optimal paths are calculated by setting the initial hydro stock equal to the observed stock at the beginning of 2000 and then letting it evolve as determined by the optimal policy. Demand and inflow realizations are taken as they in actuality occurred in each week but decisions are made under genuine uncertainty regarding the future.

The planner's output matches the observed output (the lower panel) quite well. Later, after introducing the alternative market structure, we will introduce criteria for matching the model with the data. Here, we note that the seasonal first moments (quarters of the year) for the observed historical output and social planner's output deviate on average by 5 per cent, which is less than one grid step in the planner's choice set for a significant fraction of the time. The quarters are different with respect to the match such that there seems to be some tendency for the planner to save more water during the Summer and spend more in the Winter quarters. While there is no clear systematic deviation in outputs, such a deviation is clear for the reservoir levels, as illustrated by the upper panel of Fig. 5. The market and the planner have clearly differing target levels for the reservoirs. In the first two years, the planner seeks to save more of the abundant inflow (recall that we are forcing the observed and model stocks to be equal at the start), whereas later in the sample the planner would draw down the stocks more aggressively in respond to the inflow shortage taking place in late 2002. Note that the planners differing stock levels arise not because of a systematic annual difference in usage but, rather, because of relative short and intensive 'steering' of the stocks in years 2001 and 2002-03.

The implications for prices are dramatic, see Fig. 9 (the SP price). The planner can avoid the price spike of 2002-03 by more aggressive production. Excluding the price spike, the seasonal means of predicted prices are not lower, while much more stable.

4 Market power

4.1 The Model

Using the framework introduced in section 3, we now assume that a fraction of the reservoir capacity is strategically managed. We do not seek to map the observed market characteristics such as the market shares or the ownership of capacity to market outcomes but, rather, develop a stylized, while consistent, model of market power that remains

empirically implementable in this relatively complicated dynamic market. The share for the strategic capacity, $\alpha \in [0, 1]$, is our market structure parameter for which we can search values best fitting the data in Section 4.3. We assume that the fraction α is managed by one strategic agent (single firm, or an agent for a coherent group of coordinating firms). The rest of the reservoir capacity share, $1 - \alpha$, is owned and controlled by a large number of competitive agents. Note that α is the share of the capacities (reservoir and turbine), not the share of the existing hydro stock. The small agents are nonstrategic but forward looking, e.g., an individual competitive agent has no influence on the price but its decisions are rationally based on predictions for future prices, and these are formed using information that is available to all agents. This structure for oligopolistic competition remains computationally tractable, achieves the planner's solution and monopoly as limiting cases ($\alpha = 0$ and $\alpha = 1$, resp.), and, as we will show, will reveal a quite natural pattern for market power.

To separate the state vectors, inflows, and payoffs for the strategic and nonstrategic agents, we use superscripts m and c , respectively. Competitive agents are treated as a single competitive unit so that their state, for example, is

$$s_t^c = (S_t^c, x_t, \omega_t)$$

where S_t^c is the aggregate physical stock held by the competitive agents. There are thus two physical stocks that evolve according to

$$S_{t+1}^i = \min\{\bar{S}^i, S_t^i - u_t^i + r_t^i\}, \quad i = m, c, \quad (7)$$

where the reservoir capacity is what determines the size of the strategic agent: $\bar{S}^m = \alpha\bar{S}$. Both parts of the market have their own choice sets, $u_t^i \in U^i(s_t^i)$, and inflows r_t^i .¹⁸

The division of the aggregate inflow can have important implications for the exercise of power. In principle, we would like to experiment with the correlation of inflows into the stocks S_t^c and S_t^m to study its impact on the equilibrium. Unfortunately, for computational reasons, we are able include only perfectly correlated inflows: the aggregate inflow is first drawn from the weekly distribution $G_\omega(r)$, as described earlier, and then this inflow is divided into the two stocks in accordance with α .

¹⁸For the planner's model, we did not impose any formal restrictions on spilling of water as the planner has no incentives to do so, but for the large agent this incentive is material. Therefore, we want to impose a spilling constraint (implemented as a financial penalty on water spilled over in the numerical part). We have been told that the hydro plants are monitored for spilling.

We look for a subgame-perfect equilibrium in the game between the strategic and nonstrategic agents. To save on notation, we let s_t now denote $s_t = (s_t^m, s_t^c)$. At each period, the sequence of events is

1. States $s_t = (s_t^m, s_t^c)$ are observed;
2. Strategic agent chooses u_t^m ;
3. Nonstrategic agents make the aggregate choice u_t^c ;
4. Nonhydro production clears the market: $z_t = x_t - u_t^m - u_t^c$;
5. Inflow for $t + 1$ is realized.

When we impose a Markov-restriction on strategies, this timing implies that a policy rule for the strategic agent depends on both states, $u_t^m = g_t^m(s_t)$. As said, we treat the nonstrategic agents as a single competitive unit and thus look for a single policy rule for this unit, $u_t^c = g_t^c(u_t^m, s_t)$.¹⁹ It is useful to think that the competitive agents' policy seeks to solve the planner's problem of minimizing the overall social cost of meeting current and future demand requirements, given the current and future strategic behavior of the large agent. In this sense, the competitive agents minimize the cost of market power arising from the concentration of capacity in the hands of the large agent. Solving such a resource allocation problem for the competitive agents is the appropriate objective as it will generate a policy rule that implies a no-arbitrage condition for small storage holders. Thus, no small agent can achieve higher profits by rearranging its production plan from what we describe below.

