

Market power and storage: Evidence from hydro use in the Nordic power market

Olli Kauppi and Matti Liski*

June 2008

PLEASE DO NOT CITE THE NUMBERS WITHOUT PERMISSION

Abstract

The Nordic power market presents a unique opportunity to test market power in storage behavior due to preciseness of data on market fundamentals determining hydro resource use. We develop and calibrate an aggregative hydro storage model. We find that historical market experience in 2000-2005 implies a 7.3 per cent welfare loss, or that the cost of meeting the same demand could have been 636 mill. € lower. The data suggests a behavioral pattern that we can match with a model of market power. A market structure where 30 per cent of the storage capacity is strategically managed outperforms the competitive model. Market power increases expected reservoir and price levels, and also implies a considerable increase in the price risk.

JEL Classification: L1, L9

Keywords: storage, hydroelectricity, market power

*Kauppi <kauppi@hse.fi> and Liski <liski@hse.fi> are at the economics department of the Helsinki School of Economics. Liski is faculty associate at MIT-CEEPR. Both authors are grateful for funding from the Nordic Energy Research (NEMIEC and NEECI). We are grateful for comments and very constructive help from Karl-Axel Edin, Reyer Gerlagh, Mitri Kitti, Maria Kopsakangas-Savolainen, Pauli Murto, Thomas Tangerås, Otto Toivanen, and seminar participants at Catholic University of Chile, Catholic University of Leuven, HECER-Helsinki, Research Institute of Industrial Economics (Stockholm), and University of Manchester.

1 Introduction

Market power in a storable-good market is notoriously difficult to detect because price-cost margins depend on expected future market conditions that cannot be observed ex post. There exists a well-developed theory on competitive storage that can explain some stylized patterns of price series (work by Williams and Wright is summarized in their book 1991; see also Deaton and Laroque, 1992, 1993). However, there is little work on market structure and storage and, in particular, empirical applications or tests are practically nonexistent.¹ This paper is concerned with a unique opportunity to use an electricity market as a natural laboratory for testing the nature and degree of market power in a storage market for water. The opportunity is unique since it is hard to think of other markets where the level of storage, prices, demand, production, and technical information on market fundamentals are reported with similar preciseness. We study the storage of hydroelectricity in the Nordic power market where about half of annual consumption is met by hydro. In this market, there are several hundred hydro power stations connected to a relatively tightly integrated Nordic market area covering Finland, Denmark, Norway, and Sweden. The market is operated through a common pool determining day-ahead hourly prices.

Using data on historical inflows, demands, and thermoelectric supply, we estimate how hydro producers should view these market fundamentals. We then develop and calibrate an aggregative hydro storage model. We find that historical market experience in 2000-2005 implies a 7.3 per cent welfare loss, or that cost of meeting the same demand could have been 636 mill. € lower. We estimate structurally various unobserved constraints on the hydro system and do not find evidence that such constraints could explain the deviation. However, the data suggests a behavioral pattern that we can match with a model of market power. A market structure where 30 per cent of the storage capacity is strategically managed outperforms the competitive model.

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

¹Mc Laren (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. Rothenberg and Saloner include storage as a strategic device supporting collusive oligopoly equilibria. There is also an extensive literature on exhaustible-resource oligopoly where the resource stock can be interpreted as a storage (Lewis and Schmalensee 1980, Polansky 1986, Loury 1992, and Salo and Tahvonen, 2001). However, stochastic 'harvest' and the possibility of a market level stockout are material for understanding commodity price behavior. These features are absent in exhaustible-resource literature.

that seeks to use this endowment to equalize expected prices until the next Spring. The market has also features of a traditional storage market: favorable demand-inflow realizations lead to storage demand and savings to the next year. Using the model to map the distributions of the market fundamentals to price distributions, we can find support for these interpretations. 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. Also, towards the end of the hydrological year weekly price distributions have moment properties familiar to those observed in other storable-commodity markets.

We find that market power seeks to shift available supply to the future, thereby increasing the expected reservoir levels as well as prices. The distortion in the historical reservoir development that we gauge by evaluating socially optimal policies along the historical sequence of events, can be unambiguously matched with the market power model. The market power model seems to produce a better match with the data with respect to all key variables including hydro output, prices, and reservoirs, when first moments are considered. Market power also implies an increase in the price risk.

We find that the expected cost of market power that we obtain by using the market fundamentals estimated from the data, 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 according to our calculations can occur 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 as 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. The final section concludes and discusses the shortcomings of the approach.

2 Institutions and market fundamentals

2.1 System price

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.² Market participants submit quantity-price schedules to the day-ahead hourly market (Elspot market).³ 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 departures in weekly averages. About 94% of the hydro resource stocks are located in the Norwegian and Swedish price areas, which are the least problematic of all price areas in the Table. It would be difficult to choose any other price

²For more information, see www.nordpool.com.

³The day-ahead Elspot market is the relevant spot market. While there is an after market (Elbas market) closing an hour before delivery, volumes in the Elbas market are small relative to the Elspot. In some other markets, like the California PX, the day-ahead market is called the spot market.

than the system price as the reference price for hydro storage decisions.⁴

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 hydro-electricity was generated in Norway and 33 per cent in Sweden.⁵ 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 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% of annual inflow arrives in the three months following week 18. The aggregate reservoir capacity in the market is 120 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

⁴The 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.

⁵The capacities cited here are reported by the Organisation for the Nordic Transmission System Operators (www.nordel.org) unless otherwise noted.

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 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%, 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 potential market power, it is important to understand these two interpretations. 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 using a nested fixed-point algorithm. 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 also 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-, and oil-fired plants), industrial 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. However, we are uncertain on the relevance of this phenomenon in the Nordic area.

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 $C_\omega(z)$ includes the level of fuel prices and we could also incorporate changing fuel prices explicitly. However, while an important source of uncertainty, fuel prices are not structural variables of the Nordic market in the same sense as inflow and demand are; we cannot estimate fuel price distributions with the same accuracy. Indeed, it is 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.

