ESSAYS ON ENERGY MARKETS

BY **RITVANA RRUKAJ**

Department of Business and Management Science
NHH Norwegian School of Economics

Acknowledgement

This thesis concludes an exciting journey that started in 2017. I consider myself fortunate to have been part of the Department of Business and Management Science and to have been granted the opportunity to conduct research on this topic. Above all, I would like to thank my main supervisor, Prof. Leif K. Sandal, for the innumerable hours he has spent discussing my research and guiding me along the path to attaining my PhD. With his enthusiasm and immense knowledge and his unwavering efforts to explain things clearly and simply, he helped me become more interested in my research topic, and he continuously encouraged me during my academic research and in daily life. I would like to thank my co-supervisor, Prof. Jan Ubøe, for supporting me and providing a positive environment, especially during the first two years of my studies.

I would also like to express my immense gratitude to my co-author, Prof. Frode Steen, who was very generous with his invaluable support and helpful advice. I also thank him for his interest in my research, sharing data, and providing excellent collaboration. I am grateful to David P. Byrne, from the University of Melbourne and Øystein Foros from NHH for their valuable comments, which have helped me improve the quality of my dissertation. In addition, I am also grateful to my former advisor, Nevila Mehmetaj, from the University of Shkodra "Luigj Gurakuqi," who encouraged me to apply for the PhD program abroad and provided continuous support during my journey. I am also thankful to Benjamin P. Fram for his long-lasting friendship and collaboration.

The support of the administration team at the department has also been invaluable. Special thanks to Kristin, Charlotte, Natalia, Stein, and Torill. I am grateful to them for assisting me in many different ways.

I would like to thank my colleagues and friends Aija, Andreas Ø, Atle, Christer, Christian, Evangelos T, Evan K, Henrik, Gabriel, Lars, Mostafa, Ondrej, Somayeh, Yuanming, and many more for the lunches, chats, and camaraderie. I want to thank my friends and flat mates, Huynh and Tatevik, for their steadfast support through the toughest moments during this journey and also my friends, Beatriz, Jesmina, Kyriaki, Lori, Mai, Nahid, and Rabia, for all the emotional support, entertainment, and care they provided.

I cannot end without thanking my family, whose constant encouragement and love I have relied on throughout this journey. I would like to express immense gratitude to my parents, Kadri and Vera, for always believing in me and giving me and my sisters Suela, Rajmonda, Suada, and Migena, the best education they could. I dedicate this thesis to them. I also thank Simone, who has been a constant source of support and encouragement. Understanding me best as a Ph.D. himself, Simone has been my best friend and a great companion during the last year of my Ph.D.

Contents

General Introduction

Energy markets comprise a dynamic environment involving millions of transactions, sometimes repetitive, between thousands of organizations, impacting billions of people worldwide. Energy products possess peculiar characteristics, such as non-storability, seasonality, extreme high-price volatility (power), scarcity, high costs, and limited predictability, making them very different from classic financial market products (Eydeland and Wolyniec, 2003). Technological progress has provided us with more and more tools to deal with energy management problems. We are now able to explore the energy grid, collect data, forecast demand, and hedge market risk. However, at the same time, knowing the particular characteristics that energy commodities possess, a thorough and careful analysis of their features is required. A long-term plan with efficient decision-making mechanisms could result in optimal hedging strategies and future profit-making investments in the field.

On a global scale, many countries are at the forefront of promoting cleaner energy supply, which has increased interdependencies between various energy infrastructures that generate, convert, and provide different forms of energy. This, in turn, brings out the instabilities in the energy grid and the uncertainties related to energy security, as consumers have more diversified and complicated energy usage patterns.

To ensure better management and construct more accurate predictive models, it is very important to explore these markets and provide a deeper insight into the diverse issues. Energy demand and price dynamics have been extensively studied in the literature. Many empirical studies have focused on price evolution processes and energy assets, such as power plants and gas storage (Cook and Green, 2005; Chen and Forsyth, 2007; Carmona and Ludkovski, 2010). As the literature in energy markets evolves, many studies have begun addressing other aspects, including multi-energy systems (Zhou et al., 2021), mapping techniques (Benth et al., 2008; Taylor, 2011), climate uncertainty (Blazev, 2021), artificial intelligence tools (Kempitiya et al., 2020), and other economic and financial factors (Xu et al., 2021).

The energy sector is a fast-growing and complex sector with multiple economic, political, or technological challenges. These challenges are linked to the requirement of the industry to meet a wide array of social demands, which can be related to the three tra-

ditional dimensions of a sustainable energy system: environmental sustainability (greenhouse gas mitigation), security of energy supply (diversification of energy sources and reliability of supply), and economic sustainability (a competitive energy system and affordable energy supply).

In this thesis, we explore various aspects of stochastic energy demand and price modeling. We attempt to fill some of the gaps in this research area by introducing new ideas on modeling market dynamics and predicting trends for tackling daily energy market challenges. We outline how a data-driven approach can effectively support various challenging energy market issues such as prediction, pattern recognition, competition, modeling, and so on.

Empirical applications have been the subject of many studies in applied economics. Some of these applications focus on perfectly competitive markets while others focus on markets that are characterized by monopolistic or oligopolistic behavior. Moreover, we do not know the degree of competitiveness in certain markets and would, therefore, be interested in assessing it. The identification of market structure and the measurement of the degree of competitiveness are, in fact, among the most important issues in industrial organizations. Industrial organization studies usually use a variety of indexes of monopoly as a means for identifying market structure.

The competitive environment in the retail gasoline market, particularly the Swedish gasoline market, is a main challenge that I attempted to address in Chapter 1 of this dissertation. The real gross margins of five oil brands across six different geographical markets are examined and compared to the margins in rural areas. The concentration ratio across regions is used to measure monopoly power, and rural markets, as anticipated, have gas stations with the most distance between them. Service and self-service brands have been analyzed separately. We found that a relatively stable average gross margin pattern can be observed for both service and self-service oil stations, as each brand successfully targets a relatively fixed average real gross margin value annually. These gross margin patterns were compared with the theoretical background, whereby some differences and counterintuitive aspects have been identified . The results of the empirical analysis suggest that high-margin markets have less competition and seem to indicate differences between service and self-service brands' profits.

The transport sector in many countries relies heavily on gasoline, especially in rural areas where public transportation is not well developed. Thus, predicting gasoline consumption is an important function in planning for future demand in order to address the environmental consequences of gasoline consumption, particularly with respect to the emission of greenhouse gases. Therefore, the gasoline consumption and price relationship is another main topic that I investigate in Chapter 2 of this dissertation. We develop a dynamic method to describe and model gasoline quantity patterns using price and time

as inputs. In addition, our model is capable of learning nonlinear mappings between inputs and output and is suited for gasoline consumption predictions. We perform this analysis for the Swedish gasoline market while analyzing the market dynamics of the main cities, the E6 highway, smaller cities, rural districts, and five oil brands in Sweden. The proposed model illustrates the power of the structural time series modeling in identifying the structural relationships between economic variables, underlying dynamics, and the prediction of future gasoline consumption.

Asymmetric price adjustment is a common phenomenon in several markets around the world, particularly in retail gasoline markets. Chapter 3 of this dissertation studies this phenomenon in the context of the retail gasoline market in Sweden using a dataset of daily gas station volume and prices that covers a year. We establish short-run cost pass-through for both average pump prices and volume-adjusted gasoline prices and find that the disequilibrium error takes longer to be fully corrected when we test for the existence of volume-adjusted price asymmetry compared to pump price asymmetry. This finding suggests that oil companies are more focused on pricing on days and at stations with larger sales. Estimating Non-linear Autoregressive Distributed Lags models across regions, different station service levels, and local competition levels, we find that rural areas and gasoline stations less exposed to local competition are able to impose larger and more prolonged asymmetry on retail gasoline prices. This supports the conclusion of several studies, that asymmetry in cost pass-through positively correlates with the degree of market power. The same pattern is uncovered when we estimate models for varying service levels. Full-service stations have a higher and more prolonged asymmetry in pricing than automated, self-service stations. In several studies, the former group of stations has been found to have a higher market power than the latter.

The late 1990s were followed by a large increase in the number of energy modeling studies, in particular those that aimed to forecast future electricity demand. A large number of studies focused on short-term load forecasting over a time horizon of hours. These forecasts can help market participants devise their bidding strategies in the auction-based, pool-type markets. Less attention has been given to approaches that make medium-term load forecasts, although, medium-term forecasting methods are useful for balance sheet calculations and risk management applications. Chapter 4 of this dissertation focuses on the relationship between seasonal load and prices in the New York wholesale electricity market (NYISO). In this chapter, we propose a new modeling approach that is useful in predicting the seasonal load in the NYISO, which may assist market participants in assessing market risk better and adjusting hedging positions. We first sort and filter raw market data from 2006 to 2018 into annual, seasonal (90-day), and weekly moving averages. We then develop three non-linear waveform models that relate the moving averages of load and prices while accounting for seasonality in the data. As a next step, we

extract the waveform from the filtered data and employ the Box-Jenkins methodology for the time series analysis to model the resulting noisy residuals. We conclude by performing out-of-sample forecasts of the seasonal moving average model using data from 2019 and 2020. The results provide evidence that the approach developed in this paper could forecast the relationship between load and prices in NYISO with a high degree of accuracy.

This dissertation investigates the demand and price dynamics in energy markets. It focuses on some specific energy industries, namely the gasoline and electricity industries, and attempts to provide useful pre-processing data techniques and dynamic models that can help market participants make more informed decisions.

My thesis consists of four chapters. Every chapter is structured as a self-contained chapter, with a separate bibliography. All four chapters have a common research focus: the energy field and empirical methodologies. The goal is to propose innovative modeling approaches and data pre-processing techniques that can play a key role as decision-support tools, which can estimate risk and generate recommendations for policy making. The readers will find a wide spectrum of methods and techniques across the chapters, such as non-linear data analysis, data filtration techniques, waveform functions, and predictive models.

Bibliography

[1] Blazev A. S., (2021), "Global energy market trends, " *CRC Press*.

[2] Chen Z., and Forsyth P.A., (2007), "A semi-lagrangian approach for natural gas storage valuation and optimal operation, " *SIAM Journal on Scientific Computing*, Vol.30, iss. 4, pp. 339-368.

[3] Cook W. D., and Green H. R., (2005), "Evaluating power plant efficiency: a hierarchical model, " *Computers & Operations Research*, Vol.32, iss. 4, pp. 813-823.

[4] Eydeland A. and Wolyniec K., (2003), "Energy and power risk management: New developments in modeling, pricing, and hedging," *John Wiley & Sons Inc.*.

[5] Taylor James W., (2011), "Short-term load forecasting with exponentially weighted methods, " *IEEE Transactions on Power Systems*, Vol.27, iss. 1, pp. 458-464.

[6] Benth F. E., Benth J. S., and Koekebakker S., (2008), "Stochastic modelling of electricity and related markets, " *World Scientific*, Vol.11.

[7] Kempitiya T., Sierla S., De Silva D., Yli-Ojanperä M., Alahakoon D., and Vyatkin V., (2020), "An Artificial Intelligence framework for bidding optimization with uncertainty in multiple frequency reserve markets, " *Applied Energy*, Vol.280, pp.

[8] Xu B., Fu R., and Lau C. K. M., (2021), "Energy market uncertainty and the impact on the crude oil prices, " *Journal of Environmental Management*, Vol.298, pp. 113403.

[9] Zhou Y., Yu W., Zhu Sh., Yang B., and He J., (2021), "Distributionally robust chance-constrained energy management of an integrated retailer in the multi-energy market, " *Applied Energy*, Vol.286, pp. 116516.

Chapter 1

The real effects of gross margin patterns: Evidence from self/service oil brands

Ritvana Rrukaj[a*]

[a]Department of Business and Management Science, NHH Norwegian School of Economics, 5045 Bergen, Norway

Abstract

Using descriptive techniques and detailed data, we attempt to describe the competitive environment in the Swedish gasoline market. In particular, we can draw conclusions on the local competition by constructing real gross margins and comparing these across oil brands and regions with different characteristics. The service and self-service brands have been analyzed separately. A relatively stable average gross margin pattern can be observed for both service and self-service oil stations, as each brand successively targets a relatively fixed average real gross margin value annually. The observed mean-revision patterns suggest that the real gross margin volatility and historical margin values will eventually revert to the long-run mean. In addition, the results suggest less competition in high-margin markets and seem to indicate differences between service and self-service brands' profits. Surprisingly, each brand seems to get a higher profit in Stockholm, and their margins resemble the rural district's margin levels while varying across other regional markets. Finally, the results of our empirical analysis on oil brands' margins show the ability of brands to generate stable profits over time and, in some regional markets, even to earn a large profit. These findings have important implications for potential competitors in the market and policymakers who wish to encourage effective competition.

Keywords: real gross margins, local competition, (self) service stations, price rigidity, mean-revision.

*I thank Leif K. Sandal, Øystein Foros, and Frode Steen for their useful comments and suggestions and also for providing data access.

1.1 Introduction

Although highly vertically integrated oil companies may not necessarily engage in implicit or explicit cooperation, retail gasoline prices should be impacted by the degree of competition in retailing. Besides, wholesale prices should also drive retail price fluctuations. This study explores the firms' real gross margins, which reflect the volume-adjusted retail prices minus costs. Studies on gasoline margins have not examined the real gross margins due to the station-specific gasoline volume data requirement imposed by such a study. However, having an idea of firms' real margin strategies is important for uncovering the degree of competitive market constraints.

Therefore, we attempt to contribute to the literature by using detailed Swedish daily volume and price data to examine five oil brands' real gross margins by simultaneously considering two different categories of gasoline brands: service and self-service. Specifically, we investigate the brand's real gross margins at the national level and across six different local markets between January 1, 2012 and December 31, 2012. A significant positive relationship between service brands' real gross margins and market concentration supports the idea that government incentives to support new entry in the retail market will enhance competition and consequently lead to lower margins. Similar implications would apply to the self-service market. On the contrary, obtaining significant market power across different regions, accompanied by significant differences in margin levels, suggests that government intervention to investigate anticompetitive behavior at the national level might have a limited impact. Byrne and De Roos (2019) and Foros and Nguyen-Ones (2021) are examples of studies that use descriptive techniques to study the presence of coordination and coordinated effects. However, such papers do not consider the real gross margins and the service and self-service brands separately across different geographical areas and over time. Foros and Nguyen-Ones (2021) show how the market leader's recommended prices are used to coordinate the timing and the level of industry-wide price jumps in the Norwegian gasoline market. In comparison, Bryne and De Roos (2019) document how gasoline retailers in an urban market in Australia used the price-transparency program called FuelWatch to initiate and sustain cooperation. Hence both studies examine the presence of price-leadership follower as a mechanism through which firms create focal points to originate and sustain cooperation. Foros and Nguyen-Ones (2021) focus on the systematic use of daily recommended prices. Byrne and De Roos (2019) concentrate on the frequent use of everyday station-level prices to organize oil companies. Hence, the primary distinction between our observations and those of Foros and Nguyen-Ones (2021) and Byrne and De Roos (2019) relates to what daily station-level prices represent. In contrast to these studies, we consider brand volume-adjusted prices and real gross margins across urban and rural markets facing different competition constraints. Analogously, we draw attention to how brands are more integrated in some urban and rural markets than

others. Within this line of research, this study revisits the relationship between the intensity of market competition and the firms' margins by focusing on Swedish gasoline retailing. How and why do brands' real gross margin patterns change across urban and local markets over time? Are there lower margin levels in more competitive markets? This study answers these questions through a unique empirical analysis of real gross margins of brands. We analyze a rich consecutive dataset that spans a year and contains station-specific observations from six different geographical areas.

Both tax and wholesale prices are given, which firms cannot influence. Therefore, real gross margins of brands describe the firms' price adaptation over time. Foros and Steen (2013) and Nguyen-Ones and Steen (2019) also describe the market features during the sample period. Foros and Steen (2013) examine retail pricing, vertical control, and competition in the Swedish gasoline market. Their descriptive analysis also reveals some similarities between Stockholm and the price patterns in rural areas. Later, Nguyen-Ones and Steen (2019) investigated the degree of competition in the Swedish gasoline market by estimating a structural model that considered the main features of both demand and supply. In contrast, we focus only on the brands' real gross margins and separately explore the local competition conditions for service and self-service brands.

All retail stations throughout the six different geographical areas adjust their real gross margins according to cost changes every day. The data shows similar margin patterns across brands on the long-run time interval. We find that service brands, such as Statoil, OK-Q8, and Shell, are positively correlated on time in adjusting their real gross margins due to cost changes at the national and regional level. Nationally, we observe a distinct picture for the self-service brands. St1 always charges on average ten Öre/liter higher real gross margins than Jet, and their profits never cross. The difference between the margins of St1 and Jet gets closer in Stockholm compared to other regions. Overall, the five oil brands seem to target a relatively fixed average profit value on their sales patterns' adaptation that reflects a mean-reverting process.

We consider the average distance to the nearest competitor across different regional markets to measure the competition environment. The largest span is observed in rural areas where the closest competitor is over 4 km away. Therefore, we consider the level of real gross margins in rural areas to measure the profit of isolated markets. The data shows that, surprisingly, all brands generate the highest profits in Stockholm.

Furthermore, the data shows that both the real gross margin level and its pattern in Stockholm resemble the market features of rural areas. Therefore, based on this monopolization-like pattern in Stockholm, it is reasonable to think that the markets are less competitive in Stockholm and the rural districts than other main cities in Sweden. The analysis of the brand margins at the station level also supports this finding. The spread of real gross margins between stations of the same brand is bigger in Malmö than in Stockholm and the rural areas.

Further, the data shows that service branded stations have different market positions in terms of margins across regions, and that their daily moving average of margins follows a stable trend. Among self-service brands, St1 is the leader in terms of margins in all local markets. It never lowers its margins to Jet's level, although the difference is significantly lower in Stockholm. Our rich and complete dataset allows a purely descriptive analysis for ascertaining these empirical findings.

The remainder of this paper is organized as follows. Section 1.2 provides a brief review of the empirical studies investigating the relationship between competition and prices. Section 1.3 contains the data description, and Section 1.4 presents our empirical findings and 1.5 discusses them. Section 1.6 contains the concluding remarks and discusses the potential for future research.

1.2 Review of literature on the relationship between competition and co-operation

According to Encaouna and Geroski (1986), firms in more competitive industries face higher uncertainty about their future position in the market. In that regard, they may be more concerned with ensuring short-run margins, which lead to higher responsiveness to current cost or demand shocks. By contrast, the price dynamics of firms facing less competition are more oriented toward long-run objectives than short-run costs or demand fluctuations, resulting in a higher price rigidity. Alternatively, Stiglitz (1984) argues that oligopolists may prefer to delay adjusting prices to avoid breaking tacit pricing understandings.

Maskin and Tirole (1988) posit that firms alternate between changing prices. This can work like a firm's commitment not to change prices instantaneously, as in the classic Bertrand competition model. The authors show that there exists a perfect Markov strategy that allows firms to coordinate on the monopoly price. They also show that the commitment to stick to a fixed price for an extended period can lead to cooperation and alignment. This fact has already been supported by several empirical studies that investigated the link between competition and price flexibility (Alvarez and Hernando, 2006; Cornille and Dossche, 2006). The existence of greater price stickiness in markets characterized by higher concentration levels is viewed as an inverse proxy for market competition. Furthermore, the results of these studies corroborate previous findings of the positive relationship between price change frequency and the speed of difference to the degree of competition.