Letting $v_t^m(s_t)$ denote the overall expected payoff for the strategic agent at state s_t , we see that a pair of equilibrium strategies $\{g_t^m(s_t), g_t^c(u_t^m, s_t)\}$ must solve

$$\begin{aligned} v_t^m(s_t) &= \max_{u_t^m \in U^m(s_t^m)} \{p_t u_t^m + \beta E_{s_{t+1}|s_t} v_{t+1}(s_{t+1})\}, \\ p_t &= C'_\omega(x_t - u_t^m - u_t^c) \\ u_t^c &= g_t^c(u_t^m, s_t). \end{aligned}$$

While an individual small agent takes the expected path of both stocks as given, aggregate u_t^c can be solved by minimizing the expected cost-aggregate from meeting the

¹⁹Notice that the Stackelberg timing simplifies the market clearing. Small agents' policy depends not only on the state but also on u_t^m , and so we do not have to dwell on complications caused by simultaneous moves.

demand that is not served by the large agent. Let $v_t^c(u_t^m, s_t)$ denote the value of this cost-aggregate. We define

$$\pi^c(u_t^m, u_t^c, s_t) \equiv -C_\omega(x_t - u_t^m - u_t^c)$$

as the per period payoff and note that equilibrium policy $g_t^c(u_t^m, s_t^m, s_t^c)$ solves the following recursive equation

$$v_t^c(u_t^m, s_t) = \max_{u_t^c \in U^c(s_t^c)} \{ \pi^c(u_t^m, u_t^c, s_t) + \beta E_{s_{t+1}|u_t^m, s_t} v_{t+1}^c(\tilde{u}_{t+1}^m, s_{t+1}) \},$$

where \tilde{u}_{t+1}^m is taken as given by equilibrium expectations. Having observed u_t^m , the expectation for the next period stock S_{t+1}^m is fixed by the knowledge of the inflow distribution. Similarly, for a given u_t^c , the next period competitive stock S_{t+1}^c can be estimated using the inflow distribution. Therefore, competitive agents can correctly anticipate the next period subgame (s_{t+1}^m, s_{t+1}^c) and the strategic action $u_{t+1}^m = g_t^m(s_{t+1})$. The equilibrium expectation \tilde{u}_{t+1}^m must be such that the current period action u_t^c , through the physical state equation (7) for S_{t+1}^c , fulfills this expectation:

$$\tilde{u}_{t+1}^m = E_t g_t^m(s_{t+1}).$$

In this way, competitive actions today are consistent with the next period expected subgame, without any strategic influence on the market price.

If there exists a stationary long-run equilibrium, we can drop the time index from policies and value functions. We solve the equilibrium by a long backward induction and use the first year weekly policies in the empirical application.²⁰ In this procedure, the existence of the equilibrium is not an issue.

4.2 Interpretation

We have illustrated in section 3 that the hydro market has features of an exhaustible-resource market (allocation of the Spring inflow) and a storage market (savings to the next year). In an exhaustible-resource market, market power is exercised by a sales policy that is more conservative than the socially optimal policy: sales are delayed to increase the current price²¹. In the hydro market, the seller is not free to extend the sales path in

²⁰One can in principle test if such a finite-horizon equilibrium approximates a long-run equilibrium well by simulating the long-run value functions using the finite-horizon policies, and then computing the payoffs from one-shot deviations. However, using such a test for choosing the number of needed backward-induction steps, is computationally demanding.

²¹See Hotelling (1931) for the analysis of a monopoly; Lewis and Schmalensee (1981) consider an oligopoly.

this way because of the recurrent Spring allocation which limits the length of the period over which there is scarcity of supply. In this sense, the ability to exercise market power as in exhaustible-resource models is limited. Nevertheless, the seller can shift sales to the future by storing the resource excessively to the next year, and in general such behavior is profitable because of discounting.

For illustration, suppose that all actions are made at the annual level (one period is one year), that there is no uncertainty, and that the decisions described in the previous section are made in the beginning of the year where all agents receive a deterministic annual allocation of water. It is then clear the strategic agent can reduce current supply only by saving to the next year; in equilibrium, saving takes place to the point where the current period marginal revenue equals the next period discounted marginal revenue, minus the cost from marginally reducing next year's potential for supply reduction. When the agent cannot spill water, a given stock in the hands of the strategic agent has only negative shadow price for him, as increasing the stock reduces the size of the 'sink' that is available for supply reduction. This mechanism will emerge clearly in the empirical part below.

4.3 Empirical implementation

We calibrate the market power model using the estimates for weekly inflow, demand, and thermoelectric supply, as in the model of efficient hydro use. However, we leave the strategic agent's capacity share parameter α open, and consider in next what α provides the best match with the data. We would like find to the capacity share parameter structurally, i.e., by maximizing the empirical match of the model, using the criteria discussed below, with respect to α . In principle, we follow this approach but we are limited to consider only a subset of values for α due to computational reasons. As opposed to the one-decision maker problem, the game cannot be computed using policy iteration techniques. Instead, we solve the equilibrium by straight backward induction over the weeks of 10 years. In each state, we need to solve the following fixed-point problem as part of the procedure for finding the market policy $u_t^c = g_t^c(u_t^m, s_t)$: a given u^c induces the transition of the expected stock s_{t+1}^c , which when used together with s_{t+1}^m in $\tilde{u}_{t+1}^m = E_t g_t^m(s_{t+1})$ determines the expected behavior of the large agent; in equilibrium, the assumed u^c for the state transition must be the same as the cost minimizing optimal u^c for an agent who takes the aggregate state transition as given. Since such a fixed-point may not exist on a discrete grid, we use a lexicographic criterion at each state: (i) if there

exists a unique most consistent u^c , when consistency is measured as the distance between the aggregate and private u^c , then this u^c is chosen; (ii) if criterion (i) fails, we use the Pareto criterion for choosing among the candidates. We need to apply the lexicographic procedure in approximately 5% of the states depending on the size of the strategic storage α . In total, it takes several days to solve the model on a standard desktop computer, which limits the set of parameters we can consider.

The program files for computing the model are available at the authors' webpage.