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. 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 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. The market 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 market. 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:

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

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

$$\mu_\omega > \beta\mu_{\omega+1},$$

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 storage cannot be positive and long-run price means satisfy $\mu_\omega > \beta \mu_{\omega+1}$. 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 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 and computation

In this section, we describe the inputs needed for the calibration of the model. For demand, we use weekly demand data for the Nordic market in 2000-05 as published by the Organization for Nordic Transmission System Operators. We could use a longer data for demand estimation but this would be a source of problems because of trend growth. 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.⁶ 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 week. 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 1980-1999.⁷ 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. As to the minimum reservoir level, we use a lower bound of 10 TWh for the whole Nordic system.⁸ The lower bound of the aggregate reservoir level is based on the importance of hydroelectric resources as a fast power reserve supporting the electrical system. In principle, one can

⁶Demand for electricity showed little trend growth over the sample period. Testing for normality is difficult due to the fact that the data contains only six observations for each week. Nevertheless, a Shapiro-Wilk test supports the normality assumption, rejecting it (at the five percent level) only for weeks 12, 25 and 41. On average, $W = 90.0$ and $P = 46.9$.

⁷A Shapiro-Wilk test was applied to the inflow series (1980-99) of each week of the year. Averaging over the weeks, $W = 95.5$ and $P = 52.5$. The null hypothesis of log-normality was rejected at the five percent level for two weeks (weeks 25 and 29).

⁸Bye 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.

structurally estimate the unobserved reservoir and minimum flow constraints for a given market structure. In the section after the main results, we take steps toward this goal to evaluate if such constraints can produce similar implications as market power.

For the residual demand of hydro, 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.⁹ 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 fossil-fuel fired plants. In this market the theoretical non-hydro supply curve experiences considerable shifts because of heating demand (making electricity a side product) and planned maintenance outages. If we have knowledge, e.g. from historical data, of the heating demand and maintenance decisions, the engineering supply curve can be used in analysis. However, for hydro usage decisions, we need to know how the non-hydro supply curve is linked to the state of the market, because the value of water in a given state can be evaluated only by evaluating its future value in possible future states. For this reason, we cannot avoid estimating how the available non-hydro capacity is linked to the state of the market.

Rather than using 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. 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,

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

where z_t is the thermal supply, and q_t is the vector of fuel prices. The thermal generation

⁹A 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.

is composed of all other production than hydro, including 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.

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 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 in response to the shortage of hydro after the price spike.

We solve the model using a combination of backward induction and modified policy iteration. Modified policy iteration (see Puterman 1994) algorithms are a fast and easily implementable method for solving discrete time Markov Decision problems. Modified policy iteration replaces the policy evaluation step of normal policy iteration with a fixed number of successive approximation iterations. In our particular algorithm, we iterate over the value of water in the first week of the year. Backward induction is used in the policy improvement step of modified policy iteration to compute the optimal policy for each week within the year. 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. The policy estimate thus received is then evaluated by computing its value over a fixed number of years. In this evaluation step, the end value of water is given by the current estimate for the value function. 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 value function converges.

Discount rate is 7.5 per cent throughout this paper.

3.5 Results from the model of efficient allocation

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 mean reveals 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.¹⁰ 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 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 socially optimal prices risks 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

¹⁰The 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.

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 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 explain in detail various 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 (see Table 4).

4 Market power

4.1 The Model

Given the deviations between the calibrated planner's model and observed data, we now look for potential explanations for the deviations. We first develop an alternative market structure allowing for strategic management of the hydro reservoir. We do not seek to map the observed market characteristics such as market shares or 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.

Using the framework introduced in section 3, we now assume that a fraction of the reservoir capacity is strategically managed. The share for the strategic capacity, $\alpha \in [0, 1]$, is our market structure parameter for which we can search values best fitting in Section 4.3. An oligopolistic market structure with multidimensional state and complicated uncertainties becomes quickly intractable or, at least, produces implications that are hard to test empirically. Instead, we assume that fraction α of the total reservoir capacity 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 the 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 quite a 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

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

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 .¹¹

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 to include only perfectly correlated inflows at this moment: 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 .

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 the policy rule for the strategic agent depends on both states, $u_t^m = g^m(s_t)$. As said, we treat the nonstrategic agents as a single competitive unit and thus find a single policy rule for this unit, $u_t^c = g^c(u_t^m, s_t)$.¹² 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.

¹¹For 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.

¹²Notice 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.

Thus, no small agent can achieve higher profits by rearranging its production plan from what we describe below.

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

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

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 demand that is not served by the large agent. Let $v^c(u_t^m, s_t)$ 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 $g^c(u_t^m, s_t^m, s_t^c)$ solves the following recursive equation

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

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 $u_{t+1}^m = g^m(s_{t+1})$. 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:

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

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

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 price¹³. In the hydro market, the seller is not free to extend the sales path in 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 market share parameter α open, and consider in next what α provides the best match with the data. We would like find to the market 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 $u_t^c = g^c(u_t^m, s_t)$: a given u^c induces the transition

¹³See Hotelling (1931) for the analysis of a monopoly; Lewis and Schmalensee (1981) consider an oligopoly.

of the expected stock s_{t+1}^c , which when used together with s_{t+1}^m in $\tilde{u}_{t+1}^m = E_t g^m(s_{t+1})$ 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.

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 for various market 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, which is consistent with the interpretation given in section 4.2: the steady state stock increase is a way to organize the disposal of supply not meant to reach the market. While under uncertainty the logic of market power is slightly more intricate than in the deterministic case, as will be illustrated shortly, the implication for the 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 to 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 inflow 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 best matching the historical data and 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, because market power becomes evident through this variable and also because there is a systematic discrepancy between observed reservoir development and that chosen by the social planner. Including both prices and hydro outputs in the set of variables would clearly be unnecessary if the historical 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 is¹⁴

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

One criterion for choosing the model is to find 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 very crude way to proceed is to choose $T = 312$, i.e., 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 can be very 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

¹⁴We 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. See, for example, Cochrane (2001).

allow weighting the variances of the prediction errors when choosing W in $H_T(\alpha)$. The latter shortcoming can be avoided, for example, when $T = 12$ and the statistic $H_{12}(\alpha)$ is calculated separately for each of the 24 quarters in the data. Then, we can use the two-stage GMM approach¹⁵ 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.