Another strand of literature explored the rockets and feathers phenomena in the gasoline markets. The studies investigated the factors that may explain why retail prices respond quicker to wholesale price increases and slower to wholesale price decreases. Among the potential reasons for asymmetric cost pass-through, consumer-search behavior and coordination theories have received the most attention in the literature. The latter implies that previous retail prices may be the focal points based on which retailers can collude when oil prices decrease. On the contrary, retailers would quickly raise retail prices when oil price increases

because continuing to maintain the previous retail price would be unprofitable. Several studies have empirically found consistency in the collusion theory for explaining asymmetric pricing. Using monthly state-level US data, Deltas (2008) observes that markets with higher average retail–wholesale margins experience a slower adjustment and a more asymmetric response. In addition, Deltas (2008) also shows that a higher degree of retail market power results in a slower adjustment speed. In other words, asymmetric pricing is more prevalent in less competitive gasoline markets. This finding is also supported by Verlinda (2008) and Balmaceda and Soruco (2008). Verlinda (2008) shows that the asymmetric pricing behavior falls within the geographical proximity of gasoline stations. While Balmaceda and Soruco (2008), using station-level data from Chile, find that the asymmetric responses are much more prevalent for branded gas stations than unbranded stations. Later, the findings provided by Lewis (2015) on the US gasoline retail market also suggest that the use of odd prices, particularly those ending in 5 or 9, results in higher margins and that odd prices can be an effective mechanism to sustain collusion.

On the contrary, the search-based theory (Yang and Ye, 2008; Tappata, 2009; Lewis, 2011; Cabral and Fishman, 2012) is based on the idea that consumers are more likely to intensify search following cost increase than cost decrease. Yang and Ye (2008) have indicated that cost-change search and learning may play a crucial role in the price adjustment process. Hastings and Shapiro (2013) confirm such a consumer pattern; they posit that consumers are more price-sensitive when prices increase. Consequently, knowing this purchasing behavior, retailers will be less inclined to decreasing the price following a cost decrease than increasing the price following a cost increase, which leads to asymmetric pricing.

Another explanation between competition and asymmetric pricing is the Edgeworth cycles. Noel (2007a, 2007b), using data from the Canadian gasoline market, shows that isolated and less competitive markets often exhibit a long period of price stability, while more competitive local markets tend to have short-lived price cycles.

Petrikaite (2016) shows how more information available to consumers about fuel prices facilitates switching sellers, and how high consumer search costs may result in higher gasoline market prices. In another example, Byrne and De Roos (2019) document that gasoline retailers in Australia used the price transparency program called FuelWatch to initiate and sustain cooperation. They argue that firms can undertake unilateral price announcements declaring a pricing scheme's elements, and that these elements have to comprise leadership and price match. A mutual understanding in an oligopoly might also be achieved following simple attempts to raise prices in the market, although this might not always be the case. If a firm raises prices and rivals match the new price, then the firms reach a mutual understanding and, from that moment on, they may be able to sustain tacit collusion. Pricing schemes must respond to market outcomes quickly, as they need to provide the best response in the short-term while still trying to increase profits in the long run.

Finally, two recent studies (Gautier and Le Saout, 2017; Benoit et al., 2019) have inves-

tigated whether the degree of spatial competition is a driver of price stickiness in the French retail gasoline market. Benoit et al. (2019), using daily prices for over 8000 French gas stations, find that the intensity of local competition reduces the rigidity of retail gasoline prices. They also observe that the degree of price stickiness of gas stations strongly depends on the price-setting behavior of their local rivals.

Of particular note is the fact that the intensity of gasoline competition at the station level is primarily concentrated at a specific geographical distance between a given gas station and its rivals. Since pump prices are fully transparent, both consumers and retailers observe the final prices daily and adjust their purchasing/business strategies accordingly. Therefore, gas stations are spatially dependent on their daily effort to set the level of pump prices.

Within this line of research, the current paper revisits the relationship between competition and firms' margin levels by focusing on the Swedish retail gasoline market. More specifically, this paper adds to the rich body of literature by providing evidence on the relationship between the degree of local competition as measured by station proximity to rivals and real gross margins at the station level.

1.3 Data description

We focus our research on the Swedish gasoline market, containing daily station-level microdata for five oil brands across six different local markets. All stations part of the sample are vertically integrated because the oil majors own them and directly control day-to-day pricing. In this respect, the Swedish gasoline market is similar to the Norwegian gasoline market (Foros and Steen, 2013), but distinctly different from the market in many other countries. For example, vertical integration is significantly less common in the US gasoline market (Borenstein et al., 1993).

Three major oil companies are part of our sample: Statoil Fuel & Retail AB, St1 Energy AB, and OK-Q8 AB. The first two own brands that run service (Statoil and Shell) and self-service (Jet and St1) stations. In 2012, Statoil Fuel & Retail AB was the largest retailer in the market, having a market share of 34% in gasoline volume. OK-Q8 AB and St1 Energy AB were the second and third largest retailers with 27.9% and 22.6% market share in gasoline volume, respectively.

We considered the three service and two self-service brands separately. We utilized a novel dataset containing the daily station-specific observations of gasoline volume, Rotterdam spot market price, and pump prices. This dataset is obtained from the Swedish Competition Authority covering the period from January 1, 2012 to December 31, 2012. In total, we have a balanced dataset of 366 liters in daily retail gasoline volume and price observations across 147 stations in 6 different geographical markets, that is, Stockholm, Gothenburg, Malmö, the E6 highway, smaller cities, and rural areas. Volumes are in terms of a liter and prices are in terms of Swedish kroner per liter (SEK/liter). The retail price has two components: fiscal and market. Taxes account for approximately 63% (energy tax + CO2 tax + VAT) of the

final retail price in Sweden. The Rotterdam spot market price is an input price that can be considered the price of buying gasoline on the international market. The primary cost source for oil firms is this price, as we do not observe firm-specific costs. Thus, firms are subject to the same input price. Given the richness of the data, we built a variable representing the firms' profitability, in this case, real gross margin, which is defined as the difference between the volume-adjusted price and gasoline costs. More importantly, each volume and price value is linked to a station, brand, and region that allows us to examine a brand's profitability at the national, regional, and station levels. The daily volume and price data for each station are complementary because they allow us to introduce the volume-adjusted prices in the literature and examine the real gross margin patterns. Furthermore, we considered three different price measures. 1) The daily volume-adjusted gasoline price, which is the everyday price at the pump (including taxes), adjusted by the volume sold on that day at each station. 2) The daily Rotterdam spot market price for refined products as a measure of input price. 3) We define brands' real gross margins as the volume-adjusted prices, which are less than the input cost and taxes.

1.4 Empirical analysis

The following subsections will discuss the empirical evidence separately for the service and self-service gasoline brands.

1.4.1 Service gasoline brand results

1.4.1.1 Real gross margin patterns at the national level

Our dataset has some important implications for studying firms' dynamic profit strategies. The data's high frequency implies that: 1) there is retail competition daily, 2) retailers set gross margins simultaneously each day, and 3) retailers face standard daily input cost and tax fluctuations over time. Typically, integrated players in this retail gasoline market have a very dense retail network covering the entire country. All stations of the retail network of the major oil companies in our sample are under their ownership. As such, the owners of the major oil companies are free to set the retail price. Besides, the owners can perfectly observe each other's prices each day. Our descriptive framework maps well into data-driven approaches and reveals important market issues by just observing the data. The daily real gross margins of brands are calculated at country, regional, and station level using the following formula, and consumer daily volume-adjusted price for gasoline consists of the following components:

$$P_t = PO_t + T_t + RGM_t * VAT_t \tag{1}$$

where P_t is the daily volume-adjusted price for gasoline at day t, PO_t is the input price at day t, T_t is the tax (energy tax + CO_2 tax), VAT_t is the value-added at day t and RGM_t is the real gross margin. Further, formula (1) assumes that PO_t and T_t were paid on the

same day.[1] Therefore, we should be aware that when the input prices climb for a certain period, it underestimates their gross margins, and vice-versa when the input prices decline. Of these components, tax constituted approximately 37% of the sales price, product cost 33%, gross margin roughly 5%, and VAT 25% of the retail price in December 2012.[2]. The state sets both the energy tax and the carbon tax every year. The input cost considered in our analysis comprises the cost of crude oil but does not include the transport, handling, and storage costs. Based on the information we received from the managers of the transportation and logistics department of the oil companies, the transportation cost was relatively fixed during the sample period due to the long-term contacts for transport/distribution. About 1–2% of the retail price, including taxes, reflected the transport cost, which was set to be the same across Sweden.

As a next step, we have considered the proximity of stations to measure the degree of competition across regions. Table 1 below shows the average station distance from the closest competitor across regions.

Table 1: The distance to the closest rival station across regions (in km)

Regions	Stockholm	Gothenburg	Malmö	E6	Smaller cities	Rural areas
Mean	1.1	0.9	0.9	2.5	0.616	4.7
Standard deviation	0.2	0.1	0.1	0.5	0.1	1.4

As anticipated, the distance between a station and its nearest competitor is larger in rural areas compared to other regions. Therefore, we have considered the real gross margin values in rural areas as reflecting the profits of isolated markets.

The input cost is given on each day, and it is fixed for every company. $PO_t + T_t$ is the price that a retailer has to pay to get the product at a station and then sell it. The only difference here is that T_t was constant throughout the year while PO_t is paid on a different day. Given the price formula (1), consumers will carry the tax's financial burden, increasing the input cost. Therefore, changes in RGM_t represent firms' price adaptation throughout the year and also firms' dynamic profit strategies. RGM_t in all figures that follow are expressed in SEK/liter sold. Figure 1 shows the national distribution of real gross margins of brands belonging to three service brands that run only service stations. Each color represents an oil company in our analysis. Specifically, red for Statoil Fuel & Retail AB, green for St1 Energy AB, and blue for OK-Q8 AB, unless otherwise indicated.

[1] In the literature, it is common to assume that the option value of gasoline today measured through the Rotterdam price is the daily cost.

[2] https://drivkraftsverige.se

Figure 1: The distribution of service brands' real gross margins

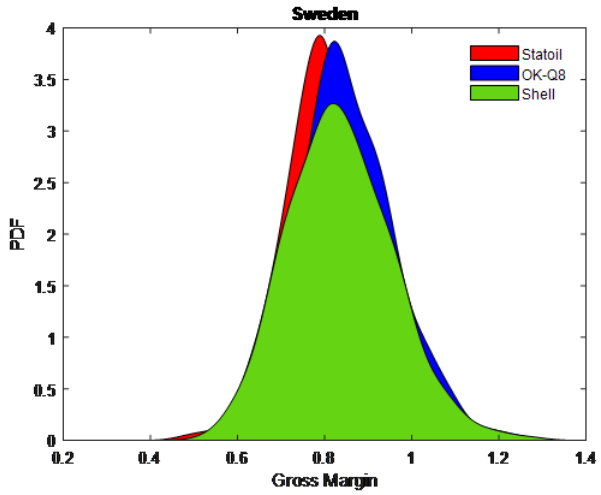

The x-axis contains the margin values, and the y-axis is the probability density function for the continuous random variable x. As can be seen, OK-Q8 charges higher real gross margin values nationally than Shell and Statoil, with an annual mean of 0.85 SEK/liter, 0.83 SEK/liter, and 0.8 SEK/liter, respectively. However, Shell's profits exhibit higher variations compared to Statoil and OK-Q8.

Next, to examine the influence of the past margin decisions on current margin values, we have calculated the average of the historical margin data for each day. Figure 2 depicts these values for the service brands nationally.

Figure 2: The daily average of the historical margin data

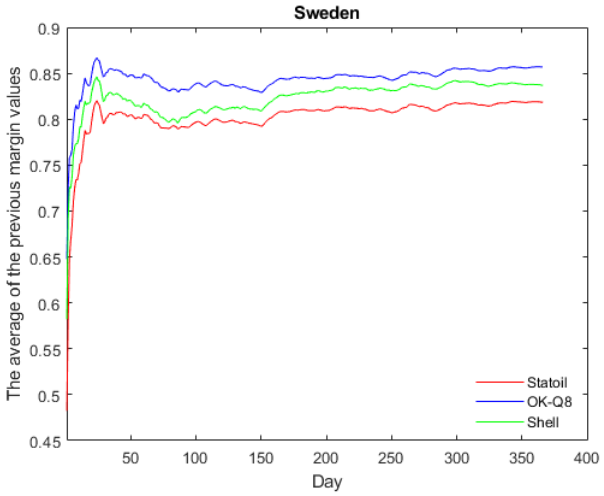

The vertical axis shows the daily average of the historical margin data, and the horizontal axis shows the days of the year. As shown in Figure 2, the mean of historical data exhibits some variations at the beginning of the year due to the lack of the previous year's data. Nevertheless, the mean is still relatively stable throughout the year. It seems like each brand

targets a fairly fixed mean value and maintains it over time. Surprisingly, their mean margins do not cross but move synchronously. In addition, Figure 2 shows that OK-Q8 targeted the highest margin mean nationally and Statoil the lowest. At the same time, Shell appears to be the middle brand in terms of margins during the entire year. This fact implies a clear margin targeting structure over time.

1.4.1.2 Real gross margin patterns at the regional level

Increased concern about gasoline pricing has increased the interest in how gasoline prices are determined and how they change in different local markets. Figure 3 shows the service brands' real gross margin distributions across six regions, including urban and rural markets in Sweden.

Figure 3: The distribution of brands' margin across regions

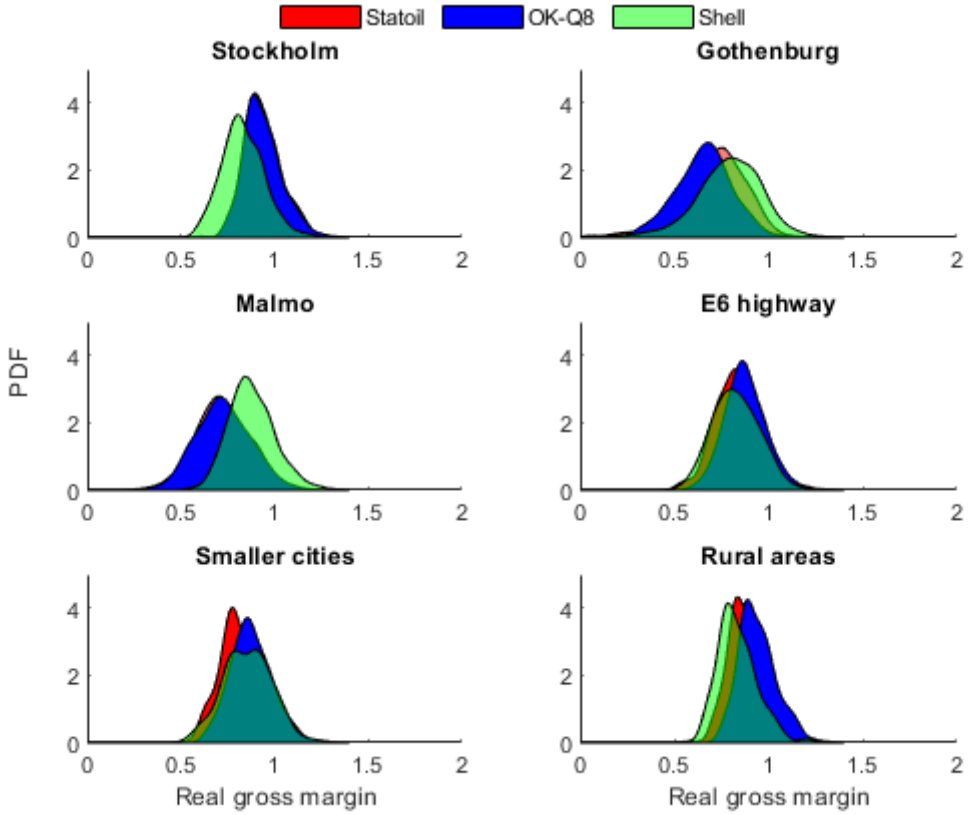

Again, the vertical axis contains the probability density function and the horizontal axis indicates the margin values. In comparison to Figure 1, Figure 3 presents a detailed picture of service brands' margin levels across local markets. Statoil and OK-Q8 are closely leading the market in terms of margins in Stockholm, while Shell sets seemingly lower margins than the other two brands in Stockholm. However, Shell is the market leader in terms of margins in Gothenburg and Malmö, followed by Statoil and OK-Q8.

Interestingly, Statoil and OK-Q8 have almost the same margin distributions with almost

identical margin means and standard deviations in Stockholm and Malmö. OK-Q8 charges higher margins than the other two brands on the E6 highway and in rural areas. While in smaller cities, OK-Q8 and Shell seem to set higher margins, although the spread of Shell's margins is more significant than that of OK-Q8. Shell has lower gross margins than its competitors in Stockholm, but it seems to have significantly higher margins in Gothenburg and Malmö. This finding is interesting because it may imply that these brands have different market positions in terms of margins across local markets.

Therefore, we calculated the brands' annual mean margin across regions. Table 2 shows how service gasoline brands are positioned across local markets. The positions highest, middle, and lowest is given based on the annual mean margin of the brands.

Table 2: Brands' position in terms of margins across six regions in Sweden

Regions	Stockholm	Gothenburg	Malmö	E6	Smaller cities	Rural areas
Highest brand	Statoil	Shell	Shell	OKQ8	OKQ8	OKQ8
Middle-brand	OKQ8	Statoil	Statoil	Shell	Shell	Statoil
Lowest brand	Shell	OKQ8	OKQ8	Statoil	Statoil	Shell

OK-Q8 has the highest mean margin in three regions, lowest in two, and medium in one region. Statoil is the leader in one region, the middle brand in three regions, and the lowest on the E6 highway and in smaller cities. In contrast, Shell seems to be the leader in two regions, the middle brand in two, and the lowest in two. Interestingly, all brands are the lowest in two local markets.

To better understand how these brands compete in those six different local markets, we have computed the daily moving average of the previous brands' gross margins at the regional level (Figure 4).