4.3.1 Simulated long-run distributions

For comparison with the social optimum, we generate the long-run weekly reservoir, price, and production moments by running the model over 2000 years using various capacity shares α . Fig. 6 depicts the long-run weekly stock levels for the social planner (SP), and for α equal to .2, .3, and .4. The expected stock levels increase monotonically with the share of the strategically managed stock. This is consistent with the interpretation given in section 4.2: the steady-state stock increase is a way to achieve the disposal of supply not meant to reach the market. Under uncertainty the logic of market power is slightly more intricate than in the deterministic case, as will be illustrated shortly, but the implication for the expected stock levels are clear.

The long-run weekly price moments are in Fig. 7, for the same parameter values. Two features can be observed. First, as expected, the price level increases with the size of the strategic agent, leading also to a more marked fall in prices at the turn of the hydrological year in the Spring. Second, for α sufficiently large, the highest expected prices are experienced earlier, before the end of the hydrological year. Our conjecture for the result is that a larger agent can follow a riskier strategy in the sense that water is withheld from the market earlier to take advantage of potential shortage of inflow during the late Summer and Fall: an inflow below expectations provides a welcome 'sink' for unused stock, so that less of the excessive saving must be carried over to the next year. On the other hand, if the inflow turns out be abundant, then the strategic agent needs to produce excessively, from his point of view, to prevent excessive storage to the next year. This latter effect tends to depress expected prices in the end of the year.

4.3.2 Matching historical data

To consider the match with the historical data, we evaluate the equilibrium policies for a given α , using the historical realizations of demands and inflows over the period 2000-05.

We set the initial hydro stock equal to the observed stock at the beginning of 2000 and then let it evolve as determined by the equilibrium policies.

We look for α that best matches the historical data. In this procedure, we use the model predictions for three variables: the reservoir levels, output, and prices. It is clearly important to include reservoir levels in the set of variables, given that imperfect competition should become evident through this variable. Recall that there is a systematic discrepancy between observed reservoir development and that chosen by the social planner (Fig. 5). Including both prices and hydro outputs in the set of variables would clearly be unnecessary if the "observed" prices were the ones computed from the estimated supply relationship using the historical outputs; in this case, there would be one-to-one relationship between outputs and prices. However, since we use the real historical prices as our observations, it makes sense to use both prices and outputs in the matching procedure to evaluate the overall performance of the model.

Let $m_t(\alpha)$ be the model prediction for a (column) vector of the three variables at t , given α . If x_t is the historical observation for the same vector, the sample mean of the prediction error is²²

$$g_T(\alpha) = \frac{1}{T} \sum_{t=1}^T (m_t(\alpha) - x_t).$$

One criterion for choosing the model is to find a value for α that minimizes the quadratic form

$$H_T(\alpha) = g_T(\alpha)' W g_T(\alpha),$$

where W is a 3×3 weighting matrix (to be discussed below). A crude way to proceed is to choose $T = 312$, i.e., to aggregate over all weeks of the six-year period to form three simple moment restrictions. When $W = I$, the statistic has a straightforward interpretation: it is the sum of three least-square errors. This statistic is misleading since it completely ignores the Markovian nature of the policy rule: the statistic should be able discriminate how well the model predicts variables as the state of the market changes. Another extreme is to let $T = 1$, which allows one to calculate the statistics $H_1(\alpha)$ for each of the 312 weeks, and then sum up these numbers (or average them). This approach would pay maximum attention to actions at individual states, but would not allow weighting the variances of the prediction errors when choosing W in $H_T(\alpha)$. The latter shortcoming can be avoided, for example, when $T = 13$ and the statistic $H_{13}(\alpha)$ is calculated separately for each of the 24 quarters in the data. Then, we can use the

²²We are abusing notation on purpose here, hopefully without a risk of confusion, in order to follow the conventions of the literature using the GMM approach.

two-stage GMM approach²³ where in the first stage α is chosen for some given W , and in the second stage, we estimate the sample variance-covariance matrix of the prediction errors associated with the chosen α to construct a weighting matrix that depends on the data.²⁴

We evaluate each model under different criteria ranging from moment restrictions for aggregated data to "path matching" using weekly data. For Table 4, we have first calculated the statistic $H_T(\alpha)$ for each model at different aggregation levels (weekly, monthly, quarterly, and semi-annual). In this calculation, we took W first as a given diagonal matrix and used the inverses of squared means of the relevant variables on the diagonal to transform the variables into comparable units.²⁵ The mean value of the statistic $H_T(\alpha)$, obtained this way, is reported in the first column for each model. The 35 per cent model provides the best score at all time aggregation levels.²⁶

For quarterly and semi-annual predictions there is enough variation to consider the variance of the sample mean and to exploit the covariance-variance properties of the data in choosing the weighting matrix for the statistics. The reported numbers are the mean values of the statistic over the 24 and 12 samples (quarterly and semi-annual aggregation, resp.). The 30 per cent model minimizes the statistic $H_T(\alpha)$ obtained this way. Note that $H_T(\alpha)$ from the 30 per cent model need not be the smallest, for example, in each of the 24 quarters, but the only the mean value of the statistic has this property. We are thus putting equal weights to the match in each of the time periods.

Our main result is that a market share of 30 per cent for the strategic agent provides the best fit with the historical data under various criteria. In Table 5, we report statistics on the entire observed and predicted price series. The average price in the sample period was 26.3 euros. The socially optimal hydro policy would have yielded a mean price of € 24.9. The 30% model outperforms the planner's model in predicting the average, variance

²³See, for example, Cochrane (2001).

²⁴In the second stage, we allow for serial correlation in the prediction errors associated with the chosen alpha by using the inverse of the estimated long-run variance matrix as the weighting matrix. The asymptotic variance matrix is computed using the quadratic spectral kernel proposed by Andrews (1991) and a bandwidth of three. The results are robust to different kernel types and to a large range of bandwidths.