We evaluate each model under different criteria ranging from moment restrictions for aggregated data to "path matching" using weekly data. Fortunately, all specifications support choosing the same α , namely $\alpha = .3$.¹⁶ For Table 4, we have first calculated the statistic $H_T(\alpha)$ for each model at different aggregation levels (weekly, monthly, quarterly, half-year, and annual prediction errors). 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.¹⁷ The mean value of the statistic $H_T(\alpha)$, obtained this way, is reported in the first column for each model. The 30 per cent model provides the best score at all time aggregation levels.

For quarterly, half-year, and 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, 12, and 6 samples (quarterly, half-year, and annual aggregation, resp.). Again, 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

¹⁵See, for example, Cochrane (2001).

¹⁶Recall that due to computational reasons we have computed only seven market share values for the strategic agent: 0, 15, 20, 25, 30, 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.

¹⁷Otherwise, 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.

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

[here we structurally estimate various unobserved constraints to distinguish the effect of potentially mismeasured constraints on the equilibrium from the effect of market power]

6 Concluding remarks

We conclude by discussing the objectives of this paper and whether these have been achieved. First, we presented a test for market power in a storage market. The hydro-electricity provided a good case for such a study because of the precise data on market fundamentals. The test involved developing a first-best and an alternative market structure, and applying both models to the historical data. We found that the alternative market structure can outperform the predicting power of the socially optimal model, but our statistical test remains rough at this point. We are also hoping to develop a more structural approach to market power, e.g., by estimating 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 as in Bajari-Benkard-Levin (2007).

Second, we conceptualized the Nordic market for hydroelectricity as partially an exhaustible-resource market and partially a storable-good market. These features are important for understanding both the price distributions and the price dynamics in a given state of the market. The relatively stylized "water-value" model, the planner's model, can well illustrate these features of the market, which have not been documented elsewhere. However, in its current state, the model remains aggregative. If one is interested in further developing a prediction tool, introducing heterogeneity in resource stock holdings and inflows would seem an important extension.

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

$$E_t \sum_{\tau=t+1}^{\infty} \varepsilon_\tau + \Delta = 0 \quad (13)$$

Changes in the optimal usage path, after the one-shot deviation from the optimal policy at t , are denoted by ε_τ , and these must in expectations sum up to the deviated amount at t (equality in (13)); assuming the opposite would imply nonoptimality of policy $g(s_t)$.

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 reduces to

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

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 (14), we have

$$\Delta < 0 \Leftrightarrow \bar{p}(s_t, \Delta) \geq E_t \beta^k \bar{p}(s_{t+k}, -\Delta) \text{ for some } k \geq 1. \quad (15)$$

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

$$\Delta > 0 \Leftrightarrow \bar{p}(s_t, \Delta) \leq E_t \beta^{k'} \bar{p}(s_{t+k'}, -\Delta) \text{ for some } k' \geq 1. \quad (16)$$

Finally, if the optimal policy is not constrained, then both (15) and (16) must hold at s_t . ■

References

- [1] Amundsen, E. and L. Bergman, Why has the Nordic electricity market worked so well?, *Utilities Policy*, Vol. 14, No. 3, (Sep. 2006), pp.148-157
- [2] Bajari, P., C. Benkard, J. Levin (2007). Estimating Dynamic Models of Imperfect Competition, *Econometrica* 75 (5) , 1331–1370
- [3] Borenstein, S., J B. Bushnell, and F A. Wolak. 2002. Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market. *American Economic Review* 92: 1376-1405.
- [4] Bye, T. and A. Bruvoll and F. Aune, The Importance of Volatility in Inflow in a Deregulated Hydro-Dominated Power Market. *Discussion Papers No. 472*. Statistics Norway, Research Department.
- [5] Cochrane, J. H. (2001). *Asset Pricing*. Princeton University Press, Princeton.
- [6] Deaton, A., and G. Laroque, On the Behaviour of Commodity Prices, *The Review of Economic Studies*, Vol. 59, No. 1 (Jan., 1992), pp. 1-23

- [7] Deaton, A., and G. Laroque, Competitive Storage and Commodity Price Dynamics, *The Journal of Political Economy*, Vol. 104, No. 5 (Oct., 1996), pp. 896-923
- [8] Lewis, T. and R.Schmalensee (1980), On oligopolistic markets for nonrenewable resources, *Quarterly Journal of Economics* 95, 475-491.
- [9] Loury, G. (1986), A theory of ‘oil’igopoly: Cournot–Nash equilibrium in exhaustible resources markets with fixed supplies. *International Economic Review* 27, 285–301.
- [10] Mc Laren, J., Speculation on Primary Commodities: The Effects of Restricted Entry, *The Review of Economic Studies*, Vol. 66, No. 4. (Oct., 1999), pp. 853-871.
- [11] David M. Newbery, Commodity Price Stabilization in Imperfect or Cartelized Markets, *Econometrica*, Vol. 52, No. 3. (May, 1984), pp. 563-578.
- [12] Polansky, S. (1992), Do oil producers act as ‘oil’ilopolists?, *Journal of Environmental Economics and Management* 23, 216–247.
- [13] Puterman, M.L. (1994). Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, Inc, New York, NY.
- [14] Rotemberg, J. J., and G. Saloner, The Cyclical Behavior of Strategic Inventories, *The Quarterly Journal of Economics*, Vol. 104, No. 1. (Feb., 1989), pp. 73-97.
- [15] Salant, S.W. (1976), Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market, *Journal of Political Economy* 84, 1079-1093.
- [16] Salo, S., and O. Tahvonen (2001), Oligopoly equilibria in nonrenewable resource markets, *Journal of Economic Dynamics and Control*, Volume 25, Issue 5, 671-702
- [17] Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989): Recursive Methods in Economic Dynamics. Harvard University Press.
- [18] Williams, J. and Wright, B. (1991) Storage and Commodity Markets (Cambridge: Cambridge University Press).
- [19] Wolfram, C.D. Measuring Duopoly Power in the British Electricity Spot Market.” *American Economic Review*, (Sept.1999), 89(4), pp. 805–26.