Figure 4: The daily average of the historical margin data across regions

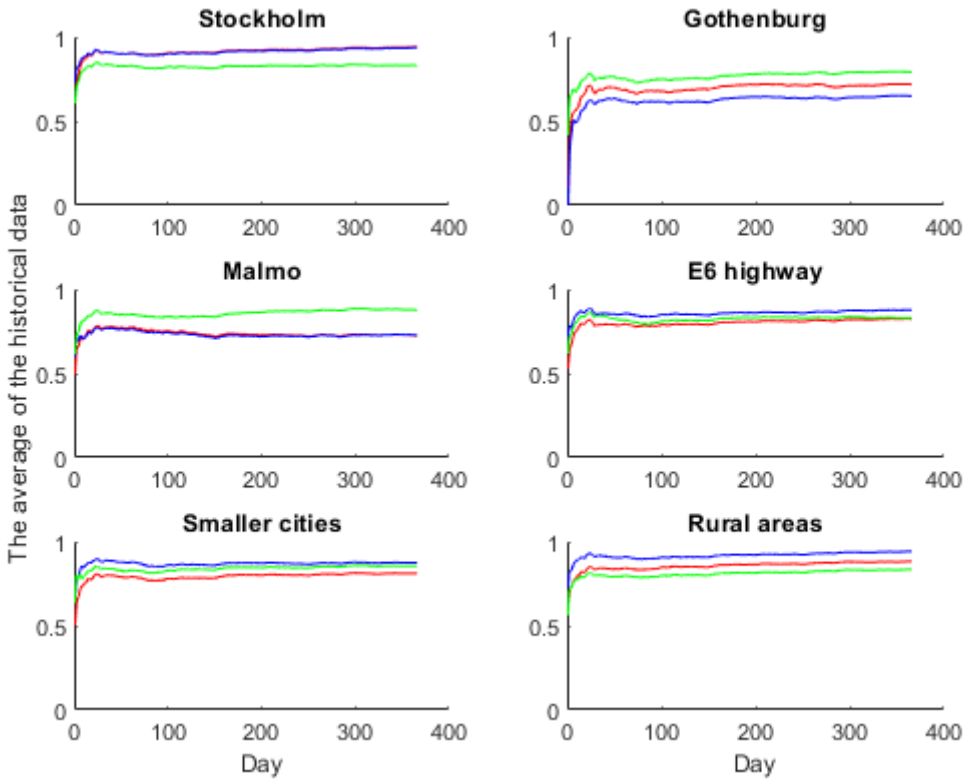

With the same reasoning as at the national level, the daily average of the historical regional margin values exhibits some variations at the beginning of the year. Again, the mean of the historical values seems to be relatively stable throughout the year in each region. Statoil and OK-Q8 closely follow each other in Stockholm, with Statoil charging a slightly higher margin mean than OK-Q8. Shell is the lowest in Stockholm. In line with previous results, we see the opposite setting in Gothenburg and Malmö, where Shell is the highest and Statoil the middle brand. In contrast, in smaller cities, rural areas, and on the E6 highway, OK-Q8 is the leader.

Each brand is the lowest brand in two regions, and this fact is also evident from Table 2. Statoil is the lowest brand on the E6 highway and in smaller cities, Shell in Stockholm and rural areas, and OK-Q8 in Gothenburg and Malmö. This fact may suggest a shared market. Another finding that Figure 4 reveals is the increased slope of the brand-running mean during the second half of the year in each region except in Malmö, where Statoil and OK-Q8 have a downward slope. Figures 2 and 4 also reveal the importance of the previous brands' margin data in the current margin decision. Given that the oil price may be random, it suggests that the adjustments may follow a certain rule. In other words, all brands try to target a fixed daily margin, but they know that the oil price is changing every day, so they have to adjust the price continuously to approach this target margin.

Moreover, the fact that brands' margins do not undercut throughout the year, either at the national or regional levels, suggests a structure in this market. In addition, we observe a higher margin mean in Stockholm and rural areas than other big cities and smaller cities. This might imply that the competition conditions in Stockholm are comparable with the competition conditions in rural markets. The latter is considered in the literature as a market that reflects monopolistic market outcomes while the former reflects competitive market features.

1.4.1.3 Real gross margin patterns at the station level

As above, we have calculated the station mean of historical margins; Figure 5 shows these values for three regions: Stockholm, rural areas, and Malmö. We have chosen these three regions to represent high margin regions (Stockholm and rural areas) and low margin regions (Malmö).

Figure 5: The mean of the historical margin values across stations and regions

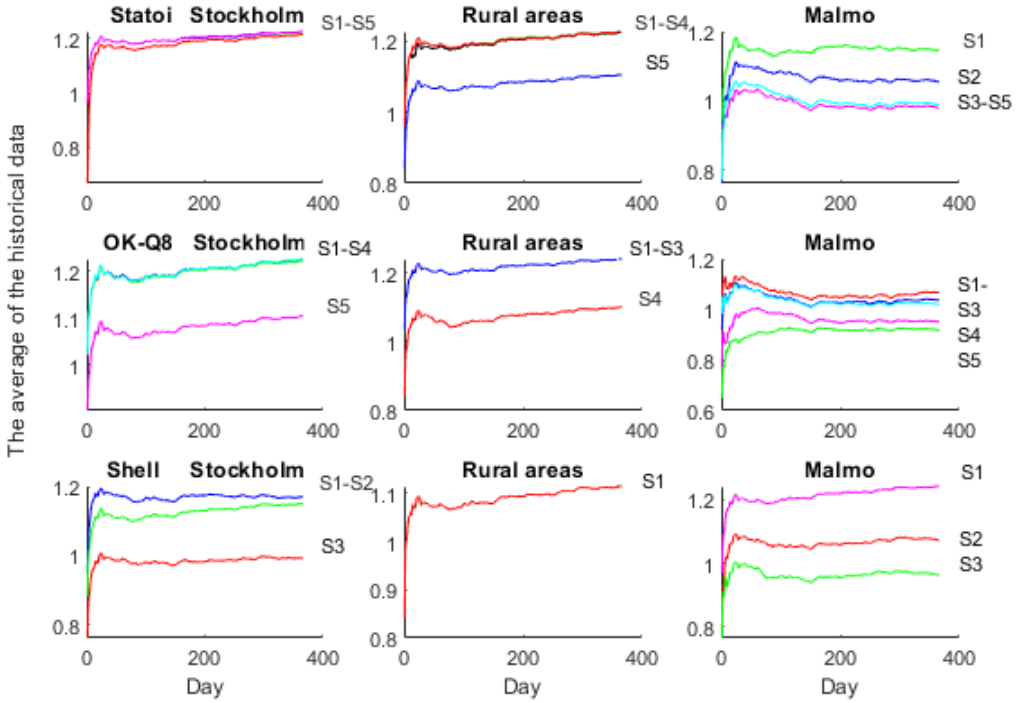

As can be seen, the spread in Malmö is much larger than in Stockholm and rural areas. This fact supports the results above, which suggest a more competitive local market in Malmö as compared to Stockholm. Similarly, as in the figures above, we observe a relatively stable mean in the station level data. It seems that each station targets a relatively fixed mean value and tries to keep it over time with slight variations of the targeted value.

1.4.2 Self-service gasoline brands' results

1.4.2.1 The gross margin patterns at the national level

We found it important to examine the self-service retail brands as there is a dearth of research focusing on these brands. Therefore, observing their real gross margins and comparing them with the service brand margins gives us insights into the market structure differences and competition conditions between (self) service brands. First, we analyzed the daily real gross margins of self-service brands at the country level; Figure 6 shows the distribution of the margins for each brand. The red indicates the brand owned by Statoil Fuel & Retail AB and green reflects the brand owned by St1 Energy AB.

Figure 6: Distribution of the real gross margins of brands

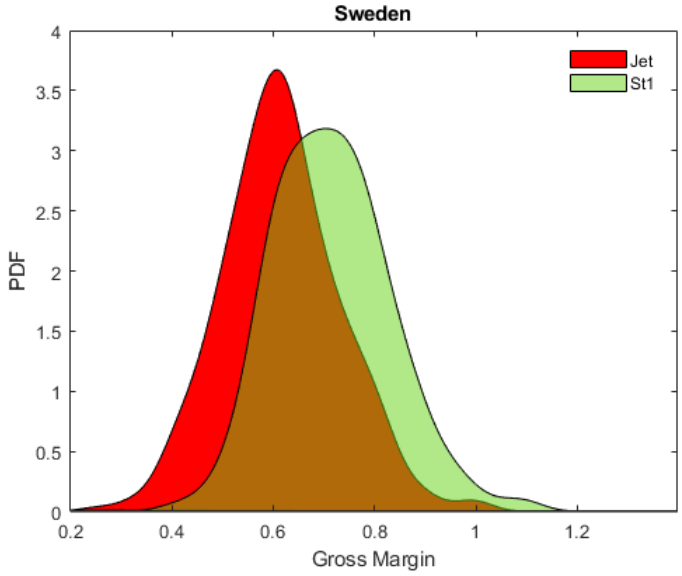

The x-axis contains the gross margin values of the self-service brands nationally, and the y-axis their corresponding probability density functions. St1's annual margin mean is seemingly bigger and exhibits a larger variation compared to that of Jet. The majority of Jet's margin values range from 0.4 SEK/liter to 0.8 SEK/liter with an annual mean of 0.61 SEK/liter, while the real gross margins of St1 range from approximately 0.4 SEK/liter to 1 SEK/liter with an annual mean of 0.71 SEK/liter. Thus, a 10 öre/liter difference exists between the two self-service brand margins on average. Figure 7 below compares the daily real gross margins between the two brands. The blue line represents the case where we assume that St1's margins are equal to Jet's margins. While, the red dots are the real combinations of their daily gross margins taken from Figure 6.

Figure 7: A comparison between St1's and Jet's margins

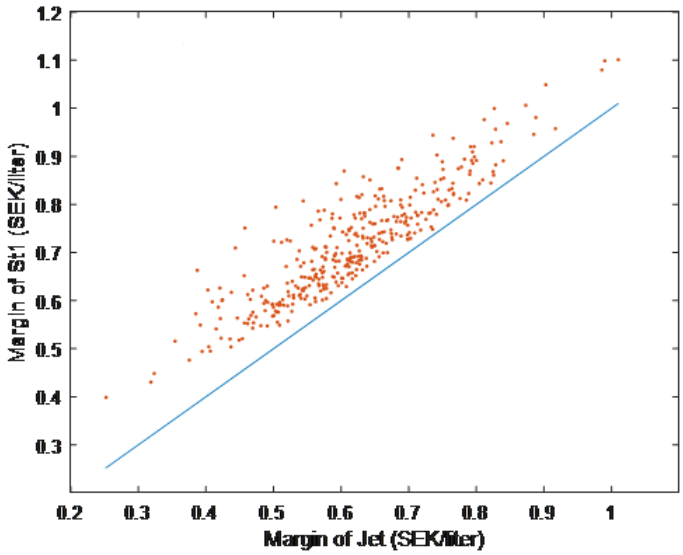

Although we noticed from Figure 6 that St1, on average, charges higher margins, Figure 7 shows that St1 always captures higher daily margins than Jet, and their margins never cross. The difference is always above the diagonal. However, the service brand margins fluctuate, as shown in Figure A1 in the Appendix. This probably simply suggests that independent of wholesale gasoline price fluctuations, Jet will always charge lower margins than St1. On the one hand, this fact can imply that St1 cannot afford to lower its margins down to Jet's margin level due to other costs that we do not have information on in our analysis.

On the other hand, this fact may suggest less competition. Brands that run self-service stations do not compete in the same manner as the service brands nationally. Indeed, Figure 7 is a strong picture if we combine it with Figure A1. Considering our visual inspection of the national margins for two different categories of gasoline brands: service and self-service, we choose to display the same graph as in Figure A1 for both categories of gasoline brands in Stockholm and rural areas. Figure A2 in the Appendix shows the results. St1's and Jet's margin values are much closer in Stockholm than in rural areas. Interestingly, Jet captures the highest margin in Stockholm, compared to other regions. This seems to imply that Jet is more competitive in Stockholm. However, does that also imply that it is not important for Jet to be cheaper in Stockholm due to the lack of fierce competition? In other words, knowing that service stations have higher prices, Jet tries to capture as much profit as possible in Stockholm. In rural districts, however, we observe a clear market leader in terms of margins: St1 in the self-service graph and OK-Q8 in the service graph.

Next, we calculated the daily mean of historical margin data of brands that run only self-service stations to examine how their past margins influence the current margin outcomes. Figure 8 shows the mean for Jet and St1 at the national level.

Figure 8: The daily average of the historical margin data

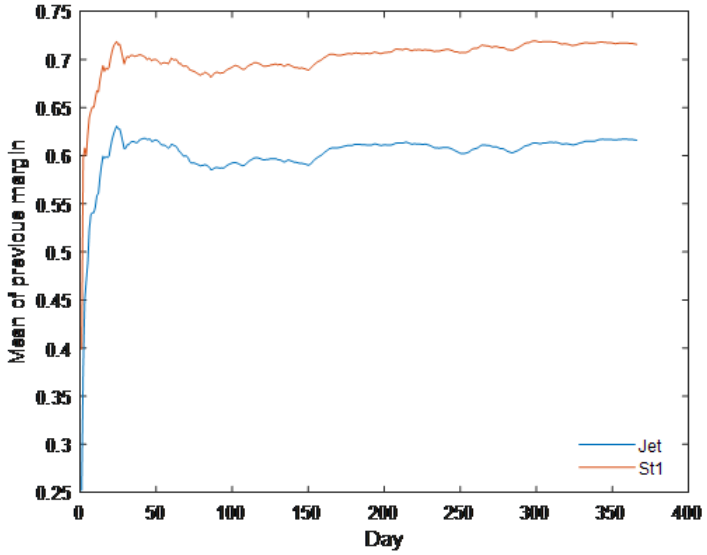

In addition, the annual average margin value that St1 and Jet seem to target can be seen in Figure 8. The variations at the beginning of the year can be due to the unavailability of the previous data. Nevertheless, approximately after January, the target seems obvious for Jet and St1 with SEK 0.6 and SEK 0.7, respectively, and each brand seems to belong somewhere above and below that targeted profit value.

1.4.2.2 Gross margin patterns at the regional level

Again, as for gasoline brands that offer full service at the station, we have observed the distribution of the daily margins of self-service brands across different local markets, such as Stockholm, Gothenburg, Malmö, the E6 highway, smaller cities, and rural areas. Figure 9 shows the distribution profiles.

Figure 9: The distribution of brand margins across regions

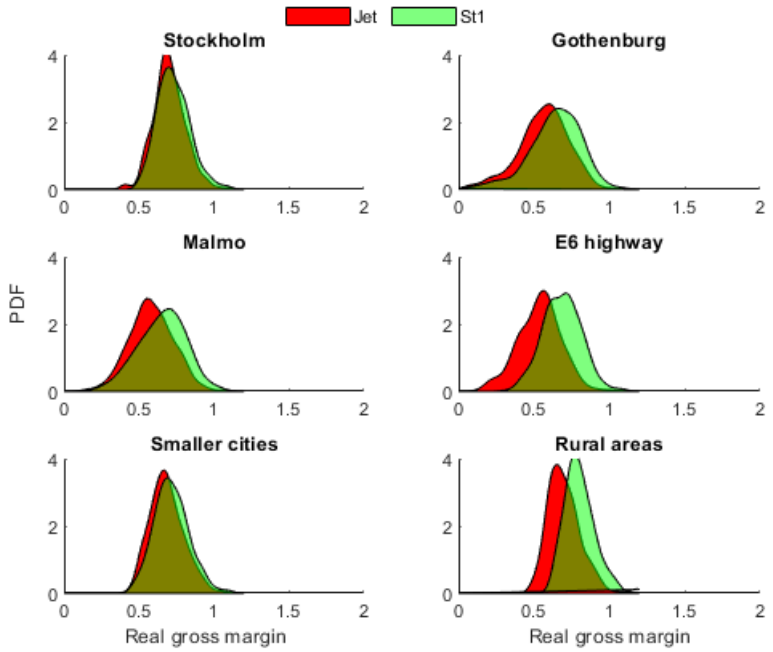

We see that St1 has a more significant mean margin across all areas. Except in Stockholm and smaller cities, we observe a ten öre/liter difference between their margins. As a next step, we have calculated each brand's daily mean of the historical margin value across regions, and Figure 10 depicts these values.

Figure 10: The daily average of the historical margin data across regions

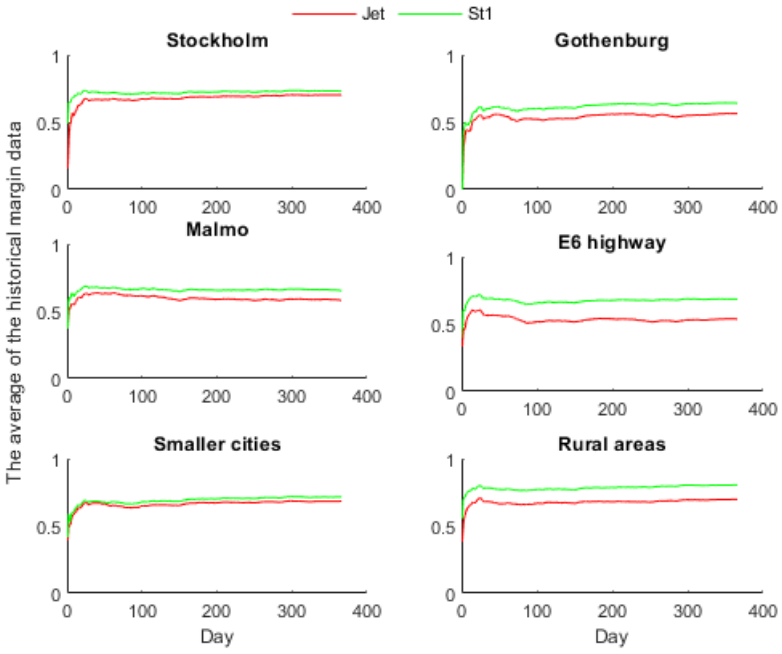

Supporting what we saw from the previous figures, the mean of St1 is seemingly higher than Jet's mean in each region. The differences are smaller in Stockholm and smaller cities but significantly bigger in Malmö, rural areas, Gothenburg, and on the E6 highway. The data shows a precise self-service brand positioning across local gasoline markets. St1 is the most expensive brand, and Jet is the cheapest brand (in daily values), regardless of their stations. We are uncertain about making some conclusions if these results suggest that St1 might enjoy some market power or that these two self-service brands seem to not actually be competing.

1.4.2.3 Gross margin patterns at the station level

For each station of St1 and Jet located in Stockholm, rural areas, and Malmö from our sample, we have calculated the daily average of the previous gross margins to analyze how important the past is on today's station profit decisions. Figure 11 shows that the spread of gross margins between Jet's stations is more significant in Stockholm than in Malmö. The same pattern can be observed for St1, indicating an opposite pattern of what we observed for the service brands (Figure 5). This means that different from the service stations, the self-service stations have a slightly more significant margin spread in regions with high gross margin values than in those with low margin values.