²⁵Otherwise, the stock variable dominates in the calculation. Correcting dimensions this way favors the hypothesis that there is no market power since the market power model is particularly good in matching the stock development.

²⁶Due to computational reasons we have computed only seven market share values for the strategic agent: 0, 20, 25, 30, 35, 40, and 50 per cent. Since we find no evidence for perfect competition, i.e., $\alpha = 0$, we do not believe that this coarse grid for α is essential for the main result.

and skewness of price. It also outperforms the other market structures in the Table, with the exception of slightly underestimating the skewness of price compared to the 40% model.

Recall that for computational reasons we did not cover a very large set of α -values, which is why a better fitting market share parameter is likely to exist. However, we do not see a large gain from this search as α has no clearly defined empirical counterpart. The objective of the analysis is to merely show that there exists some market structure with market power that has more predicting power than the socially optimal structure. While it is clear that having one more parameter to choose, cannot hurt us ($\alpha = 0$ is always a choice), it is somewhat surprising that the model prediction is better in all dimensions (price, output, stocks). In Fig. 9, we depict again the observed price, this time together with the predicted price under $\alpha = .3$ and the planner's solution. The market power model can replicate the price shock of 2002-03 quite well (the price shocks in 2003-04 originate our supply curve estimation which does not capture well the change in the available capacity of thermal; see Fig. 3). In Fig. 8, we see the systematic improvement in the reservoir match throughout the period 2000-2005.

5 Robustness analysis

In this section, we study the possibility that unobserved factors, mismeasured data and expectations, or limitations in the model structure can lie behind the pattern that we have connected to imperfect competition.

5.1 Unobserved reservoir capacity constraints

Reservoir constraints can have substantial implications for the main behavioral patterns in the market. In Figure 8, we see that the first-best reservoir levels overshoot the observed levels in years 2000-02 and then, in the latter part of the period, the deviation is to the opposite side. It seems clear that by sufficiently reducing the maximum reservoir capacity, we may obtain a better match in the years of overshooting, while a sufficient increase in the minimum capacity may improve the match for the remaining years.

We took the reservoir constraints from the data (see fn. 3.4) but now we look for constraints that maximize the model fit under the competitive behavioral assumption. We thus compute the social planner's model for all minimum reservoir levels (in TWh)

$$\underline{S} \in \{0, 1, 2, \dots, 20\},$$

and all maximum reservoir constraints

$$\bar{S} \in \{112, 113, \dots, 120\}.$$

These parameter sets were chosen based on historical reservoir levels, so the search was conducted over a range that should cover the "true" limits. After solving for the policy rules corresponding to the alternative parameterizations, we applied the policies at historically observed states, and then computed the two-stage GMM statistic for each model. The model with the lowest test score provides the best fit with the historical data.

Using this procedure we find that the best-fitting pair is $\underline{S} = 17$ TWh and $\bar{S} = 112$ TWh. These choices lead to almost identical reservoir development with that predicted by our model of imperfect competition, and since the reservoir is important for the test statistic, the model fits are indistinguishable. However, the constraint adjustments cannot explain the observed price increase. Are estimated capacity constraints consistent with data? The estimated lower limit of 17 TWh is implausibly high given the discussion in footnote 3.4. As such, the maximum capacity of 112 TWh is also off by being too low; higher actual levels have been observed since the deregulation of the Norwegian power market.²⁷ It should also be noted that the capacity used in the model may proxy limitations in the hydro system arising from regional heterogeneity, and that therefore, there may not be a single number that would be the appropriate estimate of the maximum capacity for the whole sample period.²⁸

²⁷The aggregate maximum reservoir capacity in the Nordic market was almost constant throughout the sample period, being 120.5 TWh in the beginning of 2000 and 121.0 TWh in the end of 2005 (Nordel annual statistics 2001 and 2006). In 1990-2007, the maximum observed aggregate reservoir level in the market was 115 TWh (94.5%) (Nord Pool). In Norway, reservoirs have reached a high of 97.3% (1990-2007) and in Sweden 97.7% (in 1950-2007).

²⁸Such a constraint can artificially represent the unmodelled limitations that regional heterogeneity puts on the storage behavior. For example, in the summer of 2007, the reservoir levels in Southern Norway were close to the capacity, and the local producers had to generate so much power to avoid overflow that they were unable to export all the power to other parts of the system, and the weekly average area price dropped to just 3.77 €/MWh in week 34, when the system price was 16.2 €/MWh. In general, once the transmission line from a hydro abundant region becomes congested, increasing output in that region does not affect the prices faced in the other areas. Thus, the hydro producers in the other parts of the system have no incentive to reduce their output and save more water even though hydro

The estimated 112 TWh coincides with the actual maximum level in 2000, and thus forces the reservoir path to the true level. The reduction in the initial storage level also means that the planner is carrying less water in the end of 2002, and has less hydro resources to allocate to the price spike in the winter 2002-03. In the latter part of the sample, the maximum reservoir capacity has only a marginal effect on the optimal policy until in 2005, when reservoirs again approach the maximum capacity. On the other hand, the reservoir lower limit has a very small effect to the results in the first three years of the sample period. It thus follows that one needs to adjust both the upper and lower limits for capacity to challenge the market power explanation. We find this implausible.

We also experimented with a minimum flow constraint. We considered several levels for the lower bound of hydro output, $\underline{u} \in \{0, .2, .4, \dots, 2.8\}$ (in TWh). Low levels of the minimum flow constraint have no effect on the optimal hydro policy. For high enough levels, the model fit slightly improves. The fact that the planner must be able to meet the constraint on hydro output in future periods means that the planner must have enough water in storage to meet these future obligations. In the historical simulation, this effect can be seen as a gradual build-up of storage levels throughout the sample. This more conservative hydro use policy also implies slightly higher prices during the price crisis of 2002-03. Nevertheless, the influence of the minimum flow constraint is of secondary importance when compared to the reservoir level constraints.