Inflow distribution in the Nordic market area 1980-99




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).

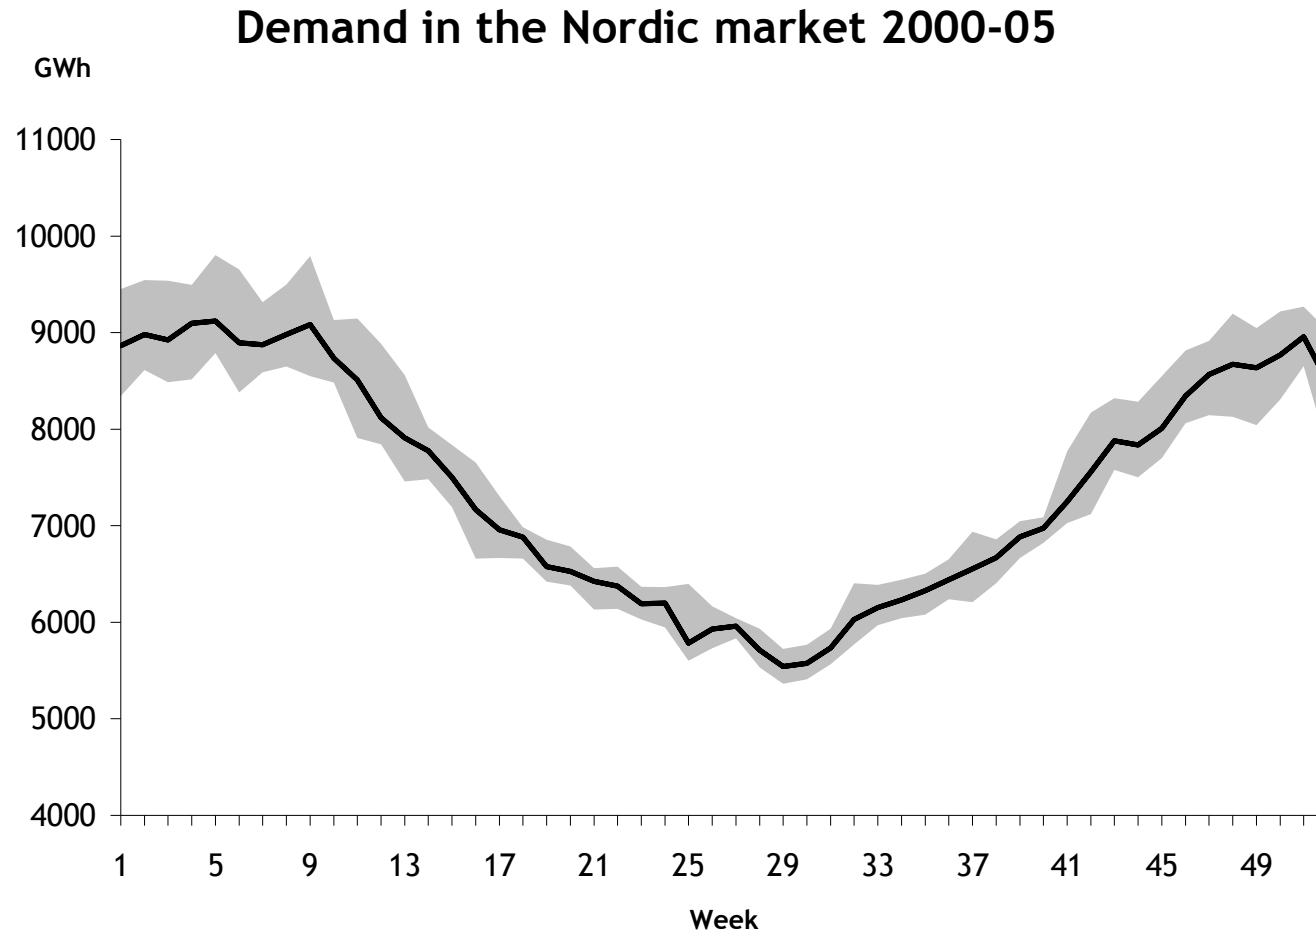


Figure 2: Mean and empirical support of demand in Nordic market 2000-05




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

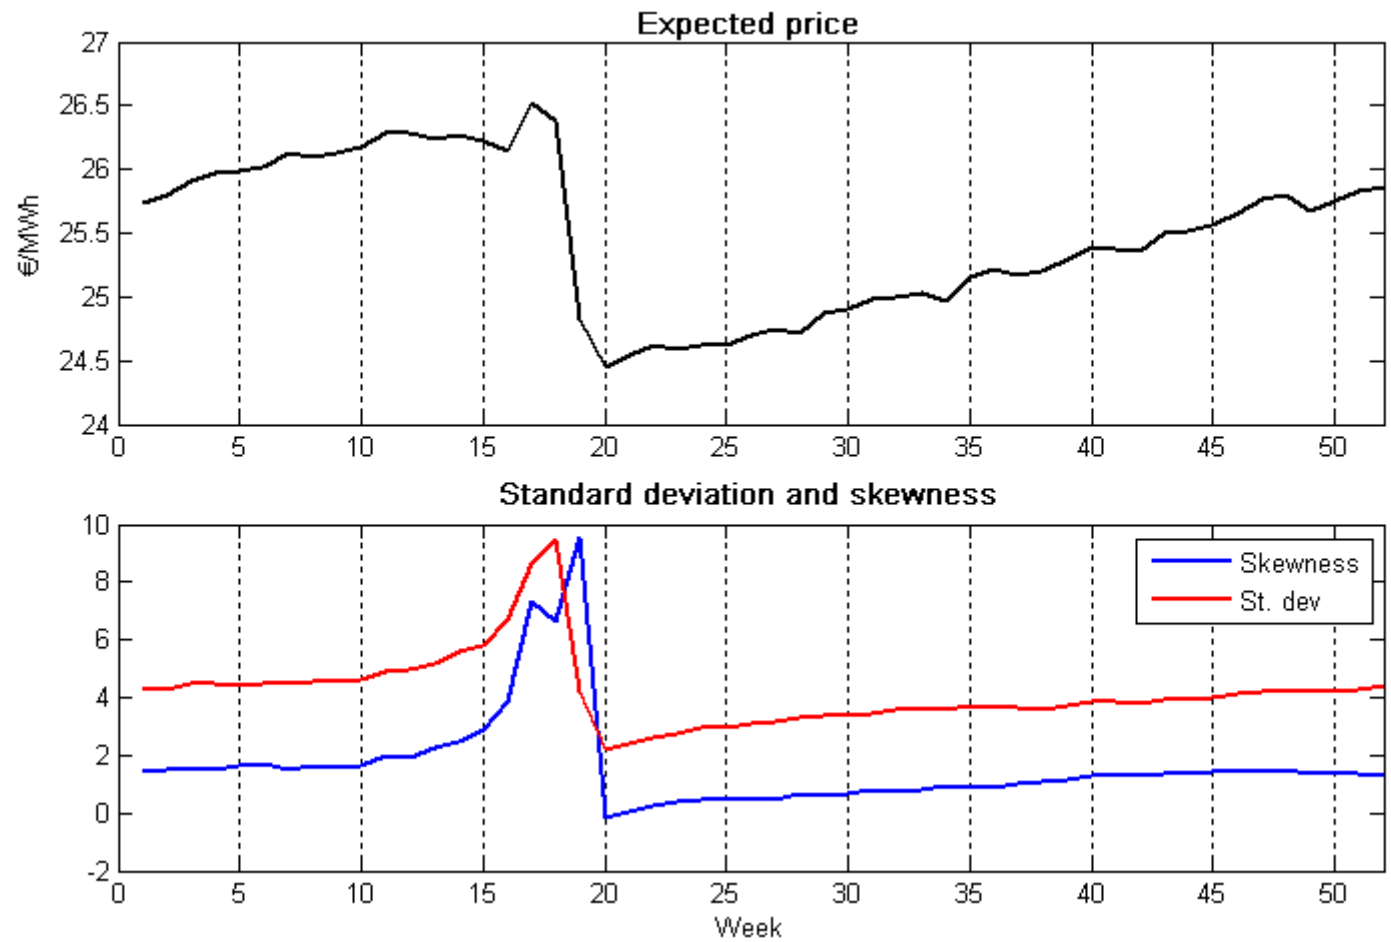


Figure 4: Simulated expected price (upper panel) and the skewness and standard deviation (lower panel) of price




Figure 5: Social planner output (blue), observed output (black), social planner reservoir levels (blue), observed levels (black).