Figure 11: The daily average of the historical margin data across stations and regions

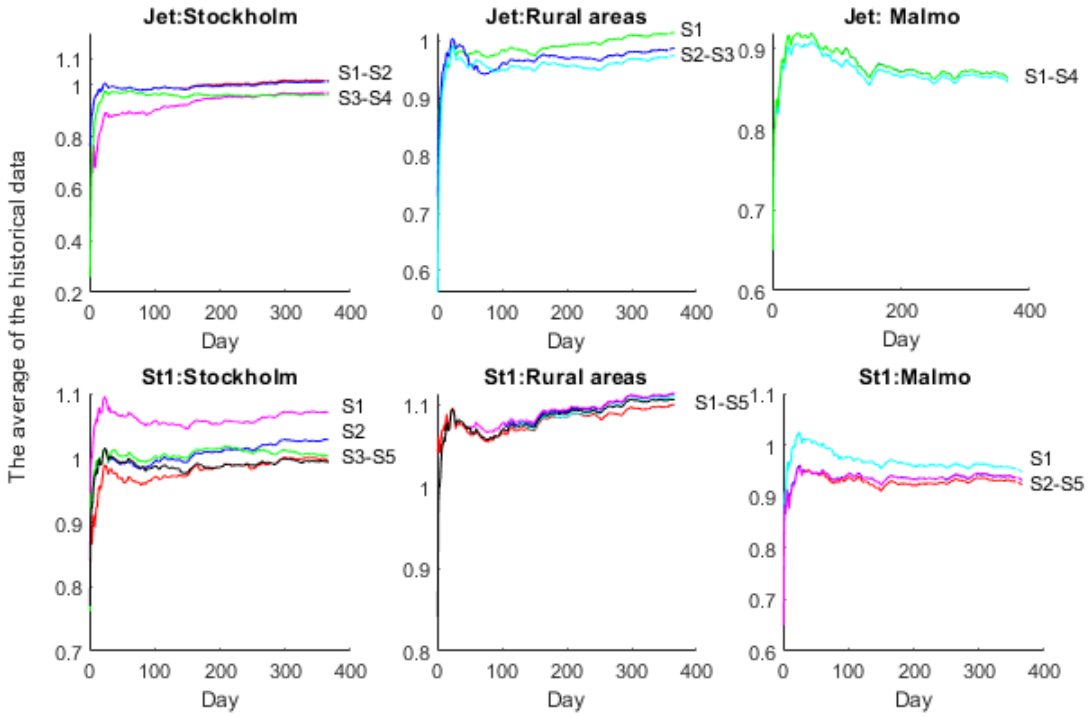

The mean of historical data for both self-service brands has a downward slope in Malmö

and an upward slope in Stockholm. Gasoline stations that belong to Jet charge slightly lower gross margins than the ones that belong to St1.

1.5 Discussion

This study observed oil companies' profit decision-making in the Swedish retail gasoline market and investigated how real gross margin values adjust as firms' responses to cost change.

We started by separately analyzing service and self-service gasoline brands and looking at their margin strategies to investigate if it can be possible that brand margins enhance, notwithstanding the evidence of regional differences in margins. We observe that all brands capture higher margins in Stockholm and rural areas and that such an increase in margins might be due to the absence of fierce competition in Stockholm. Oil brands that run only service stations seem to not compete as aggressively in Stockholm as they do in Gothenburg or Malmö. Abrantes-Metz et al. (2006), using daily price data from the retail gasoline industry in Louisville, argue that markets exhibiting high mean prices and low variance suggest the presence of conspiracies. Referring to their paper, we have calculated the mean and variance of volume-adjusted prices across regions, and Figure 12 depicts the results observed.

Figure 12: Price means and standard deviations across regions

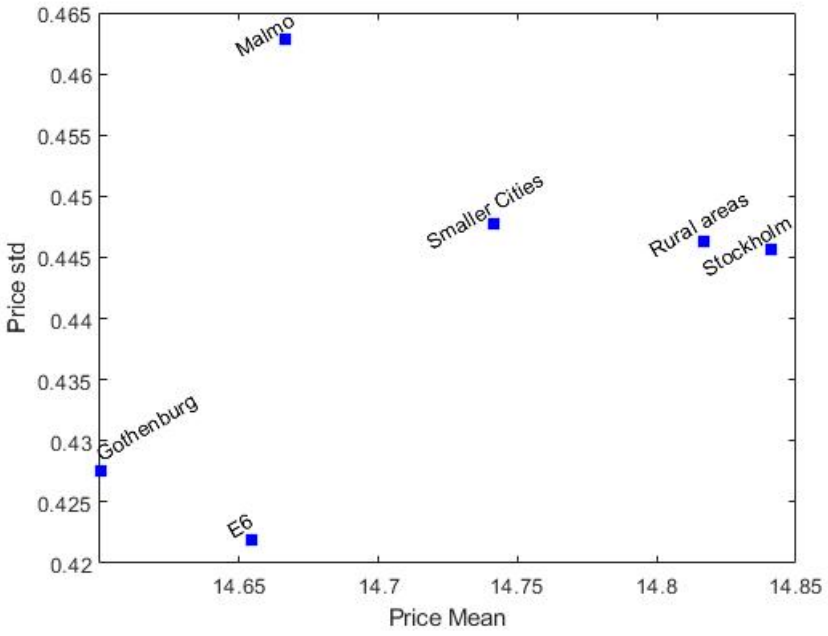

As apparent from Figure 12, the so-called "outliers" (Abrantes-Metz et al., 2006) in our case seem to be Stockholm and rural areas, which have high means and relatively low variances. Gothenburg and the E6 highway have a low mean and standard deviation. In

contrast, Malmö and smaller cities exhibit a lower price mean and a higher price variation than other regions. Table A1 placed in the Appendix contains the mean price t-tests across regions. In this case, we test whether the prices are significantly different in Stockholm compared to rural areas and other regions. We observe that the difference between the mean prices in Stockholm and rural districts is not statistically different. In contrast, the difference in mean prices in those regions is statistically different from that of the other regions, further strengthening the observed results on the similarities between Stockholm and rural areas.

A reasonable assumption is that Stockholm seems to be a less competitive market based on the observed monopolization pattern. All brands target a higher annual margin mean in Stockholm, and the difference between service and self-service brand margins gets smaller. Among service brands, OK-Q8 charges the highest margin countrywide and St1 charges the highest among self-service brands in our sample. The margin variations are more considerable in low-profit regions than in high-profit regions for both cases. Surprisingly, the margin level in Stockholm is significantly different from the other big cities in Sweden, resembling the rural district's margin levels. These results suggest less competition in Stockholm than in Gothenburg and Malmö. The price difference is much more evident between self-service brands, where St1's national margins are always higher than Jet's national margins, and they never crossed during the year. On average, Jet is cheaper than St1 across all regions, but their margin difference is much smaller in Stockholm compared to other areas. In contrast, service brands shift their roles in terms of margins across regions. The relatively stable moving average margin patterns observed across different local markets reveal the importance of the previous brand profit decision in their day-to-day decisions.

Moreover, the evident steady mean gross margin target seems to be well adapted by each brand nationally and locally. As expected, the real gross margins of service brands are seemingly higher than the real gross margins of self-service brands. This difference may arise from various reasons, such as the capitalization of providing a higher service, labor cost, covering the convenience store costs, or earning additional profits.

1.6 Conclusion

At the national level, we observe that both self-service and service oil brands seem to target a seemingly rigid mean margin over time. Significantly, the regional level analysis results show that the highest brands' real gross margin areas are Stockholm and rural areas. However, the data suggest higher competition in Gothenburg and Malmö. Results in Stockholm significantly differ from other big cities in Sweden.

Additionally, a mixed, leading role in terms of margins across different local markets is observed for the service brands. Statoil, OK-Q8, and Shell have lower margins in two regions out of six. In contrast, the data show a different reality for self-service brands. St1 always gets higher margins than Jet, independent of the station's geographical location. Service brands have on average a ten Øre/liter higher gross margin than self-service brands and

seem to compete more than the self-service brands. The stable mean margin throughout the year may reflect a lack of effective competition in Stockholm and isolated markets (rural areas). Thus, price-sensitive drivers that compare prices might indeed not exert competitive pressure on the gasoline brands in those two local markets.

However, due to the high market transparency and the technological developments, no explicit communication between companies is needed to cooperate on a strategy of fixed target profit that satisfies all market participants. During our sample period, the centralized price setting in the Swedish retail gasoline market may have contributed to a common profit strategy on rigidity in gasoline markups where major companies aim to maximize common profit for their self-service and service stations.

The analysis at the station level across regions supports our findings at the regional level. The variations between stations' margins in Stockholm and rural areas are significantly minor than in Malmö.

However, the surprising regional margin patterns and the coexistence of sticky daily mean margin observed in Stockholm should be developed further, withs special focus on the role of price asymmetry or implicit coordination.

References

[1] Abrantes-Metz, R. M. and Froeb, L. M. and Geweke, J. and Taylor, Ch. T., (2006), "A variance screen for collusion, " *International Journal of Industrial Organization*, Vol.24, iss.3, pp. 467-486.

[2] Alvarez, L. J. and Hernando, I., (2006), "Price setting behavior in Spain: Evidence from consumer price micro-data, " *Economic Modelling*, Vol.23, iss.4, pp. 699-716.

[3] Balmaceda, F. and Soruco, P., (2008), "Asymmetric dynamic pricing in a local gasoline retail market, " *The Journal of Industrial Economics*, Vol.56, iss.3, pp. 629-653.

[4] Borenstein, S. and Shepard, A., (1993), "Dynamic pricing in retail gasoline markets, " *National Bureau of Economic Research.*

[5] Borenstein, S. and Cameron, A. C. and Gilbert, R., (1997), "Do gasoline prices respond asymmetrically to crude oil price changes?, " *The Quarterly journal of economics*, Vol. 112, iss.1, pp. 305-339.

[6] Byrne, D. P., and De Roos, N., (2019), "Learning to coordinate: A study in retail gasoline, " *American Economic Review*, Vol.109, iss.2, pp. 591-619.

[7] Byrne, D. P., and De Roos, N., (2017), "Consumer search in retail gasoline markets, " *The Journal of Industrial Economics*, Vol.65, iss.1, pp. 183-193.

[8] Cabral, L. and Fishman, A., (2012), "Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment, " *International Journal of Industrial Organization*, Vol.30, iss.4, pp. 371-376.

[9] Cornille, D. and Dossche, M., (2006), "The patterns and determinants of price setting in the Belgian industry, " *National Bank of Belgium Working Paper*, Vol.82.

[10] Connor. J. M., (2005), "Collusion and Price Dispersion, " *Applied Economic Letters*, Vol.12, iss.6, pp. 225-228.

[11] De Roos, N. and Smirnov, V., (2020), "Collusion with intertemporal price dispersion, " *The RAND Journal of Economics*, Vol.51, iss.1, pp. 158-188.

[12] De Roos, N., (2004), "A model of collusion timing, " *International Journal of Industrial Organization*, Vol.22, iss.3, pp. 351-387.

[13] Eckert, A., (2013), "Empirical studies of gasoline retailing: A guide to the literature, " *Journal of Economic Surveys*, Vol.27, iss.1, pp. 140-166.

[14] Encaoua, D.and Geroski, P. and Jacquemin, A., (1986), "Strategic competition and the persistence of dominant firms: A survey, pp. 55-89.

[15] Gautier, E. and Ronan L. S., (2015), "The dynamics of gasoline prices: Evidence from daily French micro data, " *Journal of Money, Credit and Banking*, Vol.47, iss.6, pp. 1063-1089.

[16] Götz, G. and Gugler, K., (2006), "Market concentration and product variety under spatial competition: Evidence from retail gasoline, " *Journal of Industry, Competition and Trade*, Vol.6, iss.3-4, pp. 225-234.

[17] Green, E. J. and Porter, R. H., (1984), "Non-cooperative collusion under imperfect price information, " *Econometrica: Journal of the Econometric Society*, pp.87-100.

[18] Hastings, J. S., (2004), "Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in Southern California, " *American Economic Review*, Vol.94, iss.1, pp. 317-328.

[19] Hastings, J. S. and Shapiro, J. M., (2013), "Fungibility and consumer choice: Evidence from commodity price shocks, " *The quarterly journal of economics*, Vol.128, iss.4, pp. 1449-1498.

[20] Foros, Ø. and Nguyen-Ones, M.and Steen, F., (2021), "The effects of a day off from retail price competition: evidence on consumer behavior and firm performance in gasoline retailing, " *International Journal of the Economics of Business*, pp. 1-39.

[21] Foros, Ø. and Nguyen-Ones, M., (2021), "Coordinate to obfuscate? The role of prior announcements of recommended prices, " *Economics Letters*, Vol.198, pp. 109680.

[22] Foros, Ø. and Steen, F., (2013), "Vertical control and price cycles in gasoline retailing, " *The Scandinavian Journal of Economics*, Vol.115, iss.3, pp. 640-661.

[23] Ivaldi, M. and Jullien, B. and Rey, P. and Seabright, P. and Tirole, J., (2003), "The economics of tacit collusion, " *IDEI Working Paper*.

[24] Levenstein, M. C. and Suslow, V. Y., (2006), "What determines cartel success?, " *Journal of economic literature*, Vol.44, iss.1, pp. 43-95.

[25] Lewis, M. and Noel, M., (2011), "The speed of gasoline price response in markets with and without Edgeworth cycles, " *Review of Economics and Statistics*, Vol.93. iss.2, pp. 672-682.

[26] Lewis, M. S., (2015), "Odd prices at retail gasoline stations: focal point pricing and tacit collusion, " *Journal of Economics and Management Strategy*, Vol.24, iss.3, pp. 664-685.

[27] Maskin, E. and Tirole, J., (1988), "A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles, " *Econometrica: Journal of the Econometric Society*, pp. 571-599.

[28] Nguyen-Ones, M.and Steen, F., (2019), "Market Power in Retail Gasoline Markets, " *NHH Dept. of Economics Discussion Paper*, vol. 21.

[29] Noel, M. D., (2007), "Edgeworth price cycles: Evidence from the Toronto retail gasoline market, " *The Journal of Industrial Economics*, Vol.55, iss.1, pp. 69-92.

[30] Noel, M. D., (2007), "Edgeworth price cycles, cost-based pricing, and sticky pricing in retail gasoline markets, " *The Review of Economics and Statistics*, Vol.89, iss.2, pp. 324-334.

[31] Petrikaite V., (2016), "Collusion with costly consumer search, " *International Journal of Industrial Organization*, Vol.44, pp. 1-10.

[32] Rotemberg, J. J. and Saloner, G., (1986), "A supergame-theoretic model of price wars during booms, " *The American economic review*, Vol.76, iss.3, pp. 390-407.

[33] Salvanes, K. G. and Steen, F. and Sørgard, L., (2003), "Collude, compete, or both? Deregulation in the Norwegian airline industry, " *Journal of Transport Economics and Policy (JTEP)*, Vol.37, iss.3, pp. 383-416.

[34] Sen, A., (2003), "Higher prices at Canadian gas pumps: international crude oil prices or local market concentration? An empirical investigation, " *Energy Economics*, Vol.25, iss.3, pp. 269-288.

[35] Steen, F. and Nguyen-Ones, M., (2019), "Market Power in Retail Gasoline Markets, " *NHH Dept. of Economics Discussion Paper*, Vol.21.

[36] Stiglitz, J. E., (1984), "Theories of wage rigidity, " *National Bureau of Economic Research*.

[37] Tappata, M., (2009), "Rockets and feathers: Understanding asymmetric pricing, " *The RAND Journal of Economics*, Vol.40, iss.4, pp. 673-687.

[38] Verlinda, J. A., (2008), "Do rockets rise faster and feathers fall slower in an atmosphere of local market power? Evidence from the retail gasoline market, " *The Journal of Industrial Economics*, Vol.56, iss.3, pp. 581-612.

[39] Young, D. P., (1997), "Dominant firms, price leadership, and the measurement of monopoly welfare losses, " *International Journal of Industrial Organization*, Vol.15, iss.5, pp. 533-547.

1.7 Appendix

Figure A1: The margins of service brands in Sweden

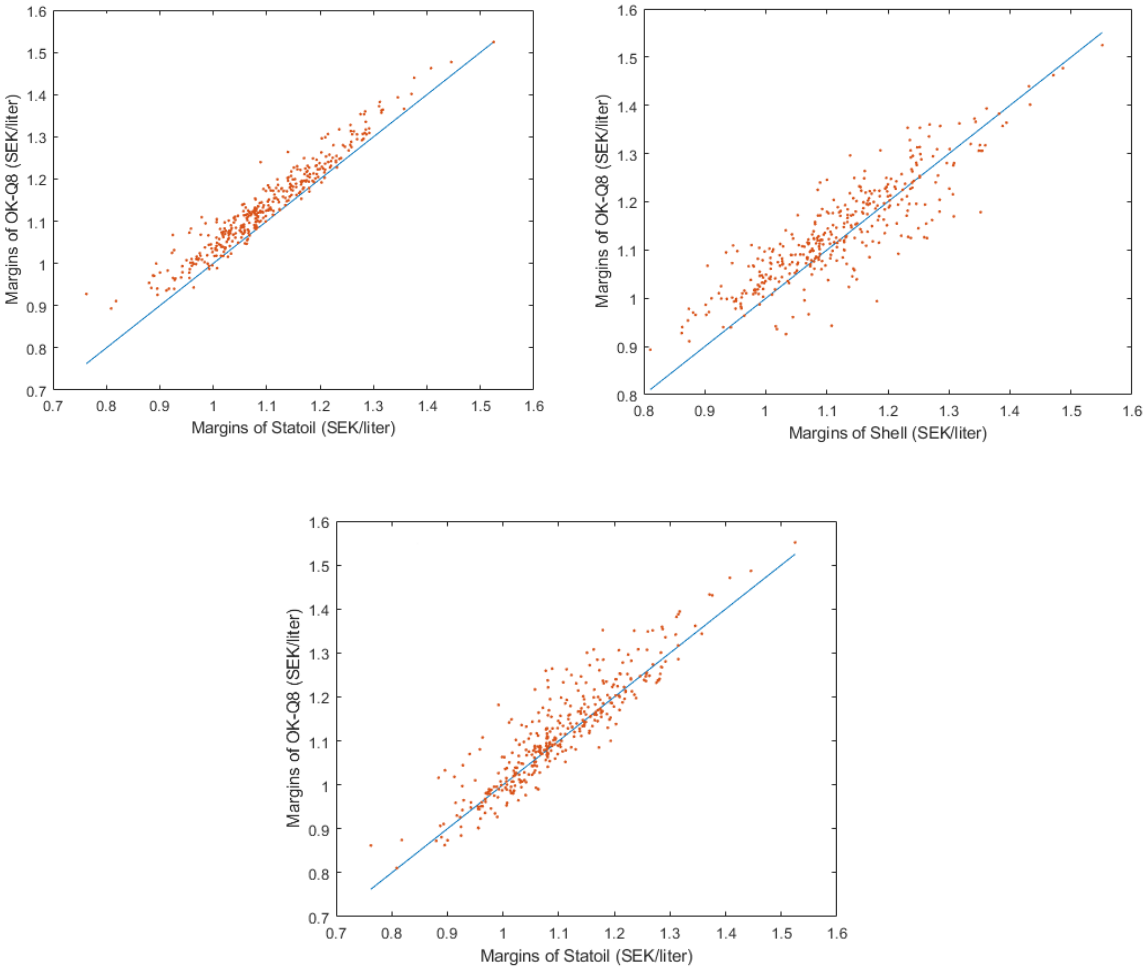

Figure A2: The brands' margins in Stockholm and rural areas

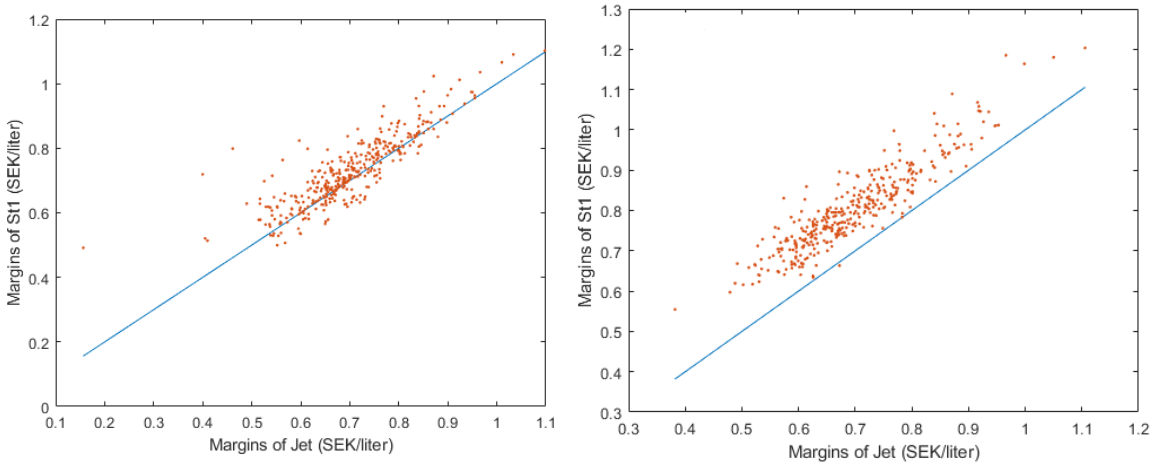

(a) Self-service brands

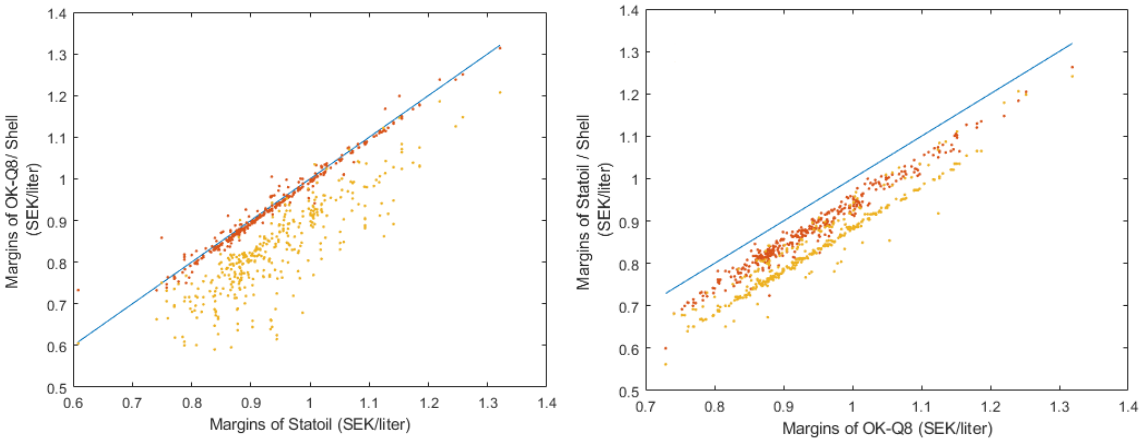

(b) Service brands

Table A1: The t-tests of the difference between mean prices across regions.