5.2 Fuel price uncertainty

We took the oil price, which was the only statistically significant fuel price in the nonhydro supply, as an average price from 2000-05. Due to the curse of dimensionality, we could not solve the model of strategic hydro use with stochastic price, but we can solve the planner's model under this assumption.²⁹ We can therefore evaluate whether the fuel price changes can explain the discrepancy between the first-best and observed behavior. To this end, we assume a Markov process for the price. The price belongs to a finite set,

output in the congested region is very high. Economically, transferring water from the congested area to the other regions would improve welfare. In our model, where all reservoirs are aggregated into a single storage, such uneven distribution of inflow has no similar consequences.

²⁹The inclusion of oil price in the state increases computation time approximately linearly in the number oil price states. This is due to the fact that the most time-consuming part of the algorithm is taking the expected value over the reservoir state transitions. These transitions depend on the stochastic inflow process and the current estimate of the planner's hydro use policy. Since the policy is dependent on the current oil price, the expectations must be computed for each possible current oil price state; hence the linear increase in computation time.

roughly consistent with the empirical support from 2000-05. To be more specific,

$$\begin{aligned} p_t^{oil} &\in \{10, 12, 14, \dots, 80\}. \\ p_t^{oil} - p_{t-1}^{oil} &\in \{-6, -4, -2, \dots, 6\} \end{aligned}$$

The transitions are assumed to follow a normal distribution, the mean (.08) and standard deviation (1.46) of which are estimated from actual weekly oil price changes in 2000-05.

The fit of the planner's model's predicted price path improves with the inclusion of the oil price state, but not the fit of the reservoir levels. The price effect is most pronounced in 2004-05, when the price of Brent roughly doubled from its level at the end of 2003. While the price prediction becomes generally more accurate, it does not replicate the observed price spike of 2002-03. Indeed, the predicted prices in 2002-03 are lower in the new model than in the benchmark planner's model. Uncertainty over future input prices increases the planner's incentive to store more water in the relatively water abundant years 2000-01. This increased storage is then used during 2002-03 not to alleviate price pressure due to high input prices, but due to the scarcity of water.

5.3 Discounting

We have used a discount factor that corresponds to an 8 per cent annual discount rate. Holding wealth in hydro stocks is relatively risky, justifying a rate above the risk-free rate, although we are unaware of prior studies elaborating what discount rates should be applied in this context.

To test which interest rate is supported by the historical data, we evaluated the alternative planner's models in the same way we compared the different α -values. Using the historical demand and inflow realizations, we simulated the price, reservoir and hydro output paths for all discount rates (percentages) in the range $\{2, 4, \dots, 20\}$. We then computed the GMM test statistic for all models, using quarterly averages as observations. The test score is lowest for the model with 12 per cent discounting.

As one would expect, increasing the interest rate decreases the expected level of reservoirs. Earlier, we have shown that in the strategic model, raising the market share of the large agent increases the expected reservoir level. This effect is much stronger than the one from decreasing the interest rate in the planner's model. For example, even at a 2 per cent interest rate the planner's expected reservoir level is lower than in the strategic model with $\alpha = .3$ and the interest rate at 7 per cent. Since expected price increases at

the rate of interest through the hydrological year, a higher discount rate implies higher output in the Spring and Summer periods and lower output in late Fall and Winter. A higher interest rate also increases the weekly standard deviation of price in virtually all weeks of the year, the exception being the weeks immediately following the start of the spring inflow, when variation is lowest and differences between discount rates are very small. Price skewness, on the other hand, is more variable when interest rates are low. In particular, prices are more positively skewed before the spring inflow and less skewed in the summer for low discount rates.

In the historical simulation, higher discounting lowers the reservoir level in every week of the sample period. Yet, even a very high discount rate does not explain the low storage levels in 2000. During the winter of 2002-03 the low reservoir levels due to higher discounting force the planner to use less water, thus causing a more pronounced price spike than in the benchmark model. Prices at the peak are, however, probably depressed by the high discount rate. That is, the planner is more willing to take losses in the future than now, and will therefore use water more aggressively. After the price crisis, higher discounting leads to slower build-up of the reservoirs.

5.4 Expectations

We also considered the possibility that our assumptions about the parameters of the demand and inflow distributions might be off the mark.

Given the low levels of storage in the early part of the sample, one possible source of bias could be too high expectations of future inflows. If expected inflow in the benchmark model is underestimated, then the high realizations in 2000 are seen as more valuable, and storage is higher. If, in reality, the expectations were higher than in the model, this could create a shortage of water in 2002-03, which could replicate the price spike. On the other hand, higher expected inflow should also lead to more aggressive use of water during a stock-out, as the producers believe that their storages will be soon replenished. This should then bring down the price spike in the simulation results. To test for these hypotheses, we increased the expected inflow mean by 5 per cent, leaving the variance of the inflow distribution unaltered. As expected, the change in expectations induces the planner to use more water in the early part of the sample than before, but the pattern is quite different from the actual hydro output. More specifically, compared with historical output, the new simulation results still overestimate the level of reservoirs in the water abundant year 2000, but a steadily decreasing reservoir level thereafter, so that storage

is significantly lower than in reality in the summer of 2002 before the price crisis. The shortage of water then causes the prices to peak at a higher level than before, but the price spike is not as pronounced as in the fringe models, for example. After the shortage, reservoir levels are built up too slowly compared to the actual pace. Overall, it seems that changing the inflow expectations in the described manner will bring about only a modest improvement in the model fit, at best.

The demand support used in the benchmark model was based on the empirical support of demand in 2000-05. To be specific, the week-specific lower bound of the demand space was set at the grid point below the observed minimum demand in that week, and the upper bound similarly at the grid point just above the observed maximum. We analyzed the sensitivity of the simulation results to the choice of demand support by considering mean-preserving spreads of demand uncertainty. We first decreased the week-specific lower bounds and increased the upper bounds by two standard deviations each. The probabilities formerly assigned to the lowest and highest demand levels were spread to cover the new support according to the original distributional assumption. That is, the mean and variance of the demand distributions were not changed.