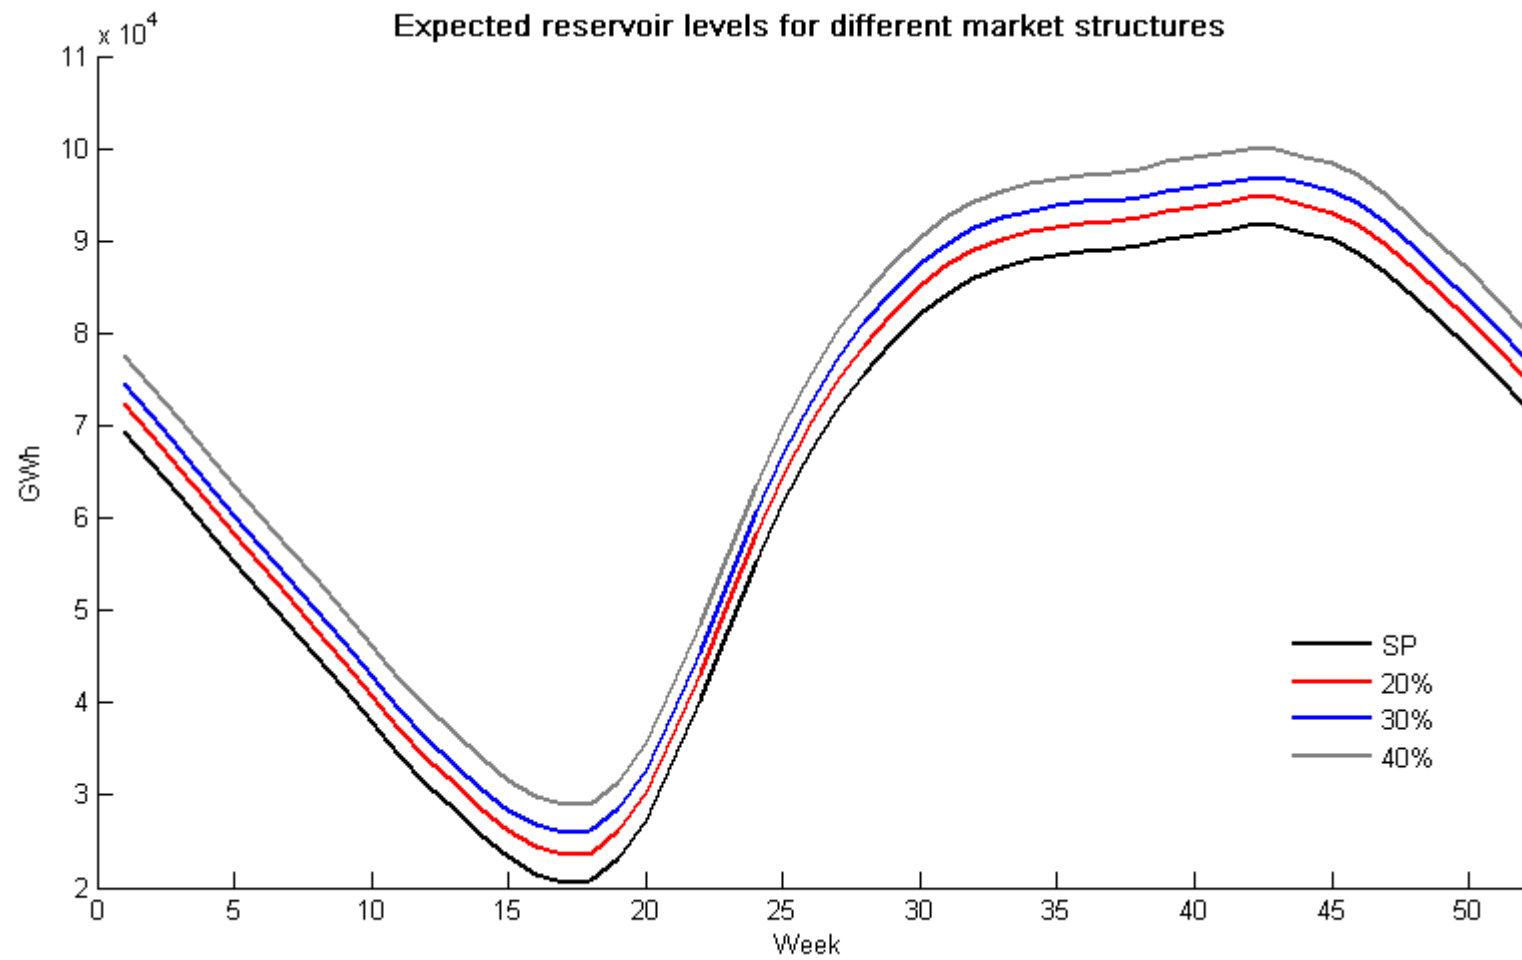


Figure 6: Simulated expected reservoir levels for different market structures.