Regions	Mean Vol. weighted price	Std. Dev.	t-tests
Stockholm — rural areas	14.8407 — 14.8168	.4455 — .4462	0.7242
Stockhom — Gothenburg	14.8407 — 14.6004	.4455 — .4275	7.444***
Stockholm — Malmo	14.8407 — 14.6665	.4455 — .4627	5.189***
Stockholm — E6	14.8407 — 14.6543	.4455 — .4218	5.813***
Stockholm — Smaller Cities	14.8407 — 14.7412	.4455 — .4476	3.016***
Rural areas—Gothenburg	14.8169 — 14.6004	.4462 — .4275	6.699***
Rural areas — Malmo	14.8169 — 14.6665	.4462 — .4627	4.475***
Rural areas — E6	14.8169 — 14.6543	.4462 — .4218	5.064***
Rural areas — Smaller Cities	14.8169 — 14.7412	.4462 — .4476	2.291**
Gothenburg — Malmo	14.6004 — 14.6665	.4275 — .4627	-2.005**
Gothenburg — E6	14.6004 — 14.6543	.4275 — .4218	-1.715*
Gothenburg — Smaller Cities	14.6004 — 14.7412	.4275 — .4476	-4.349***
Malmo — E6	14.6665 — 14.6543	.4627 — .4218	0.372
Malmo — Smaller Cities	14.6665 — 14.7412	.4627 — .4476	-2.219**
E6 — Smaller Cities	14.6543 — 14.7412	.4218 — .4476	-2.701***

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Chapter 2

Stochastic modeling of quantity and price dynamics in the Swedish gasoline market

Ritvana Rrukaj[a*], and Leif K. Sandal [a]

[a]Department of Business and Management Science, NHH Norwegian School of Economics, 5045 Bergen, Norway

Abstract

Medium-term gasoline demand shifts are required for the efficient management of consumers' responsiveness to price changes. A dynamic model that has performed well with a medium-term prediction is the model that describes the seasonal effects noticeable in the data. This study develops a model that takes into account trend and seasonality of the five-week moving averages of gasoline quantity and quantitative adjusted prices. The analysis includes three modeling steps: a pre-processing step where the raw data is transformed to moving averages, a model specification step, and the residuals modeling analysis step. The leftover residuals from the model have been fitted by two different normal distributions. The resulting approach is promising, providing a helpful tool in understanding the gasoline quantity-price relationships nationally and regionally in the retail gasoline market in Sweden.

Keywords: Gasoline quantity, price variations, cyclicality, stochastic modeling

*We thank Frode Steen for his helpful comments on this chapter.

2.1 Introduction

Gasoline is one of the most heavily traded oil-derived products worldwide. Economists and environmentalists have been interested in understanding the factors that determine gasoline demand, especially in the last 30 years. Oil is a necessity in daily life and is one of the main drivers of countries' economies. Gasoline prices fluctuate due to occasional shortages in available refining capacity and the uncertainty in the global oil market. Measuring consumers' responses to such price variations is essential for understanding the potential macroeconomic impacts of future petroleum disruptions and developing sustainable policies to respond appropriately to urban and global environmental challenges.

Different econometric methods have been explored in the retail gasoline industry to shed light on the gasoline demand and price relationship. Some approaches focus on the path estimation of price elasticity of gasoline demand using aggregate data while others use disaggregate data.[1] A substantial body of literature has provided the elasticity estimates at the country level (e.g., Baranzini and Weber, 2013; Liu, 2014; Rivers and Schaufele, 2015; Coyle et al., 2012; Crotte et al., 2010). In reality, however, consumers make gasoline purchase decisions daily, responding directly to the gasoline pump prices on the day of the purchase. However, models that use monthly, quarterly, or annual gasoline consumption and price data across states or large geographic areas necessarily aggregate these different purchase decisions, likely resulting in masking consumer responsiveness to local price changes. Some notable studies that have provided gasoline demand elasticities using micro-level data are Nguyen-Ones and Steen (2020), Levin et al. (2017), and Coyle et al. (2012).

The gasoline demand elasticity values documented across several gasoline markets worldwide provide us with information about the gasoline demand responses due to price changes. However, at the same time, how this relationship between gasoline demand and prices is evolving remains to be explored. This study develops a stochastic modeling approach to describe the medium-term relationships between gasoline quantity and volume-adjusted prices over time. We first construct a five-week moving average (MA) of the daily gasoline volumes and volume-adjusted prices. We then build an explicit model describing the seasonal effects of the gasoline quantity development. After removing the deterministic components trend (and seasonality), the received stochastic residues of gasoline quantity are modeled via two different normal distributions. We utilize daily gasoline price and volume data from the Swedish retail gasoline market at the station level.

Given the detailed data on gasoline quantity and price, we use the quantity-adjusted prices instead of the average prices typically used in the literature. Another contribution of this paper is to provide evidence regarding medium-term gasoline quantity and the price

[1] Dahl and Sterner (1991), Graham and Glaister (2002), and Lipow (2008) provide summaries of the results regarding gasoline demand elasticities.

correlational relationship, which is provided nationally and across six regional markets. In addition, a similar modeling approach is used to understand the relationship between gasoline sales and the prices of five oil brands. The R-squared values that the stochastic model yielded are large in all the analyses. They imply that the Swedish consumers exhibit similar gasoline purchase behavior across different local markets and that our dynamic model has been able to describe the stochastic development of the gasoline quantity.

Establishing these facts is important for reacting quickly to the increased concerns about the contribution of road transport to climate change and the growing demand for electric vehicles. Although advanced dynamic models are not used in the gasoline demand literature, they have been extensively included in the modeling of commodity prices in order to describe the strong seasonal effects that electricity prices possess.[2] Finally, it is worth mentioning that this approach for mid-term prediction of future price and quantity relationships can be extended to handle risks in other markets, such as agriculture and power.

This paper is organized as follows. Section 2.2 presents a literature review of gasoline demand modeling. Section 2.3 provides a brief explanation of the data used in this paper from the Swedish Competition Authority (SCA) and the Swedish retail gasoline market. Section 2.4 offers a more in-depth explanation of how the dynamic model and the residual fitting model are built and how they are used to explain the gasoline demand–price relationship. Section 2.5 presents and discusses the results and Section 2.6 concludes the paper.

2.2 Gasoline demand literature

This paper's contribution is related to but also differs from the previous research on gasoline demand estimation. We start by discussing the existing methodology and empirical literature in this field, and then we explain how new insights can be gained by applying our dynamic approach. Studies before the 1970s, that directly focused on analyzing gasoline elasticities, are minimal, while in the last four decades, the price elasticity of gasoline demand has been extensively studied. Besides, a wide range of studies has been conducted to explain how gasoline demand is related to price, income, and other variables, such as population, tax, public transportation, vehicle characteristics, and so on. However, the results vary between these studies, depending on the estimation technique, data used, and period. Studies on gasoline demand estimation also differ by the functional form introduced in their approaches. The log-linear demand model that has been applied in most studies utilizes monthly, quarterly, or annual gasoline consumption and average gasoline prices across states or large geographic areas. Such models are called "static" demand models, where the parameters on price represent an estimation of demand elasticity (Wadud et al., 2010). In cases where demand and price changes vary over time, as in the real gasoline market, these static models can be appropriate in capturing the immediate and long-term responses of a price change. In studies that assume that gasoline demand may take time to adjust to price fluctuations,

[2]Lucia and Schwartz (2002), Ghiassi et al. (2006), Keles et al. (2012), and Nepal (2015).

authors often apply a more flexible approach by including lagged values of prices and other control variables in the demand equation. The outcome of these "dynamic" methods is the elasticity of demand, which is the sum of the parameters on current and lagged prices (Dahl and Sterner, 1991; Lin and Prince, 2013). Levin et al. (2017) argue that the higher the levels of data aggregation, the more inelastic the price elasticity of gasoline demand. This argument makes sense since the aggregated data reflects the overall market rather than the micro demand from consumers facing pump prices from various gasoline stations. An exception is Nguyen-Ones and Steen's (2020) structural model analysis, which estimates the downstream competition level in gasoline markets by utilizing daily station-level demand and price data. Such data have played an important role in observing gasoline demand and prices at a higher frequency than the vast majority of gasoline studies that rely on a more aggregated market analysis.

Another strand of research on gasoline markets focuses on using the co-integration and ordinary least squares methods, which rely on macroeconomic control variables such as inflation (Hughes, 2008). Others rely on average income (Wadud et al., 2010; Baranzini and Weber, 2013; Coglianese et al., 2017), which is linked to aggregate gasoline demand, producing overall market elasticities. Using cointegration techniques, Schleiniger (1995) finds that price changes do not explain variations in demand and that income per capita has a significant impact on demand dynamics. Typically, the overall market elasticities presented in the literature are lower than one. Dahl and Sterner (1991) and Graham and Glaister (2002) observe lower short- and long-term price elasticities than income elasticities. Bons et al.'s (2008) meta-analysis reported a mean price elasticity of gasoline demand of -0.34 in the short-run and -0.84 in the long-run. Wang (2009) finds elastic gasoline demand in Perth, Australia. Moreover, he finds a price elasticity of -6.20 for a sample of stations with their nearest competitor within 4.2 km and -18.77 for a sample of stations located right next to their nearest competitor.

Despite the uncharted gap in gasoline demand modeling using high-frequency volume data at the station level, other gasoline consumption areas have been studied in detail. Much research has focused on estimating the price and income elasticities of demand.[3] Eckert (2013) suggests a need for more studies in demand estimation using high-frequency, station-level data. This paper fills this gap in the literature and observes how price and gasoline consumption relate to time.

[3]Espey's (1998) meta-analysis examined 41 articles, Graham and Glaister's (2004) examined 69 articles, Dahl and Sterner's (1991) over 100 articles, Goodwin et al. (2004) 240 articles, and Eckert's (2013) employed more than 75 empirical studies of gasoline retailing.

2.3 Data and the Swedish gasoline market

2.3.1 Data overview

The gasoline volume and retail price data used to examine the gasoline quantity–price relationship have been obtained from the SCA covering the period January 1, 2012 to December 31, 2012. The dataset comprises a time series of morning and afternoon gasoline prices and daily gasoline volume observations from 190 stations across Sweden (for more details about the market structure during the sample period, see Ganslandt and Ronnholm, 2014. For more information on the data, see the report on the Swedish gasoline market by Foros and Steen, 2013 and the article by Nguyen-Ones and Steen, 2020. We consider everyday prices the mean price between the two price observations for each day. A station must have demand and price observations available for each day in 2012 to be considered. Thus, this requirement allows for the same number of stations included in the national, regional, and brand-wide gasoline consumption analysis. This requirement should also reduce the mismatches on daily quantity and the volume-adjusted daily price for gasoline that can result from non-synchronous timing.

We have excluded the data of 43 stations because the volume and price values were missing for more than five consecutive days during a year, which may happen for several reasons, such as stations were closed or were misreporting issues. In other instances, we have substituted the missing observations by taking the simple average of the two preceding and two succeeding averages of gasoline volume and retail prices. Four big fuel companies dominated the Swedish gasoline market during our sample period: Statoil Fuel & Retail AB, Preem AB, OK-Q8 AB, and ST1 Energy AB. In 2012, they controlled more than 99% of the market: Statoil Fuel & Retail AB held 34.86% of the volume, OK-Q8 AB held 27.93%, ST1 Energy AB held 22.62%, and Preem AB held 14.22%, leaving only 0.36% of the volume to others. All stations that belonged to Preem AB had missing consecutive data for up to 18 days. Thus, we have removed this company from our analysis. Statoil Fuel & Retail AB and OK-Q8 AB have significantly higher shares compared to their rivals. Another structural change in this market is that many stations have been transformed into self-service retail stations. In 1999, automated stations had a 32% market share, while in 2013, it was up to 62% (Foros and Steen, 2013). We utilize consecutive daily volume and price measurements at the station level across regions, that is, Stockholm, Gothenburg, Malmö, the E6 highway, small cities, and rural areas, and across brands, that is, Statoil, OK-Q8, Shell, Jet, and St1. It is worth noting the importance of high-frequency gasoline micro-data because such data are rarely seen in the gasoline literature.

2.3.1.1 Gasoline volume characteristics

The absence of high-frequency gasoline volume data in the literature gives our paper a considerable advantage concerning modeling the relationship between gasoline quantity and

prices using a dynamic model. Figure 1 illustrates the daily gasoline volumes at the station level using a histogram.

Figure 1: The histogram of gasoline volume at the station level

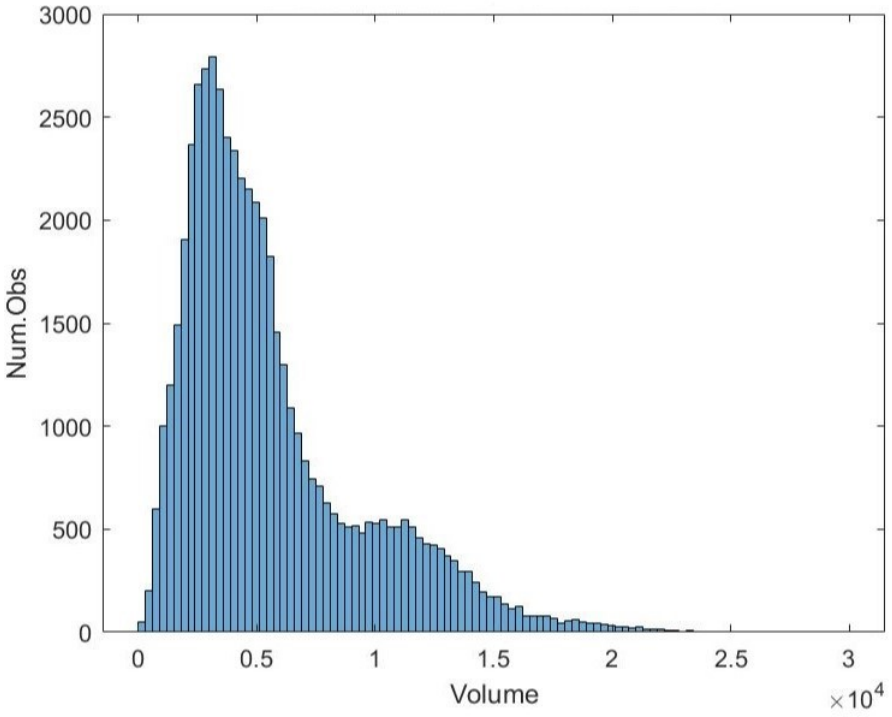

We also investigate the gasoline volume patterns at the regional level and for the five oil brands in the Swedish gasoline market; Table 1a shows the summary statistics across regions. As we have a different number of stations for each region, we cannot have a regional comparison between the mean volumes. From the number of observations, we see that rural areas have the highest number of observations, resulting in the highest number of stations in our sample. The standard deviation is highest for the E6 highway and Gothenburg and the lowest in rural areas and Malmö. Indicating that the variation of volume observations from the sample mean of the E6 highway and Gothenburg is more significant than in other regions.

Table 1a: Descriptive Statistics of **Gasoline Volume (liters)**

	Region					
Regions	**Stockholm**	**Gothenburg**	**Malmö**	**E6**	**Smaller Cities**	**Rural Areas**
Mean	7165	6586	5629	6790	6017	3022
St. Dev	3730	4448	3299	4475	3343	2078
Minimum	21.54	564.8	36.12	26.65	54.52	96.53
Maximum	23701	28242	20568	29833	21609	17504
Num. Obs	8052	7320	8418	8418	9516	12078

Note:Sample period, hourly observations, 2012:01:01-2012:12:31

The minimum and maximum values illustrate the highest and the lowest value in each of these regions. The lowest station volume data is observed at a station in Stockholm, while the highest value is observed at a station on the E6 highway. The latter pattern makes intuitive sense because the E6 highway is one of the main highways in Sweden, which runs through Malmö and Gothenburg. In contrast, the lowest volume value in Stockholm could be due to an unfavorable station location, or fierce competition. Table 1b contains the summary statistics of daily station volume data across brands.

Table 1b: Brand

Brand	**Statoil**	**OK-Q8**	**Shell**	**Jet**	**St1**
Mean	5478	4581	3682	10806	3746
St. Dev	3270.4	2654.5	1673.2	4062.7	1653.7
Minimum	330	321	26	91	21
Maximum	21713	21433	13092	29833	15200
Num. Obs	13542	8784	7686	10248	13542

Note:Sample period, hourly observations, 2012:01:01-2012:12:31

Statoil and St1 have the highest number of stations in our sample, while Shell has the lowest, that is, 21 stations. Jet has the highest standard deviation and St1 the lowest. The minimum value of daily gasoline volume is observed at a station owned by St1, while the highest volume value is observed at a station owned by Jet. As mentioned earlier, these two brands ran only self-service stations. Utilizing the above volume datasets, we have calculated the average daily gasoline sales in Sweden, across six different geographical areas and five oil brands. Figure 2 displays three plots at those analysis levels. As can be seen from Figure 2, gasoline consumption in Sweden appears to be relatively stable, with a few spikes before public holidays. The first spike is observed before the Easter holiday, the second before the national day of Sweden, the third is before Midsummer, and the fourth spike is before Christmas Eve. It is not surprising to observe the most significant jump in gasoline consumption before Midsummer (the time of the year around the summer solstice, which is a big festival on the Swedish calendar and leads to a mass exodus from cities such as Stockholm and Gothenburg to the countryside, where people go to their summerhouses). Figure 2 reflects this jump in gasoline consumption before June 21, 2012 when gasoline sales increased approximately twice before the Midsummer day and dropped almost twice on the Midsummer day.