Expanding the demand support has no effect on the historical simulation paths. We also experimented by altering the demand space in the high-season (weeks 45-10) only. This, had no effect on the simulation results, either. Adjusting the demand space by four standard deviations has a small but almost indiscernible effect on the results. This change is virtually the same whether the demand supports are expanded for all weeks of the year, or for the high-season only.

5.5 Thermal capacity and price cap

In the current model, the thermal supply curve is assumed to represent all power sources other than hydro. Based on information in the Nordel annual statistics, the aggregate capacity from all non-hydro sources including imports was approximately 8 TWh per week in the sample period. The highest observed output from these sources during the same time was 5.5 TWh. The model has no explicit constraint on thermal capacity. This does not, however, mean that we assume an infinite supply of "thermal" power. Instead, one may interpret the supply exceeding the thermal capacity as stemming from elastic demand. After all, the hydro producers are interested only in their residual demand. An assumption about demand elasticity during an extreme power shortage is always going to be ad hoc, since we have not observed such a situation in practice.

Nevertheless, we experimented by constraining thermal capacity to be less than 5.8 GWh a week - a rather stringent condition given the theoretical maximum capacity. The capacity constraint must be paired with either elastic demand or with a penalty for lost load. The value of lost load (VOLL) has been estimated to be 2000 €/MWh in the Nordic market. We used this figure as the price cap. Thus, up until thermal capacity, supply is determined by the estimated thermal supply curve as before, but at and beyond 5.8 GWh supply is flat. In the planner's case, this means that the planner incurs a cost of 2000 € per each MWh of load that it can not supply.

The results are practically identical to the case, where supply is unconstrained. We have not surveyed yet, what effect the price cap would have in the strategic model.

6 Conclusions

We have developed a structure that can be used to interpret market data in an attempt to make a distinction between behavioral patterns arising from imperfect competition and those arising from fundamental factors in the institutional and economic environment. We found evidence supporting the conclusion that imperfect competition can explain the main behavioral patterns in the market outcome in years 2000-05. The data period includes an extraordinary period allowing us to identify the pattern for market power. But in expected terms the welfare losses are extremely small. The hydro resource allocation in the Nordic power market follows surprisingly closely the first-best outcome. Our framework can be used to test if some unobserved factors or mismeasured data could lie behind the results. We found no such evidence. However, there is a number of factors whose effect on the market outcome we cannot evaluate using this approach. We conclude by discussing such factors.

We evaluated the efficiency of the long-run outcomes, which is a natural starting point for the analysis because there is no basis for evaluating the short-run outcomes without the knowledge of longer-term benchmarks. However, the hydro capacity is very different from thermal and other capacity forms also in the short-run. It allows rapid adjustments of usage, thereby potentially exploiting constraints in the transmission system or those that the other production forms face (see Hoel 2004). If the hydro capacity commands an extra, potentially state-dependent, short-run return because of its special nature, the long-run allocations are also altered. Hopefully the data allows an evaluation of potential short-run inefficiencies in the near future.

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, and then using the estimated policies to simulate hydro resource values. These values could in principle be used in estimation of structural parameters of the market (Bajari, Benkard and Levin (2007)). Such an indirect approach, rather than our direct approach, is perhaps more natural when data is not constraining the choice of the approach.

Finally, our benchmark for efficiency analysis was obtained using a risk-neutral decision maker. Behavior under risk neutrality and various constraints in the environment can show resemblance to behavior arising from pure risk aversion. There are reasons to believe in risk aversion in the hydro resource use. Large players may want to avoid extreme outcomes (e.g., stockouts) to avoid creating political pressure on the market institution. Alternatively, the regulatory constraints are likely to be 'soft' in this market in the sense that capacity usage is affected by regulatory communications in states where the electricity system is under stress. One may want to consider if risk-aversion changing pricing rule for the state-dependent resource can explain some of the deviations we discovered.

7 Appendix: proof of Proposition 1

Proof. We can take Δ as the smallest deviation allowed by the action space such that $d(s_t) \in U(s_t)$. The properties of optimal prices follow from non-optimality of one-shot deviations described by $d(s_t)$. By the optimality of $g(s_t)$,

$$\pi(g(s_t)) + \beta E_t v(g(s_t)) \geq \pi(d(s_t)) + \beta E_t v(d(s_t)) \quad (8)$$

$$\iff$$

$$\pi(g(s_t)) - \pi(d(s_t)) \geq \beta E_t v(d(s_t)) - \beta E_t v(g(s_t)). \quad (9)$$

Recall that

$$\pi(g(s_t)) - \pi(d(s_t)) = -C_\omega(x_t - g(s_t)) + C_\omega(x_t - d(s_t)).$$

As in text, we can define $\bar{p}_t = \bar{p}(s_t, \Delta)$ such that

$$-\bar{p}(s_t, \Delta)\Delta = \pi(g(s_t)) - \pi(d(s_t)), \quad (10)$$

Note then that

$$\beta E_t v(d(s_t)) - \beta E_t v(g(s_t)) = E_t \sum_{\tau=t+1}^{\infty} \beta^{\tau-t} \{\pi(d(s_\tau)) - \pi(g(s_\tau))\} \quad (11)$$

$$= E_t \sum_{\tau=t+1}^{\infty} \beta^{\tau-t} \bar{p}(s_\tau, \varepsilon_\tau) \varepsilon_\tau \quad (12)$$

where changes in the optimal usage path, after the one-shot deviation from the optimal policy at t , are denoted by ε_τ .