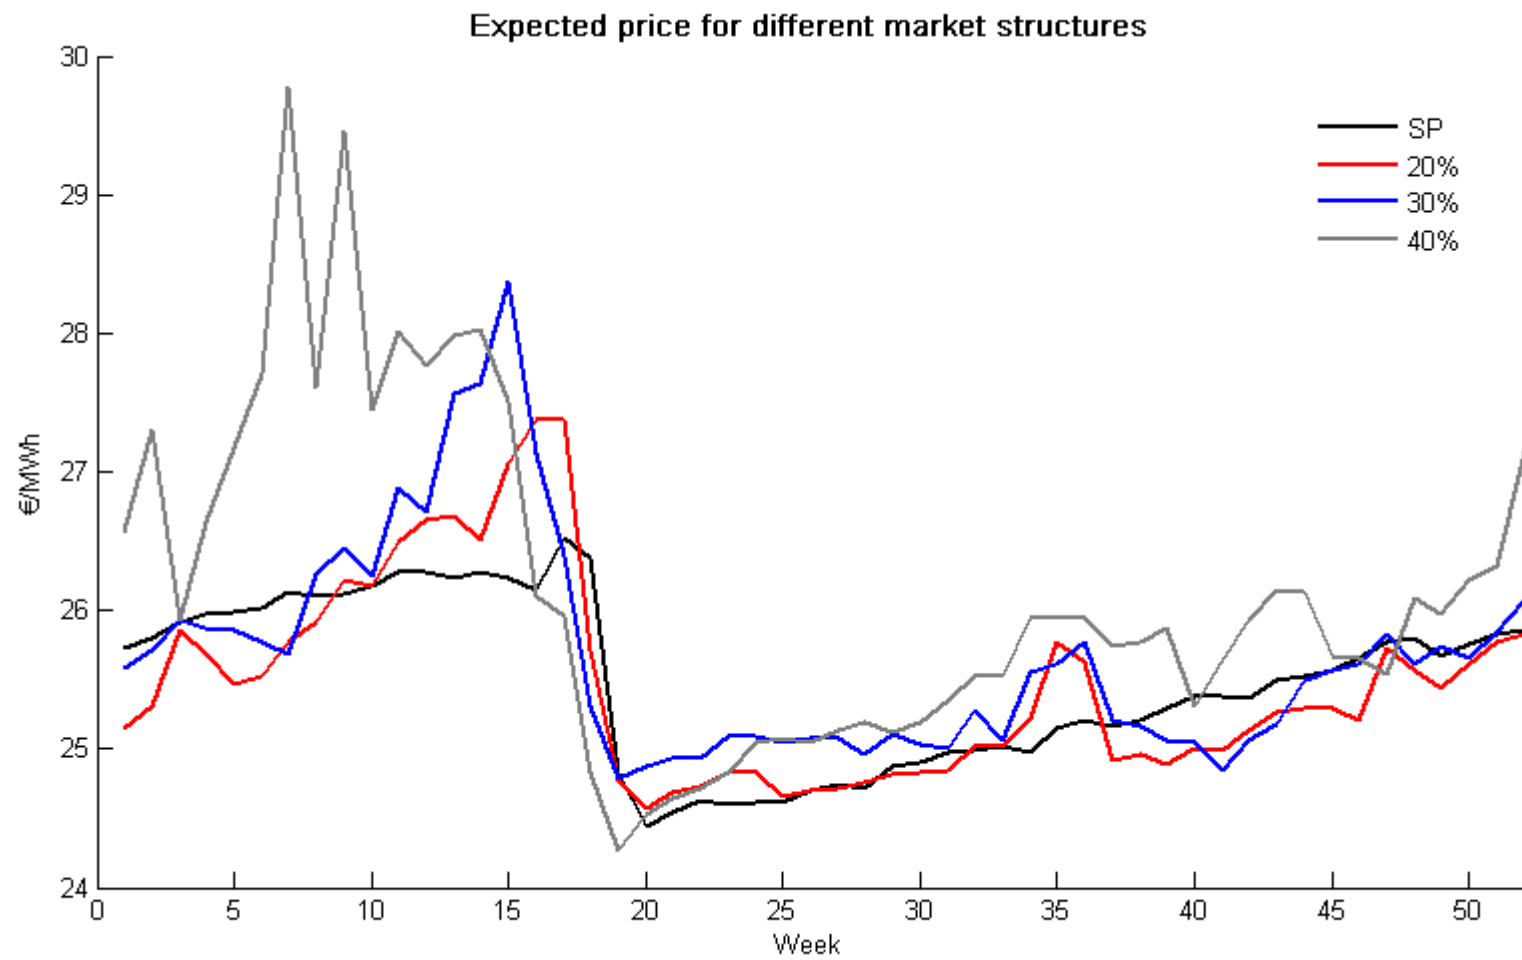


Figure 7: Simulated weekly price expectations under different market structures.