Figure 2: Gasoline consumption during 2012

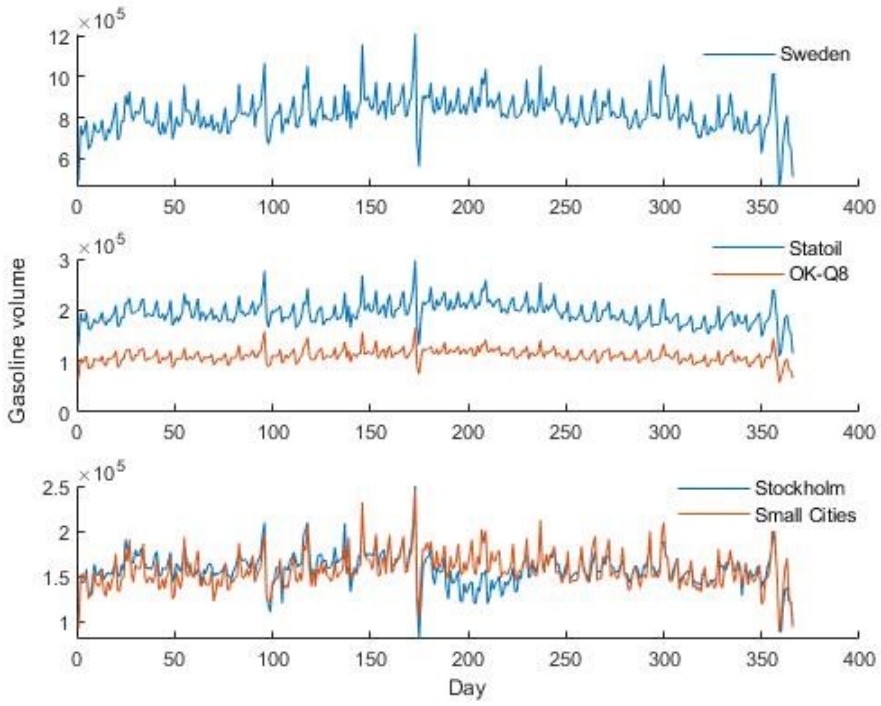

However, with different ranges, daily gasoline consumption seems to exhibit similar patterns on all plots. The highest volume values happened on the same days, and all spikes and drops appear synchronized. People seem to have similar behavior regarding filling up on gasoline across local markets, and also, each brand seems to face similar gasoline sale fluctuations at the national level. This tendency also makes intuitive sense because gasoline consumption is typically the highest before public holidays when people drive more, causing the big spikes that we see in Figure 2 and less during regular days. In the days following the increase, the demand decreases, causing significant declines.

2.3.1.2 Price characteristics

The distribution of price data is illustrated by the histogram in Figure 3. The x-axis contains the bins with price ranges, while the y-axis contains the number of price values that fall in each bin. The daily gasoline prices in Sweden vary from SEK/liter 13.5 to SEK/liter 16. The most significant variation of the pump prices between days is less than SEK/liter 3 throughout the year. The vast majority of bins have more than 1500 observations, and these bins range from SEK/liter 14.3 to SEK/liter 15.3.

Figure 3: A histogram of prices at the station level

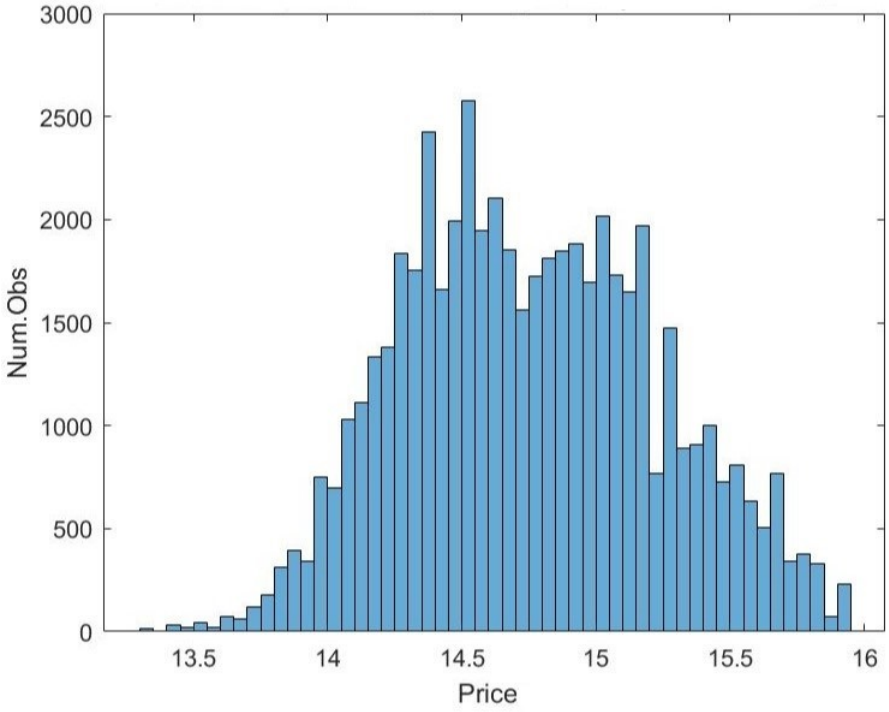

Table 2a and Table 2b below show a summary of the statistics of the pump price data at the station level for each region and each brand in the Swedish gasoline market during the sample period. The standard deviation is relatively small and almost the same across local markets. The lowest price value is observed in Stockholm and Gothenburg, while the highest value is observed in rural areas.

Table 2a: Descriptive Statistics of **Gasoline Prices (SEK/liter)**

	Region					
Regions	**Stockholm**	**Gothenburg**	**Malmö**	**E6**	**Smaller Cities**	**Rural Areas**
Mean	14.84	14.64	14.68	14.72	14.77	14.84
St. Dev	0.46	0.45	0.48	0.49	0.46	0.45
Minimum	13.3	13.3	13.4	13.4	13.44	13.68
Maximum	15.93	15.93	15.94	15.94	15.93	15.95
Num.Obs	8052	7320	8418	8418	9516	12078

Note:Sample period, hourly observations, 2012:01:01-2012:12:31

Again, similar to what is illustrated in Table 2a, the standard deviation is relatively small and very similar from brand to brand. The minimum pump price value is observed at stations of OK-Q8, Jet, and St1, while the highest pump price value is observed at a station of OK-Q8. Statoil and St1 have more stations nationally in our sample. Overall, the histogram of pump prices and the tables with the summary statistics show that the amplitude

Table 2b: Brand

Company	**Statoil**	**OK-Q8**	**Shell**	**Jet**	**St1**
Mean	14.83	14.85	14.83	14.58	14.71
St. Dev	0.46	0.46	0.47	0.45	0.47
Minimum	13.44	13.3	13.45	13.3	13.3
Maximum	15.93	15.95	15.94	15.68	15.78
Num.Obs	13542	8784	7686	10248	13542

Note:Sample period, hourly observations, 2012:01:01-2012:12:31

of the pump price variations is relatively smaller throughout the year in the Swedish gasoline market compared to other gasoline markets (West, 2004; Foros and Steen, 2013; Baranzini and Weber, 2013).

2.3.1.3 Volume-adjusted prices

In our analysis, we introduce the daily volume-adjusted prices calculated at the national, regional, and brand levels as follows:

$$VAP_t = \frac{\sum P_{i,t} \times Q_{i,t}}{Q_t}$$

where $P_i,_t$ is the pump price at station i and day t, $Q_{i,t}$ is the quantity sold at station i and day t, and Q_t is the total gasoline quantity sold on day t. In the numerator, we are summing the multiplication of daily pump prices and gasoline volumes across stations. We believe that volume must be taken into consideration because VAP_t reflects sales. Hence, if a station in the sample has a high price and does not sell much, then this price does not affect the VAP_t because the gas station has to sell quite a lot before it influences the price. Figure 4 shows three plots of VAP_t for each analysis level.

Figure 4: The daily volume-adjusted price

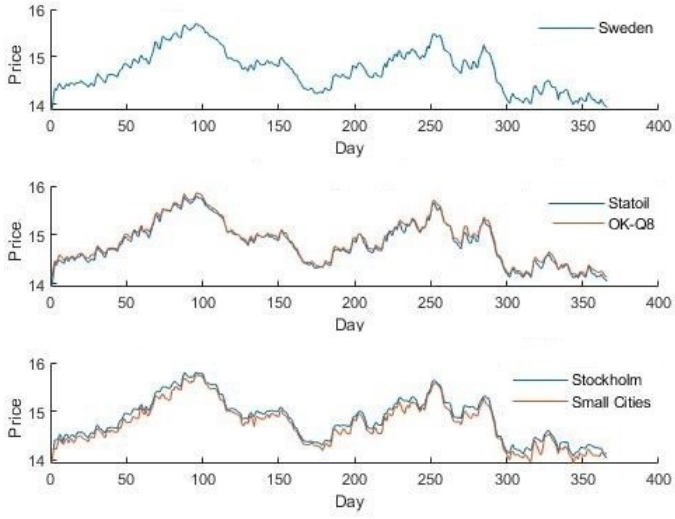

Figure 4 displays the seasonal pattern over the yearly cycle of the gasoline price. Prices rose sharply from January to April, and at the beginning of May, they rapidly decreased until the end of the month. Then, the cycle seems to repeat with much more volatile gasoline prices during the winter season. Gasoline retail prices are generally lower during winter. The seasonal patterns can also be seen in the Rotterdam spot prices, acting as the main driver of gasoline price variations. Another factor affecting the seasonal patterns might be the governmental regulations that change seasonally. These regulations require different gasoline formulas due to environmental considerations. During warmer seasons, the additives necessary for the summer formula increase the production costs for refineries and pump prices. Gasoline volume sold in the summer is less likely to evaporate in warm weather. Refiners replace cheaper and more evaporative gasoline components with less evaporative but more expensive components. On the contrary, colder weather can lower retail prices because the winter formula used is less expensive to produce.[4] In contrast to the non-apparent seasonal fluctuations seen in the volume data (Figure 2), there are evident seasonal fluctuations in gasoline prices (Figure 4).

2.3.2 The Swedish retail gasoline market

The fuel retail trade in Sweden was characterized by competition between a limited number of companies during 2012. These companies had their retail network spread throughout the country, covering extensive territory. Besides, pricing was largely transparent to consumers, and companies' retail price changes often occurred in parallel. In the years before 2012, the Swedish retail gasoline market had undergone several structural changes in terms of company acquisitions and mergers, which led to boosting the market concentration (Ganslandt and Ronnholm, 2014). To more closely examine consumer behavior in this market, we have analyzed the gasoline quantity–price relationship for each day of the week in each region. Figure 5 shows the percentage of gasoline volume sold by day of the week across six different geographical areas: Stockholm, Gothenburg, Malmö, the E6 highway, smaller cities, and rural areas. Smaller cities include cities with more than 33,000 and less than 80,000 inhabitants, while rural areas comprise less than 10,000 inhabitants. Each color in Figure 5 represents the percentages of gasoline volume sold in a particular region each day of the week during 2012. The gasoline volume by day of the week is calculated as a percentage of the total volume over the week across areas. If the same gasoline quantity is sold every day of the week in a region, it will be approximately 14% each day (black dashed line). On Fridays, the average volume rises from 14% to 16.5% in rural areas to 16% in smaller cities, continuing with a 1.5% increase in local markets such as the E6 highway, Malmö, and a lower rise in Gothenburg (1.3%) and Stockholm (1.1%). In Stockholm, the highest gasoline sales happened on Fridays and the lowest on Saturdays. Figure 5 clearly shows that gasoline sales spiked on Fridays in

[4] https://www.caranddriver.com/news/a15339380/the-vapor-rub-summer-versus-winter-gasolineexplained ; https://www.epa.gov/gasoline-standards/gasoline-reid-vapor-pressure

all local markets.

Figure 5: Volume share by day of the week

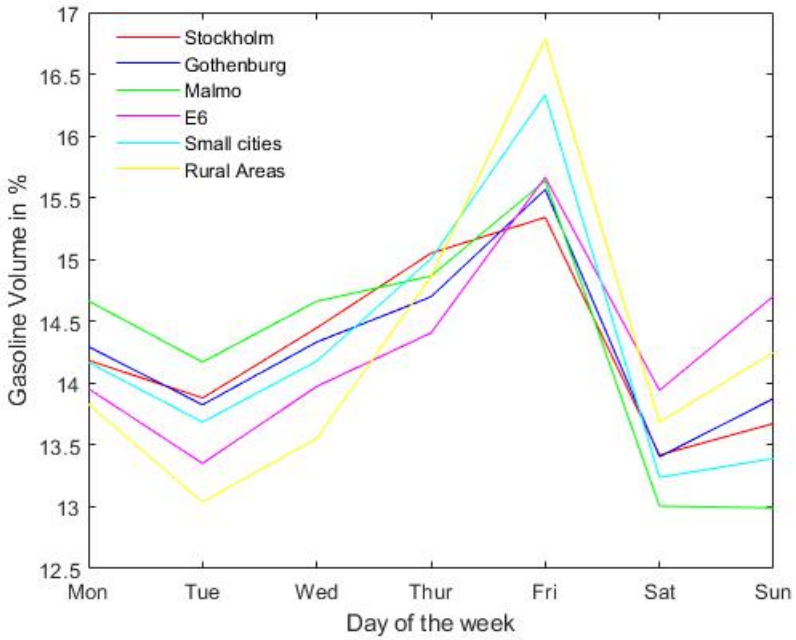

Consumers seem to fill up their cars on Fridays in all areas more than on other days of the week. In the rural areas, Fridays' volumes are approximately 18% higher than the daily average of 14% and are significantly different from Tuesdays' gasoline volumes. The gasoline consumption pattern that we observed in Figure 5 can be explained by the fact that people in Sweden tend to fill their car tanks on the last working day, before the weekend. People who live in cities plan to spend their weekends in the countryside, which forces them to fill their car tanks before long trips. We consider this pattern as the "weekend" effect.

A second reason can be related to the technological aspect, that is, the car tank capacity. The car tank has a certain capacity, depending on the type of car. Therefore, the capacity limitation does not allow people to fill their car tanks more than the specified amount of gasoline. A car delivers on average 16.6 km/liter and has a fuel tank capacity of 43–65 liters of gasoline. Ohman and Lindgren (2003) report that, on average, Swedish employees travel about five kilometers a day to and from their workplace. Moreover, they document that every third employee must cross at least one municipality border to get to work. However, the exact mileage delivered by a vehicle may differ based on several factors, including the type of vehicle, driving style, quality of fuel, and the condition of various components. In about a week, tank capacity will substantially fall, and it will be time for drivers to fill up their gas tanks again. On average, 14.3% of the sales take place on Sundays in rural areas (green line), 13.8% in Stockholm (red line), 14.7% on the E6 highway (black line), and so on. On Mondays, gasoline sales are close to the daily mean in almost all the regions, while gasoline sales drop on Tuesdays, meaning that people do not purchase much gasoline on Tuesdays compared

to the other days of the week. Not surprisingly, the average gasoline volume on Sundays on the E6 highway is significantly higher compared to other days except for Fridays, and in other regions due to the "weekend" effect. The dramatic decrease in gasoline consumption on Saturday suggests that a substantial number of consumers filled their cars on Fridays. Finally, we can conclude that consumers pursue similar gasoline consumption patterns across different geographical areas. In Figure 6, we observe how volume-adjusted prices change within a week across the same local markets. Besides the differences in geographical zones, other factors contribute to regional differences in gasoline prices, such as distance from the supply location, local competition, and operating costs. The same colors represent the same regions as in Figure 5.

Figure 6: The regional average volume-adjusted price by day of the week

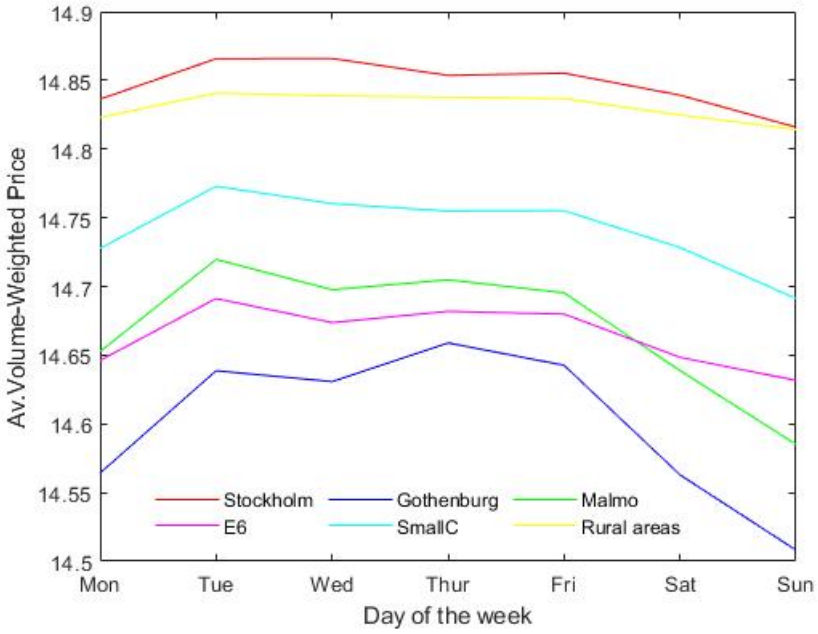

Price fluctuations are similar across regions. Stockholm and rural areas seem to be the local markets where filling up a car with gasoline costs the most while Gothenburg is the cheapest across all week days. In Figure 7, we present the price differences between Gothenburg and other regions in percentage. The average volume-adjusted prices by day of the week in Stockholm and rural areas are 2% higher than the average volume-weighted prices in Gothenburg on Mondays. The lowest price differences can be observed between Gothenburg and Malmö and between Gothenburg and the E6 highway, where the average volume-weighted prices in Malmö and the E6 highway are less than 1% higher than in Gothenburg. Surprisingly, although we would have anticipated lower prices in Stockholm than in isolated markets such as rural areas, the data shows no significant variations between the average weekly volume-adjusted prices in Stockholm and rural regions. Figure 7 indicates less competition in Stockholm.