Combining (9), (10) and (12) implies that one-shot deviations satisfy

$$-\bar{p}(s_t, \Delta) \Delta \geq E_t \sum_{\tau=t+1}^{\infty} \beta^{\tau-t} \bar{p}(s_\tau, \varepsilon_\tau) \varepsilon_\tau.$$

But when Δ is the smallest deviation allowed by the grid for actions (same for all periods), the condition implies

$$-\bar{p}(s_t, \Delta) \Delta \geq E_t \beta^k \bar{p}(s_{t+k}, -\Delta) (-\Delta) \text{ for some } k \geq 1. \quad (13)$$

Now, if $g(s_t)$ is constrained from above (i.e., there is no $d_t > g(s_t)$ such that $d_t \in U(s_t)$), then only $\Delta < 0$ is feasible, and, by (13), we have

$$\Delta < 0 \Leftrightarrow \bar{p}(s_t, \Delta) \geq E_t \beta^k \bar{p}(s_{t+k}, -\Delta) \text{ for some } k \geq 1. \quad (14)$$

On the other hand, if $g(s_t)$ is constrained from below, then only $\Delta > 0$ is feasible, and we have

$$\Delta > 0 \Leftrightarrow \bar{p}(s_t, \Delta) \leq E_t \beta^{k'} \bar{p}(s_{t+k'}, -\Delta) \text{ for some } k' \geq 1. \quad (15)$$

Finally, if the optimal policy is not constrained, then both (14) and (15) must hold at s_t . ■

References

- [1] Amundsen, E. and Bergman, L. (2002), "Will Cross-ownership Re-establish Market Power in the Nordic Power Market?", *The Energy Journal*, 23, 73-95.
- [2] Amundsen, E. and Bergman, L. (2006), "Why Has the Nordic Electricity Market Worked So Well?", *Utilities Policy*, 14, 148-157.
- [3] Bajari, P., Benkard, C., Levin, J. (2007), "Estimating Dynamic Models of Imperfect Competition", *Econometrica*, 75, 1331-1370.
- [4] Benkard, C. (2004), "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft", *Review of Economic Studies*, 71, 581-611.

- [5] Borenstein, S., Bushnell, J. and Wolak, F. (2002), "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market", *American Economic Review*, 92, 1376-1405.
- [6] Bushnell, J. (2003), "A Mixed Complementarity Model of Hydro-Thermal Electricity Competition in the Western U.S.", *Operations Research*, 51, 81-93.
- [7] Bushnell, J., Mansur, E., and Saravia, C. (2008), "Vertical Arrangements, Market Structure and Competition: An analysis of Restructured U.S. Electricity Markets", *American Economic Review*, 98, 237-266.
- [8] Bye, T., Bruvoll A. and Aune, F. (2006), "The Importance of Volatility in Inflow in a Deregulated Hydro-Dominated Power Market", (Discussion Papers No. 472, Statistics Norway, Research Department).
- [9] Cochrane, J. (2001) *Asset Pricing* (Princeton, NJ: Princeton University Press).
- [10] Crampes, C. and Moreaux, C. (2001), "Water Resource and Power Generation", *International Journal of Industrial Organization*, 19, 975-97.
- [11] Deaton, A., and Laroque, G. (1992), "On the Behaviour of Commodity Prices", *The Review of Economic Studies*, 59, 1-23.
- [12] Deaton, A., and Laroque, G. (1996), "Competitive Storage and Commodity Price Dynamics", *The Journal of Political Economy*, 104, 896-923.
- [13] Doraszelski, U. and Pakes, A. (2007), "A Framework for Applied Dynamic Analysis in IO", in M. Armstrong and R. Porter (eds.), *Handbook of Industrial Organization*, Vol. III (Amsterdam: North-Holland) 1887-1996.
- [14] von der Fehr, N-H., Amundsen, E. and Bergman, L. (2005), "The Nordic Market Design: Signs of Stress?", *The Energy Journal*, European Energy Liberalisation (Special Issue), 26, 71-98.
- [15] Garcia, A., Reitzes, J., and Stacchetti, E. (2001), "Strategic Pricing when Electricity is Storable", *Journal of Regulatory Economics*, 20, 223-47.
- [16] Genc, T. and Thille, H. (2008), "Dynamic Competition in Electricity Markets: Hydropower and Thermal Generation", (Mimeo, University of Guelph).

- [17] Hoel, M. (2004), "Electricity Prices in a mixed thermal and hydropower system", (Mimeo, University of Oslo).
- [18] Hortaçsu, A. and Puller, S. (2008), "Understanding Strategic Models of Bidding in Deregulated Electricity Markets: A Case Study of ERCOT", forthcoming *RAND Journal of Economics*.
- [19] Lewis, T. and Schmalensee R. (1980), "On Oligopolistic Markets for Nonrenewable Resources", *Quarterly Journal of Economics*, 95, 475-491.
- [20] McLaren, J. (1999), "Speculation on Primary Commodities: The Effects of Restricted Entry", *The Review of Economic Studies*, 66, 853-871.
- [21] Newbery, D. (1984), "Commodity Price Stabilization in Imperfect or Cartelized Markets", *Econometrica*, 52, 563-578.
- [22] Pesendorfer, M., and Schmidt-Dengler, P (2008), "Asymptotic Least Squares Estimators for Dynamic Games", *Review of Economic Studies* (forthcoming).
- [23] Puterman, M. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (New York, NY: John Wiley & Sons, Inc).
- [24] Reiss, P. and Wolak, F. (2007), "Structural Econometric Modeling: Rationales and Examples from Industrial Organization", in J. Heckman and E. Leamer (eds.), *Handbook of Econometrics*, Vol VI, (forthcoming).
- [25] Rotemberg, J. and Saloner, G. (1989), "The Cyclical Behavior of Strategic Inventories", *The Quarterly Journal of Economics*, 104, 73-97.
- [26] Ryan, S. (2006), "The Costs of Environmental Regulation in a Concentrated Industry" (Mimeo, Massachusetts Institute of Technology).
- [27] Scott, T. and Read, E. (1996), "Modelling Hydro Reservoir Operation in a Deregulated Electricity Market", *International Transactions in Operational Research*, 3, 243-253.
- [28] Puller, S. (2007), "Pricing and Firm Conduct in California's Deregulated Electricity Market", *Review of Economics and Statistics*, 89, 75-87.
- [29] Stokey, N., Lucas, R. and Prescott, E. (1989) *Recursive Methods in Economic Dynamics* (Cambridge, MA: Harvard University Press)