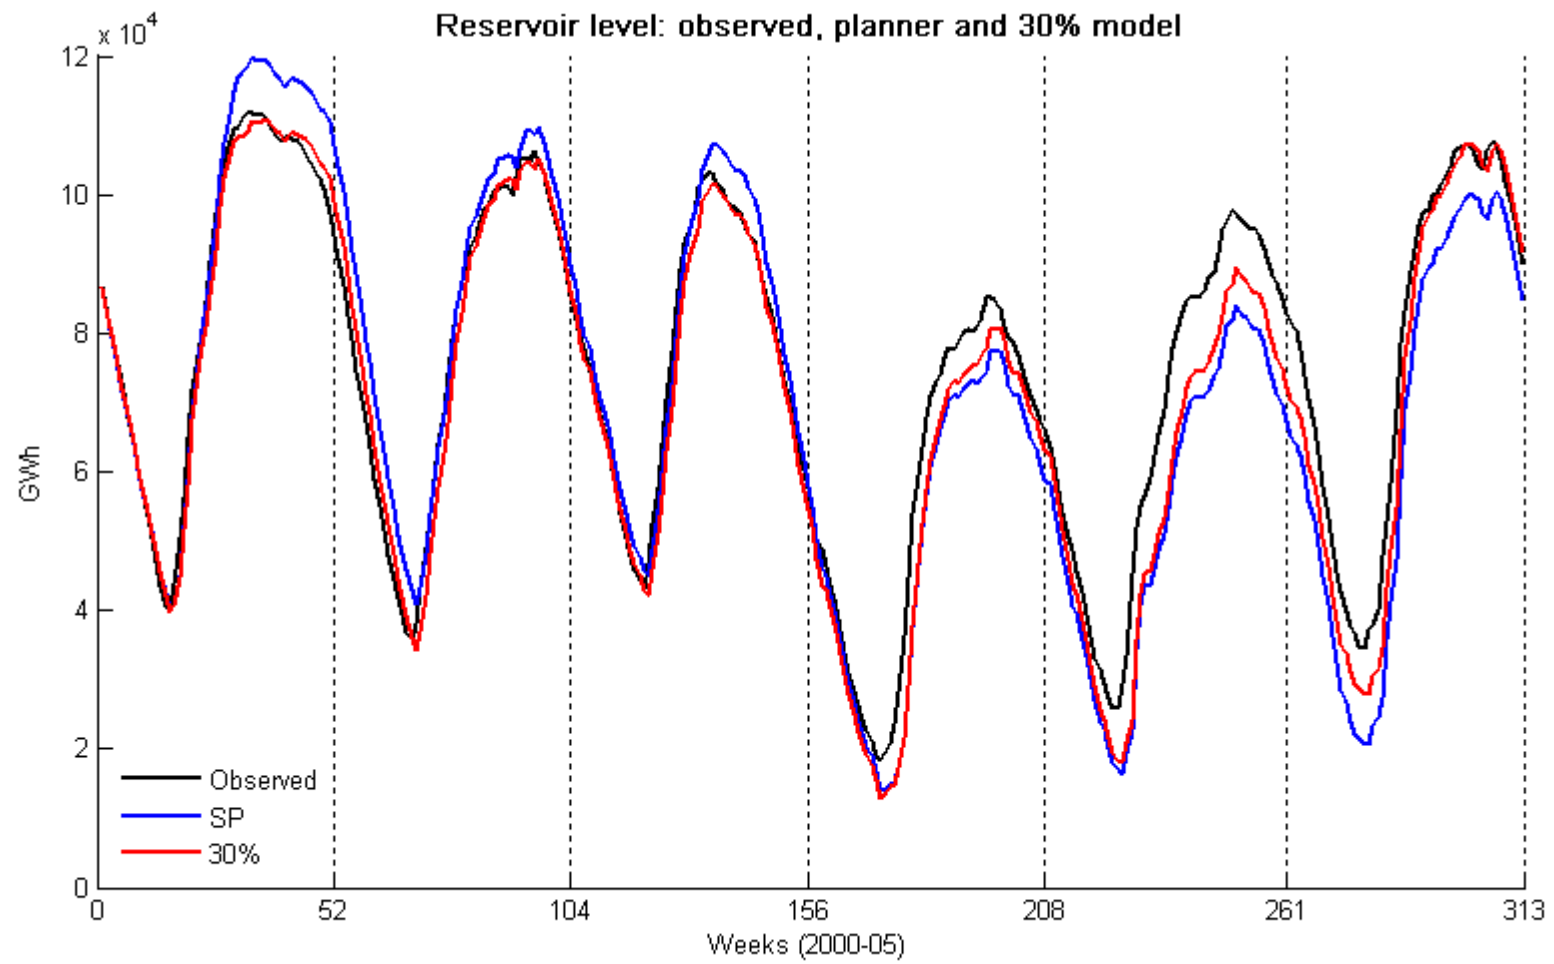


Figure 8: Historical, the planner's, and market power (30%) storage levels.