Figure 7: The prices differences between Gothenburg and other regions

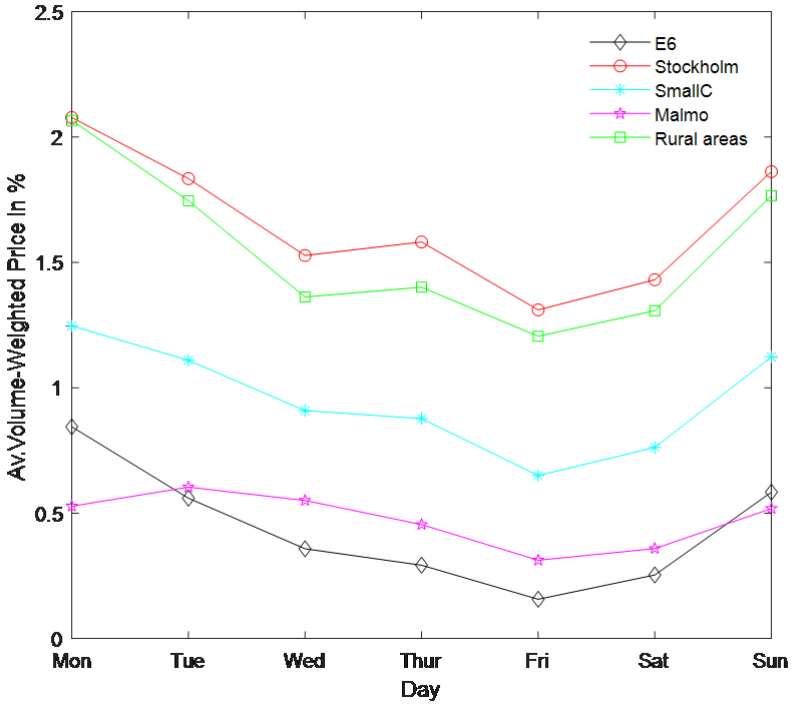

Tuesdays are days with higher gasoline prices in all regions except Gothenburg, while Sundays are days with the lowest gasoline prices. This pattern may be a result of several factors. First, gasoline demand is relatively low on weekends due to the "weekend" effect. Second, the financial markets and Rotterdam spot market are closed on Saturdays and Sundays. Thus, no new information arrives on those days. Third, knowing that sales is lower on Sundays, oil companies may attract consumers by keeping prices down. In the Norwegian market, although the price cycle amplitude is much larger than in the Swedish gasoline market, a few consumers still shift their purchases on low price days (Foros and Steen, 2013). Weekly-synchronized cycles in other markets might benefit consumers by making it easier for them to shift gasoline purchases strategically on days with relatively low prices (Noel, 2002; Doyle et al., 2010; Lewis and Noel, 2011; Noel and Chu, 2015; Siekmann, 2017; Noel, 2019). Noel et al. (2019) observed the gasoline market in Perth, Australia, where declines always occurred on Wednesdays and peaks occurred on Thursdays. In the Swedish gasoline market, we observed a different behavior. The average volume-adjusted price peaks always occurred on Tuesdays except in Gothenburg, and the declines always occurred on Sundays. The average volume-adjusted price cycles in each region in Sweden as in other markets (Foros and Steen, 2013; Noel, 2019) lasted precisely a week, and Figure 6 clearly illustrates the weekly pattern. To closely examine the effect of the average volume-adjusted price cycles on the day-of-the-week sales in each region, we compared Figures 5 and 6. Figure 5 reports pure volume measurements, while Figure 6 shows prices that contain some volume effects.

On Fridays, we see that companies try to keep the price high in all regions, because they generally expect high demand. The same pattern can be noted on Sundays across all areas when price is the lowest of the week and sales are down as well. This pattern might have more to do with time than prices, suggesting that if the oil companies expect that the Friday pattern is strong, they will keep the price as high as possible as long as it does not affect customers. On the contrary, they may also try to lower the price when the expected volume is low to capture "random" customers. As a next step, we have calculated the total variation of the average volume-adjusted prices for each day of the week across regions. Figure 8 shows the price variations together with the histogram of the volume-adjusted prices by day of the week in smaller cities.

Figure 8: Price distribution by day of the week in smaller cities

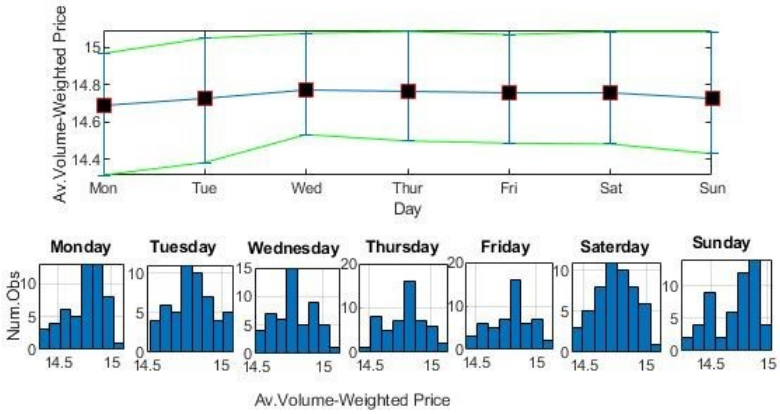

As can be seen, the spread of gasoline prices by day of the week in smaller cities is insignificant, and the average volume-adjusted price is on the lower half of the spread except for Mondays, which indicates that there are more cases with low prices than high prices in those days. It is worth noting that Figure 8 does not really show the price dispersion across weekdays and how representative the point of an average price is. To see how much price values deviate from the mean, we have calculated the coefficient of variation by the day of the week across all local markets. The coefficient of variation is expressed as a percentage in Figure 9.

Figure 9: The coefficient of variation of price values across regions

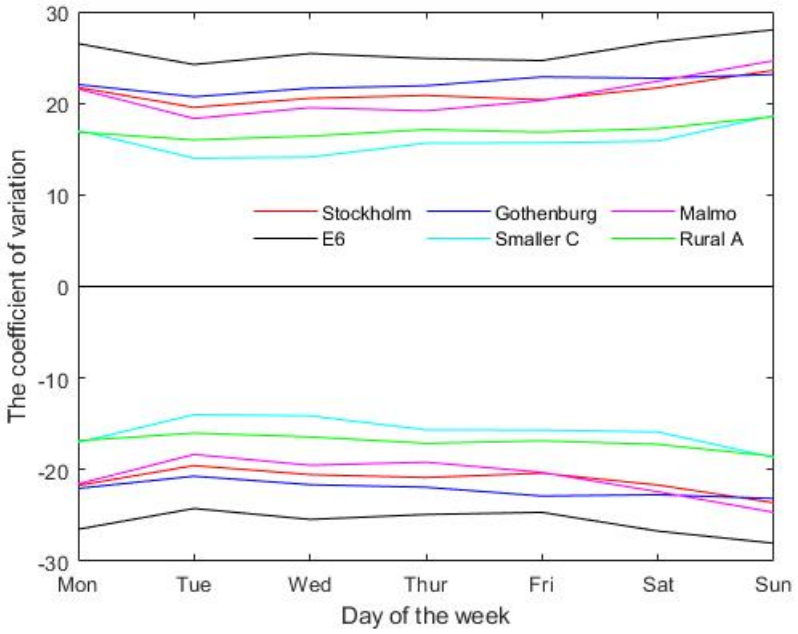

As can be seen, prices for the E6 highway and Gothenburg are more volatile. Indicating that the price dispersion is lower for the high price regions than for the low price regions. We have observed some extreme price values throughout the year in Stockholm, but they are not frequent because they disappear when considering the coefficient of variation. The volume-adjusted prices seem to be less volatile during weekdays than weekends.

2.4 Methodology

In this study, as a first step, we calculated the daily volume-adjusted prices nationally, regionally, and across brand. The stations with low sales do not significantly affect the daily volume-adjusted prices. Hence, the everyday gasoline prices that we acquire will represent an effective price per liter of gasoline sold on a specific day in Sweden, across local markets and brands. The total daily gasoline consumption and the volume-adjusted prices will be used as input data in the next step of our analysis, that is, the pre-processing step. In order to smoothen the input data, we have pre-processed them. More specifically, the first transformation on the data starts by calculating a five-week MA for both gasoline quantity and prices. To calculate the five-week moving averages, we have chosen a window size of 35 days. The formula of the moving averages for gasoline quantity and prices is as follows:

$$Moving\ Average\ \Phi_t = \frac{1}{w} \sum_{s=t-(w-1)}^{t} \phi_t$$

where Φ_t denotes either the five-week MA of gasoline quantity or prices, ϕ denotes the daily gasoline volume or price observations, w denotes the MA window, and t represents the time of observation. This pre-processing step smoothens the data and assembles the short-run dynamics of Swedish gasoline demand and price into medium-run dynamic measures.

We then specify a dynamic model that maps the 35-day MA of gasoline quantity. We ran a non-linear fitting model to find the internal parameters used to transform the data. Transformation applies the parameters to data. The non-linear model estimates the interception and the other parameters in the following formula:

$$Y_t = \beta_1 + \beta_2 t + \beta_3 t^2 + \beta_4 \sin(\frac{2\pi t}{t_1})x_t + \beta_5 \cos(\frac{2\pi t}{t_2})t + \beta_6 \frac{\sin(\frac{2\pi t}{t_3})}{x_t} + \beta_7 \cos(\frac{2\pi t}{t_3})x_t^2 + \beta_8 \sin(\frac{2\pi t}{t_4})x_t + z_t \quad (1)$$

where Y_t denotes the five-weeks moving averages of demand at time t scaled by 10^5, x denotes the five-weeks moving averages of volume-weighted gasoline prices at time t , t is a numeric time indicator of the observation, and $t_1 = 190.5, t_2 = 104.5, t_3 = 199.3$, and $t_4 = 99.9$. Our model contains periodic terms with different periods $(t_1, ..., t_4)$, but it is an aperiodic function with many quasi-periodic properties. Hence, it has flexible base functions that can mimic complex signals. z_t is a vector of residuals, and the eight β-terms are parameters that need to be decided. The fixed parameters β_1, β_2, ..., β_8 are used to transform (i.e., map) the five-week moving averages of gasoline volume.

Finally, we took the residuals from formula (1) and fitted them with the sum of two normal distributions with different mean and standard deviations. The model used to map the residuals is as follows:

$$z_t = \delta_1 * N(\delta_2, \delta_3) + \delta_4 * N(\delta_5, \delta_6) + \varepsilon_t.$$

where, z_t are the residuals at time t and δ_1 to δ_6 are the parameters to be estimated. The parameters δ_2, and δ_5 represent the mean of the normal distributions, while δ_3, and δ_6 represent the sigma values, and ε_t represents the white noise. Again, the δ-parameters are estimated using the non-linear model. By applying our modeling mechanism, every time we get new observations, the whole process can be repeated. Similarly, we have transformed the gasoline volume data for six local markets and five fuel brands in the Swedish gasoline market. It is worth highlighting that the same advanced model is used to map the 35-day MA of gasoline quantity at the country level, across regions, and across brands.

2.5 Results

Figure 10 shows the five-week MA of gasoline quantity. Figure 11 shows the five-week MA of the volume-adjusted prices at the national level.

Figure 10: Five-week MA of volume

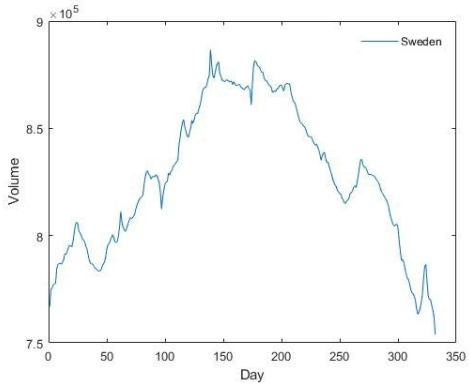

Figure 11: Five-week MA of prices

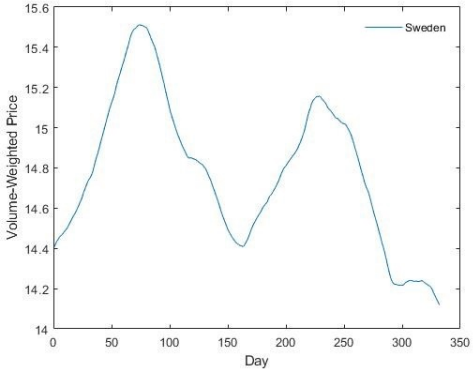

Using the five-week MA of volume data as Y_t, and the five-week MA of prices as x_t in the formula (1), we have estimated all the β_1, β_2, ...,β_8 parameters. Table 3 presents the values of β that best suited the five-week MA of gasoline quantity at the national, regional, and company level by using the same model specification.

All parameters, β_2 through β_8 measure the strength of various medium and long-run dynamics in the five-week MA of gasoline quantity. In almost every region, the values for β_2 are positive, and the values for β_3 are negative. The four parameters that measure the relationship between the five-week MA of volume and gasoline prices β_5, β_6, β_7, and, β_8, are highly statistically significant at the 1% significance level across all regions. The R-squared values presented in Table 3 indicate that our advanced model generally performs well and offers a perfect fit across areas. The model seems to perform particularly well in region 4 (the E6 highway), in region 5 (smaller cities), and in region 6 (rural areas), the last of which exhibits the highest R-squared value. Regions 4 and 5 show R-squared values of over 0.95, while regions 1, 2, and 3 indicate R-squared values that exceed 0.84. Moving to the results that we acquired by running the model on the 35-day MA of gasoline sales across oil brands, we see that the sign of β_2 is positive, and the sign of β_3 is negative across all brands. All coefficients are statistically significant at the 1% significance level, which provides strong evidence of the presence of seemingly quasi-cyclical patterns in gasoline sales. We have removed the parameters β_4 and β_5 from formula (1) because they did not improve the fit of the five-week MA of gasoline sales for Shell and Statoil, respectively. Statoil exhibits the highest R-squared values (0.979). Moreover, the R-squared values for automated brands such as Jet and St1 are also substantially high: both brands showed R-squared values over of 0.90.

As a final step, we took the residuals from the dynamic model in all analyses, and we fitted them with two different normal distributions. Figure 12 presents the empirical kernel distribution of the residuals at the national level versus the combined normal distribution.

Table 3: Five weeks model parameters at national, regional and brandwide level

β	**Sweden**	**Stockholm**	**Gothenburg**	**Malmö**	**E6**	**Small Cities**	**Rural Areas**	**Statoil**	**OK-Q8**	**Shell**	**Jet**	**St1**
β_1	7.941	1.697	1.359	1.325	1.319	1.486	0.763	1.778	1.047	0.733	2.945	1.352
	(0.023)	(0.011)	(0.005)	(0.001)	(0.005)	(0.005)	(0.006)	(0.006)	(0.004)	(0.001)	(0.011)	(0.005)
β_2	0.009	-0.001	-0.001	-	0.004	0.001	0.004	0.003	0.001	0.001	0.001	0.001
	(0.0003)	(0.0001)	(7.2E-05)		(8.2E-05)	(6.9E-05)	(7.6E-05)	(7.9E-05)	(5.1E-05)	(2.3E-05)	(0.0001)	(6.8E-05)
β_3	-2.8E-05	1.4E-06	7.1E-07	-3.9E-07	-1.3E-05	-4.5E-06	-1.1E-05	-1.1E-05	-4.6E-06	-3.7E-06	-5.1E-06	-2.7E-06
	(7.9E-07)	(3.5E-07)	(1.9E-07)	(2.3E-089)	(2.3E-07)	(1.8E-07)	(2.1E-07)	(2.1E-07)	(1.3E-07)	(6.8E-08)	(3.9E-07)	(1.8E-07)
β_4	-0.025	-0.024	-0.019	-0.017	-	-0.007	0.039	0.022	-0.004	-	-0.032	-0.011
	(0.005)	(0.002)	(0.001)	(0.0005)		(0.001)	(0.001)	(0.001)	(0.0008)		(0.002)	(0.001)
β_5	-0.0002	-0.0002	-0.0001	-3.2E-05	0.0003	-0.0001	6.7E-05	-	3.4E-05	1.8E-05	-0.0001	-7.3E-05
	(3.2E-05)	(1.5E-05)	(8.4E-06)	(6.5E-06)	(1.5E-05)	(8.1E-06)	(5.9E-06)		(4.4E-06)	(1.7E-05)	(7.9E-06)	(7.9E-06)
β_6	-3.329	-1.487	-1.201	-1.069	-0.711	-1.1	2.1	0.8	-0.563	-0.28	-2.06	-1.005
	(0.318)	(0.141)	(0.075)	(0.031)	(0.041)	(0.073)	(0.079)	(0.081)	(0.053)	(0.012)	(0.155)	(0.071)
β_7	0.002	0.001	0.001	0.001	0.0005	0.0006	-0.002	-0.001	0.0003	0.0001	0.001	0.0006
	(0.0003)	(0.0001)	(8.4E-05)	(3.6E-05)	(1.2E-05)	(7.9E-05)	(8.5E-05)	(8.2E-05)	(5.7E-05)	(3.3E-06)	(0.0001)	(7.8E-05)
β_8	-	0.004	0.001	0.0007	-0.004	0.0005	-0.002	-0.0008	-0.0006	-0.0008	0.002	0.0005
		(0.0002)	(0.0001)	(7.7E-05)	(0.0002)	(0.0001)	(0.0001)	(9.9E-05)	(7.3E-05)	(5.6Ee-05)	(0.0002)	(0.0001)
N	332	332	332	332	332	332	332	332	332	332	332	332
R^2	0.954	0.878	0.885	0.848	0.963	0.962	0.978	0.979	0.965	0.966	0.916	0.918

Notes: All coefficients in the table are significant at 1% level. Standard errors are presented in parentheses below parameter estimates.

Figure 12: The fitted distribution of the residuals at national level

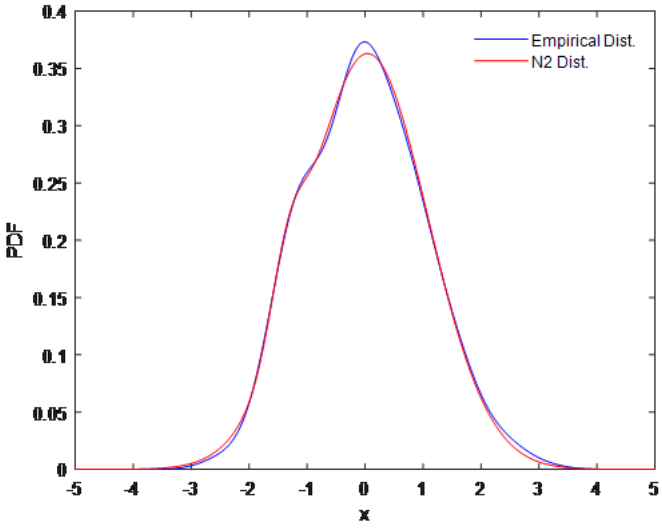

Figure 13 illustrates the empirical distributions of the residuals and the combined normal distributions across local markets. The observed distribution of the residuals in each area is characterized by a different color, while the combined normal distribution is the red curve in each plot. For example, the yellow line in the first plot indicates the empirical kernel distribution of the residuals after mapping the five-week MA of gasoline quantity in Stockholm, and the red line is the combined normal distribution. The second plot presents the empirical kernel distribution (blue) of the residuals fitted by two normal distributions (red) in Gothenburg and so on.