- [30] Williams, J. and Wright, B. (1991) *Storage and Commodity Markets* (Cambridge: Cambridge University Press).
- [31] Wolfram, C. (1999), "Measuring Duopoly Power in the British Electricity Spot Market", *American Economic Review*, 89, 805–26.




Figure 1: Inflow energy in the Nordic market area in 1980-99. Sources: Norwegian Water Resources and Energy Directorate (www.nve.no), Swedenergy (www.svenskenergi.se) and Finland's environmental administration (www.ymparisto.fi).




Figure 2: Mean and empirical support of demand in Nordic market 2000-05.




Figure 3: Observed (solid line) and estimated (dashed line) system price 2000-05. Estimation based on historical output levels.




Figure 4: Simulated expected price (upper panel) and the skewness and standard deviation (lower panel) of price.




Figure 5: Upper panel: observed (solid line) and social planner's (dashed) reservoir levels. Lower panel: observed (solid line) and social planner's (dashed) hydro output




Figure 6: Simulated expected reservoir levels for different market structures.




Figure 7: Simulated weekly price expectations under different market structures.




Figure 8: Historical, the planner's, and market power (30%) storage levels.




Figure 9: Historical, the socially optimal, and the market power (30%) price.

Quarter	Sweden	Finland	E-Denmark	W-Denmark	Norway 1	Norway 2
Q1	2.0	2.6	8.2	5.2	1.5	1.7
Q2	7.5	8.1	21.1	6.8	4.0	2.7
Q3	6.2	12.9	24.6	6.5	2.8	4.8
Q4	2.5	4.3	14.9	10.8	1.4	2.1
All	4.6	7.0	17.2	7.5	2.5	2.8

Table 1: Average weekly area price deviations from the system price 2000-05 (Source: Nord Pool)

	Denmark	Finland	Norway	Sweden
Total generation	37.3	73.4	125.2	146.5
Hydro power	0.0	12.7	124.1	67.8
Other renewable power	5.8	2.0	0.3	1.9
Thermal power	31.5	58.8	0.8	76.7
- <i>nuclear power</i>	<i>0.0</i>	<i>21.8</i>	<i>0.0</i>	<i>66.6</i>
- <i>CHP, district heating and condensing power</i>	<i>29.4</i>	<i>26.3</i>	<i>0.1</i>	<i>5.8</i>
- <i>CHP, industry</i>	<i>2.1</i>	<i>10.7</i>	<i>0.4</i>	<i>4.3</i>
- <i>gas turbines, etc.</i>	<i>0.0</i>	<i>0.0</i>	<i>0.3</i>	<i>0.0</i>

Table 2: Average production levels (TWh) by technology in the Nordic market 2000-05

Panel A: First stage results (dependent variable log of system price)

	(1)	(2)	(3)
Oil price	0.0199**	0.0197**	0.0205**
	(0.0017)	(0.0016)	(0.0018)
Coal price	-0.0019		
	(0.0014)		
Reservoir level	-0.0280**	-0.0290**	
	(0.0015)	(0.0015)	
Hydro output			-0.0006**
			(0.00004)
Observations	300	313	313
R-squared	0.71	0.70	0.59

Panel B: Second stage results (dependent variable total thermal output in GWh)

	(1)	(2)	(3)
ln(price)	1200.9**	1185.4**	1254.1**
	(43.3)	(43.2)	(47.9)
Oil price	-27.3**	-24.0**	-24.5**
	(1.8)	(1.7)	(1.8)
Coal price	7.1**		
	(1.4)		
Observations	300	313	313

Table 3: Results of the 2SLS thermal supply estimation. The standard errors (in parentheses) have been corrected for heteroskedasticity and autocorrelation. The regression also includes monthly dummy variables. Statistical significance is marked with (**) at the 1% level and (*) at the 5% level.

Weeks	SP		20%		25%		30%		35%		40%		50%	
1	1.21	-	0.82	-	0.68	-	0.55	-	0.35	-	0.66	-	0.91	-
4	1.20	-	0.80	-	0.66	-	0.53	-	0.34	-	0.64	-	0.89	-
13	1.14	12.8	0.75	10.5	0.61	9.0	0.48	7.7	0.27	17.6	0.57	25.6	0.78	42.9
26	1.06	73.6	0.67	61.4	0.53	48.7	0.40	47.8	0.21	85.7	0.47	144.2	0.56	314.3

Table 4: Goodness-of-fit tests. The first column for each model reports the H-statistic (divided by 10^5) from the first stage of the estimation. The second column reports the H-statistic from the second stage (for applicable models). The second stage weighting matrix is heteroskedasticity and autocorrelation consistent.

	Observed	SP	20%	30%	40%	50%
Mean price (€/MWh)	26.3	24.9	25.2	26.4	28.0	31.0
Standard deviation	11.9	7.5	8.3	10.6	16.6	28.7
Skewness	2.5	0.9	0.9	1.4	2.3	5.4
Total cost (bn.€)	9.3	8.7	8.8	9.2	9.8	10.9
Welfare loss (bn.€)	0.64	0	0.14	0.57	1.16	2.26

Table 5: Price and cost statistics for the historical series and model predictions. The estimates of total cost are based on the estimated thermal supply.