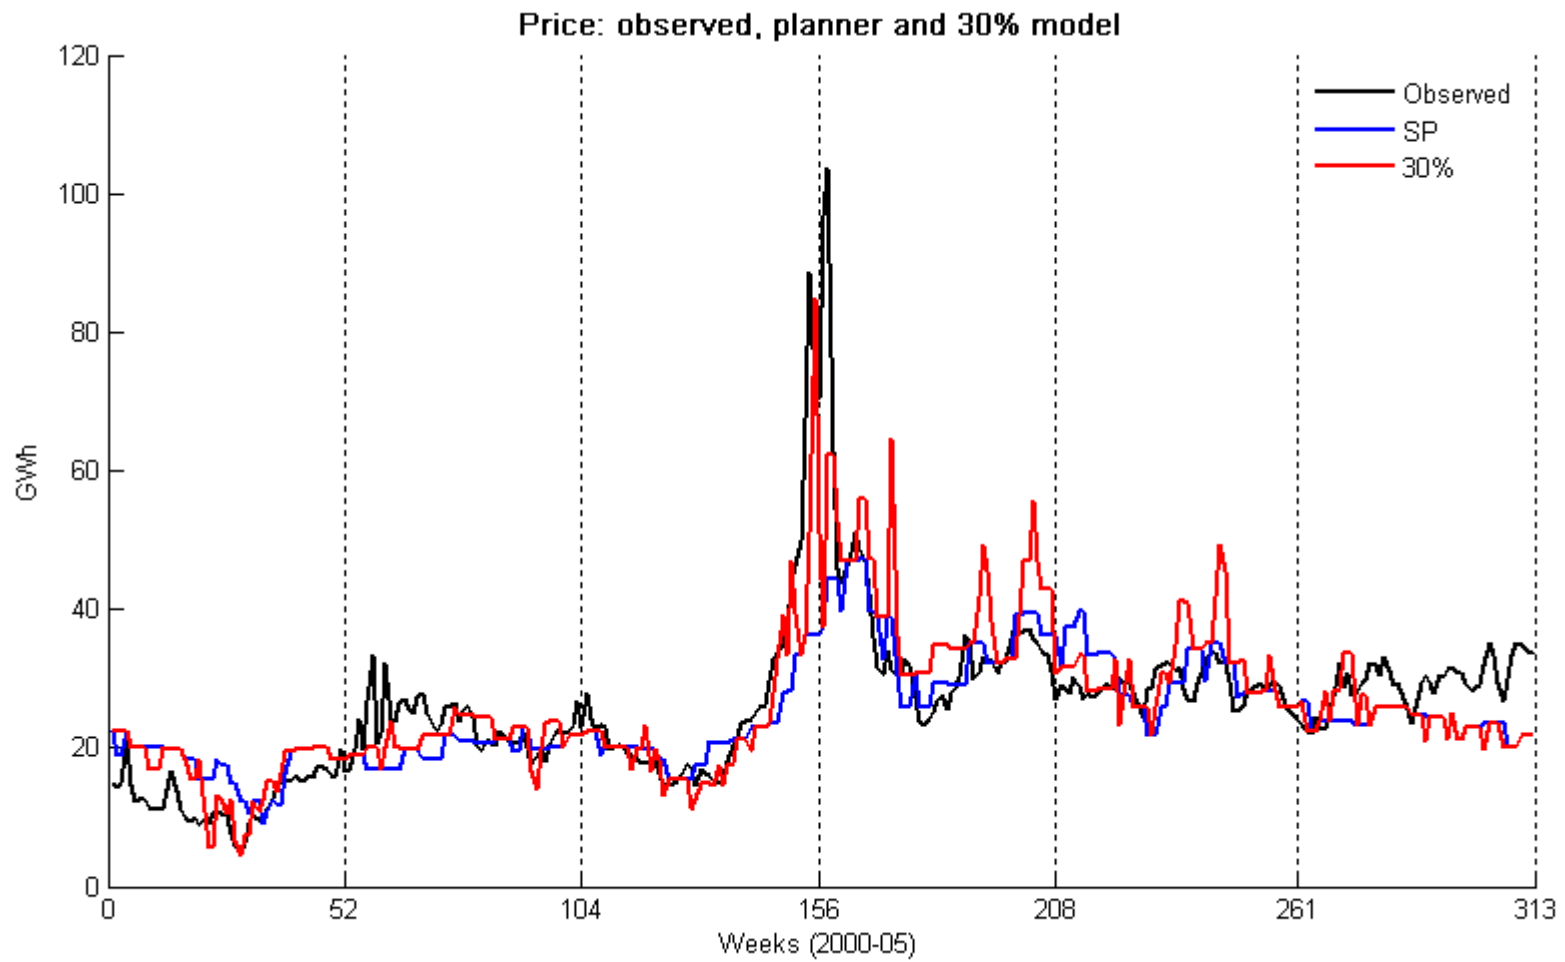


Figure 9: Historical, the socially optimal, and the market power (30%) price.

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	0.0187**	0.0185**	0.0187**
	(0.0018)	(0.0016)	(0.0018)
Coal price	-0.0013		
	(0.0014)		
Reservoir level	-0.0282**	-0.0291**	
	(0.0015)	(0.0015)	
Hydro output			-0.0006**
			(0.00004)
Observations	300	313	313
R-squared	0.70	0.69	0.58

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

	(1)	(2)	(3)
ln(price)	1192.7**	1177.8**	1252.4**
	(43.8)	(43.3)	(48.6)
Oil price	-25.5**	-22.7**	-23.2**
	(1.6)	(1.5)	(1.6)
Coal price	6.3**		
	(1.5)		
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		15		20		25		30		40		50	
1	1.21	-	1.00	-	0.82	-	0.68	-	0.55	-	0.66	-	0.91	-
4	1.20	-	0.98	-	0.80	-	0.66	-	0.53	-	0.64	-	0.89	-
13	1.14	28.0	0.93	21.5	0.75	16.4	0.61	12.2	0.48	8.2	0.57	15.9	0.78	21.1
26	1.06	9.5	0.84	7.3	0.67	5.8	0.53	4.3	0.40	3.2	0.47	5.9	0.56	10.4
52	0.94	5.8	0.73	4.2	0.58	3.3	0.46	2.4	0.35	1.7	0.37	3.4	0.48	4.3

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).

	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.