Figure 13: The fitted distribution of the residuals across regions

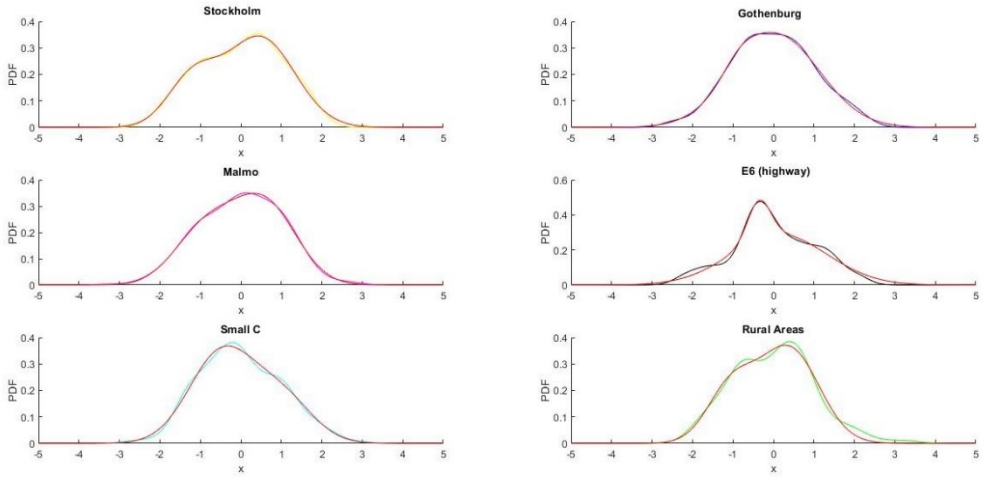

Finally, we again fitted the residuals of the dynamic model for each brand with two different normal distributions. Figure 14 shows the distributions of Statoil, OK-Q8, Shell, Jet, and St1.

Figure 14: The fitted distribution of the residuals across fuel brands

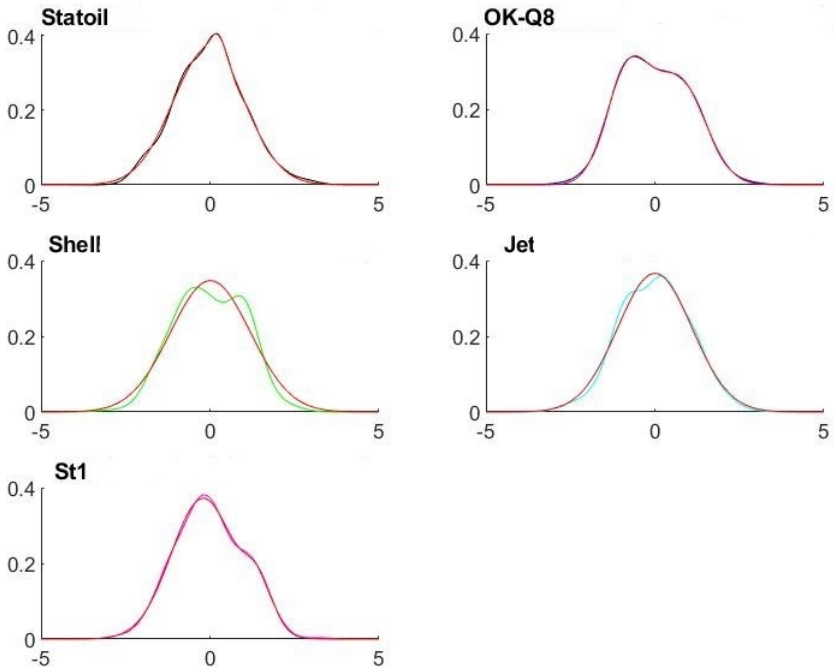

Again, the red curve is the combined normal distribution, and the other colors present the empirical kernel distribution of each brand. The leftover residuals ε_t are white noise with no structural information at all market-analysis levels developed in our study.

2.5.1 Discussion of results and implications

This study developed a modeling approach that can potentially be utilized by market participants to estimate future medium-term gasoline quantities and vice-versa. Our model consisted of three modeling steps: a pre-processing step, where the raw data is transformed to moving averages; model specification, where flexible expressions are incorporated to map the MA of gasoline quantity; and modeling the residuals from the mapping process. This study employed these three modeling techniques to model the medium-run relationship between gasoline quantity and prices in the Swedish gasoline market. We begin by calculating a five-week MA of the daily gasoline quantity and volume-adjusted price observations. These resulting data series are quasi-periodic signals with unpredictable fluctuations in amplitude and frequency. After creating the five-week moving averages of Swedish gasoline quantity and prices, we developed a dynamic model that expresses the gasoline volume as a function of time, prices, and interaction terms between time and prices. As a final modeling step, we model the leftover residuals by summing two different normal distributions. The R-squared values that our model yielded are high at national, regional, and brand levels, and all coefficients are highly significant.

While there are similarities between standard econometric studies and the dynamic modeling approach we present in this paper, we believe that there are several key differences. First, it is common to mathematically transform non-stationary data into stationary data before carrying out an econometric analysis; for instance, in the econometric analysis of gasoline time-series, it is common to log transform the raw quantity and price data. In contrast, our approach utilizes non-stationary data. To smoothen the raw data, we converted them into five-week moving averages. Second, we perform additional analysis on our model residuals, which is uncommon in most econometric studies. By carrying out this extra step, we could fit our model residuals despite having initially high R-squared model values. Third, while most econometric studies intend to identify causality in a given economic system, our dynamic model reveals a highly detailed image of the dynamics between gasoline quantity and prices in Sweden, across local markets and brands in the Swedish retail gasoline market during the sample period. In addition, our model can provide a long-run price elasticity of demand that varies over time. The calculated price elasticity of gasoline demand ranges from -0.17 to 0.16 with an average of -0.02, indicating low responsiveness of gasoline's quantity demanded relative to changes in its price. We considered it as long-run elasticity due to the data smoothening in the pre-processing step.

Furthermore, due to the non-linear nature of our model, this analysis reveals many additional details and dynamics about the mid-term relationship between quantity and prices, which are not easily visible using standard linear regressions. Finally, our results can be back-transformed, while the results obtained from statistical modeling do not have this option.

2.6 Conclusion

This paper presented a new approach to modeling the gasoline quantity and price dynamics in the Swedish retail gasoline market. We observe from the data that the quantity for gasoline exhibits similar patterns nationally, locally, and across brands. During 2012, the highest quantity of gasoline was sold on Fridays in each region in Sweden, while the lowest gasoline sales happened on Sundays. Our analysis demonstrates that consumers in Sweden do not strategically choose the day of the week when purchasing gasoline. In all regions, Tuesdays were the most expensive days, and in the majority of areas, Sundays were the cheapest days of the week to fill gasoline. Surprisingly, the data also shows that gasoline is more expensive in highly populated and less populated regions (Stockholm and rural areas), where the competition is expected to be high (Stockholm) and low (rural areas). In addition, gasoline in Sweden is more expensive in regions with fewer price fluctuations (Stockholm and rural areas) and cheaper in areas with the frequent price fluctuations (Gothenburg). These results are consistent with what was found in Valadkhani (2013). However, several other studies have argued that consumers may shift their gasoline purchases to known low-price days (Noel, 2018; Noel, 2012; Noel and Chu, 2015). The results of this paper suggest that

Swedish consumers do not do this.

Similar patterns have also been observed in the Norwegian gasoline market (Foros and Steen, 2013). While this paper provides new insights into national and regional gasoline consumption patterns and the impact of the volume-adjusted prices, it also highlights the need to use daily consecutive quantity–price data in the analysis of the gasoline market and represents an exception in the literature. The absence of high-frequency volume data at station level in the literature has prevented many studies from examining the relationship between gasoline quantity and price over time.

Given the gasoline price values, decision-makers can apply our approach to calculate likely shifts in gasoline sales. Moreover, the national and regional results can enhance the understanding of the Swedish gasoline retail market and benefit the country's competition authority, in cooperation with the government, to devise tax policies and other local policy instruments. Overall, our approach allows for understanding the underlying dynamics between gasoline quantity and prices and provides the potential for predicting future quantity–price relationships for specific local markets, brands, and the Swedish gasoline industry.

References

[1] Ackah, I., and Adu, F., (2013),"Modelling gasoline demand in ghana: a structural time series analysis," *International Journal of Energy Economics and Policy*, Vol. 4, iss. 1, pp. 76-82.

[2] Anderson, Th.W., and Hsiao, C., (1981), "Estimation of dynamic models with error components, " *Journal of the American statistical Association*, Vol.76, iss.375, pp. 598-606.

[3] Atkinson, B., (2009), "Retail gasoline price cycles: Evidence from Guelph, Ontario using bi-hourly, station-specific retail price data, " *The Energy Journal*, Vol. 30, iss. 1.

[4] Baranzini, A., and Weber, S., (2013), "Elasticities of gasoline demand in Switzerland, " *Journal of Energy policy*, Vol. 63, pp. 674-680.

[5] Borenstein, S., and Shepard, A., (1993), "Dynamic pricing in retail gasoline markets, " *National Bureau of Economic Research.*

[6] Brons, M., and Nijkamp, P., and Pels, E., and Rietveld, P., (2008), "A meta-analysis of the price elasticity of gasoline demand. A SUR approach," *Journal of Energy Economics*, Vol. 30, iss. 5, pp. 2105-2122.

[7] Brons, M., and Nijkamp, P., and Pels, E., and Rietveld, P., (2006),"A meta-analysis of the price elasticity of gasoline demand. A system of equations approach, " *Tinbergen Institute Discussion Paper.*

[8] Byrne, D.P., and De Roos, N., (2019), "Learning to coordinate: A study in retail gasoline, " *Journal of American Economic Review*, Vol. 109, iss. 2, pp. 591-619.

[9] Brynjolfson, Erik and Chris Kemerer (1996), "Network Externalities in Microcomputer Software: An Econometric Analysis of the Spreadsheet Market," *Management Science*, Vol. 42, iss. 12, pp. 1627-47.

[10] Byrne, D.P., and de Roos, N., (2017), "Consumer search in retail gasoline markets, " *The Journal of Industrial Economics*, Vol. 65, iss. 1, pp. 183-193.

[11] Castanias, R., and Johnson, H., (1993), "Gas wars: Retail gasoline price fluctuations, " *Journal of The Review of Economics and Statistics*, pp. 171-174.

[12] Chandra, A., and Tappata, M., (2011), "Consumer search and dynamic price dispersion: an application to gasoline markets, " *The RAND Journal of Economics*, Vol. 42, iss. 4, pp. 681-704.

[13] Chesnes, M., (2016), "Asymmetric pass-through in US gasoline prices, " *The Energy Journal*, Vol. 37, iss. 1.

[14] Clemenz, G., and Gugler, K., (2009), "Locational choice and price competition: some empirical results for the Austrian retail gasoline market, " *Booktitle:Spatial Econometrics*, pp. 223-244.

[15] Coglianese, J., and Davis, L.W., and Kilian, L., and Stock, J.H., (2017), "Anticipation, tax avoidance, and the price elasticity of gasoline demand, " *Journal of Applied Econometrics*, Vol. 32, iss. 1, pp. 1-15.

[16] Coyle, D., DeBacker, J., and Prisinzano, R., (2012), "Estimating the supply and demand of gasoline using tax data, " *Journal of Energy Economics*, Vol. 34, iss. 1, pp. 195-200.

[17] Crotte, A., and Noland, R.B., and Graham, D.J., (2010), "An analysis of gasoline demand elasticities at the national and local levels in Mexico, " *Journal of Energy Policy*, Vol. 38, iss. 8, pp. 4445-4456.

[18] Dahl, C.A., (1979), "Consumer adjustment to a gasoline tax, " *The review of economics and statistics*, pp. 427—432.

[19] Dahl, C., and Sterner, T., (1991), "Analysing gasoline demand elasticities: a survey " *Energy Economics*, Vol. 13, iss. 3, pp. 203-210.

[20] Dahl, C.A., (2012), "Measuring global gasoline and diesel price and income elasticities, " *Journal of Energy Policy*, Vol. 41, pp. 2-13.

[21] De Roos, N., (2004), "A model of collusion timing, " *International Journal of Industrial Organization*, Vol. 22, iss. 3, pp. 351-387.

[22] Eckert, A., (2013), "Empirical studies of gasoline retailing: A guide to the literature, " *Journal of Economic Surveys*, Vol. 27, iss. 1, pp. 140-166.

[23] Eckert, A. and West, D.S., (2004), "Retail gasoline price cycles across spatially dispersed gasoline stations, " *The Journal of Law and Economics*, Vol. 47, iss. 1, pp. 245-273.

[24] Espey, Molly., (1998), "Gasoline demand revisited: an international meta-analysis of elasticities, " *Journal of Energy Economics*, Vol. 20, iss. 3, pp. 273-295.

[25] Foros, Ø., and Steen, F., (2013), "Vertical control and price cycles in gasoline retailing, " *The Scandinavian Journal of Economics*, Vol. 115, iss. 3, pp. 640-661.

[26] Foros, Ø., and Steen, F., (2013), "Retail Pricing, Vertical Control and Competition in the Swedish Gasoline Market, " *Report for the Swedish Competition Authority.*

[27] Haucap, J., and Møller, H.C., (2013), "The effects of gasoline price regulations: experimental evidence, " *Available at SSRN 2482288.*

[28] Higgins, E.J., and Peterson, D.R., (1999), "Day-of-the-week autocorrelations, cross-autocorrelations, and the weekend phenomenon, " *Journal of Financial Review*, Vol. 34, iss. 4, pp. 159-170.

[29] Hughes, J., and Knittel, C.R., and Sperling, D., (2008), "Evidence of a shift in the short-run price elasticity of gasoline demand, " *The Energy Journal*, Vol. 29, iss. 1, pp. 113-134.

[30] Graham, D.J., and Glaister, S., J., (2002), "The demand for automobile fuel. A survey of elacticities " *Journal of Transport Economics and Policy*, Vol. 36, iss. 1, pp. 1-25.

[31] Kleit, A.N., (2003), "The economics of gasoline retailing: Petroleum distribution and retailing issues in the US, " *American Petroleum Institute.*

[32] Lewis, M.S., (2009), "Temporary wholesale gasoline price spikes have long-lasting retail effects: The aftermath of Hurricane Rita, " *The Journal of Law and Economics*, Vol. 52, iss. 3, pp. 581-605.

[33] Lewis, M., and Noel, M., (2011), "The speed of gasoline price response in markets with and without Edgeworth cycles, " *Journal of Review of Economics and Statistics*, Vol. 93, iss. 2, pp. 672-682.

[34] Lewis, M., (2012), "Price leadership and coordination in retail gasoline markets with price cycles, " *International Journal of Industrial Organization*, Vol. 30, iss. 4, pp. 342-351.

[35] Lewis, M., (2015), "Odd prices at retail gasoline stations: focal point pricing and tacit collusion, " *Journal of Economics & Management Strategy*, Vol. 24, iss. 3, pp. 664-685.

[36] Li, S., and Timmins, C., and Von Haefen, R., (2009), "How do gasoline prices affect fleet fuel economy?, " *American Economic Journal: Economic Policy*, Vol. 1, iss. 2, pp. 113-137.

[37] Lin, C.Y.C., and Prince, L., (20139, "Gasoline price volatility and the elasticity of demand for gasoline, " *Journal of Energy Economics*, Vol. 38, pp. 111-117.

[38] Liu, W., (2014), "Modeling gasoline demand in the United States: A flexible semiparametric approach, " *Journal of Energy Economics*, Vol. 45, pp. 244-253.

[39] Nguyen, M., and Steen, F., (2018), "Measuring Market Power in Gasoline Retailing: A Market-or Station Phenomenon?, " *NHH Dept. of Economics Discussion Paper 6.*

[40] Murat, Y.S., and Ceylan, H., (2006), "Use of artificial neural networks for transport energy demand modeling, " *Journal of Energy policy*, Vol. 34, iss. 17, pp. 3165-3172.

[41] Nalca, A., and Cai, G. G., (2008), "Who Benefits From Banning Discriminatory Wholesale Pricing When Retailers Can Price Match, " *Available at SSRN 3299781.*

[42] Noel, M.D., (2007), "Edgeworth price cycles, cost-based pricing, and sticky pricing in retail gasoline markets, " *Journal of The Review of Economics and Statistics*, Vol.89, iss. 2, pp. 324-334.

[43] Noel, M.D., (2007), "Edgeworth price cycles: Evidence from the Toronto retail gasoline market, " *The Journal of Industrial Economics*, Vol. 55, iss. 1, pp. 69-92.

[44] Noel, M.D., and Chu, L., (2015), "Forecasting gasoline prices in the presence of Edgeworth Price Cycles, " *Journal of Energy Economics*, Vol. 51, pp. 204-214.

[45] Noel, M.D., (2019), "Calendar synchronization of gasoline price increases, " *Journal of Economics & Management Strategy*, Vol. 28, iss. 2, pp. 355-370.

[46] Pock, M., (2010), "Gasoline demand in Europe: New insights, " *Journal of Energy Economics*, Vol. 32, iss. 1, pp. 54-62.

[47] Polemis, M.L., and Tsionas, M.G., (2016), "An alternative semiparametric approach to the modelling of asymmetric gasoline price adjustment, " *Journal of Energy Economics*, Vol. 56, pp. 384-388.

[48] Ramsey, J., and Rasche, R., and Allen, B., (1975), "An analysis of the private and commercial demand for gasoline, " *The Review of Economics and Statistics*, pp. 502-507.

[49] Rivers, N., and Schaufele, B., (2015), "Salience of carbon taxes in the gasoline market, " *Journal of Environmental Economics and Management* , Vol.74, pp. 23-36.

[50] Schleiniger, R., (1995), "The Demand for Gasoline in Switzerland- In the Short and in the Long Run, " *Institute for Empirical Research in Economics. University of Zurich.*

[51] Scott, K.R., (2012), "Rational habits in gasoline demand, " *Journal of Energy Economics*, Vol. 34, iss. 5, pp. 1713-1723.

[52] Scott, K.R., (2015), "Demand and price uncertainty: Rational habits in international gasoline demand, " *Journal of Energy*, Vol. 79, pp. 40-49.

[53] Valadkhani, A., (2013), "Seasonal patterns in daily prices of unleaded petrol across Australia, " *Journal of Energy policy*, Vol. 56, pp. 720-731.

[54] Van Meerbeeck, W., (2003), "Competition and local market conditions on the Belgian retail gasoline market, " *Journal of De Economist*, Vol. 151, iss. 4, pp. 369-388.

[55] Wadud, Z., and Graham, D.J., and Noland, R.B., (2010), "Gasoline demand with heterogeneity in household responses, " *The Energy Journal*, Vol. 31, iss. 1, pp. 47-74.

[56] Wadud, Z., and Noland, R.B., and Graham, D.J., (2010), "A semiparametric model of household gasoline demand, " *Journal of Energy Economics*, Vol. 32, iss. 1, pp. 93-101.

[57] Wadud, Z., and Graham, D.J., and Noland, R.B., (2009), "A cointegration analysis of gasoline demand in the United States, " *Journal of Applied Economics*, Vol. 41, iss. 26, pp. 3327-3336.

[58] Walsh, K., and Enz, C.A., and Canina, L., (2004), "The impact of gasoline price fluctuations on lodging demand for US brand hotels, " *International Journal of Hospitality Management*, Vol. 23, iss. 5, pp. 505-521.

[59] Wang, Z., (2008), "Collusive communication and pricing coordination in a retail gasoline market, " *Journal of Review of industrial organization*, Vol. 32, iss. 1, pp. 35-52.

[60] Wang, Z., (2009), "Station level gasoline demand in an Australian market with regular price cycles, " *Australian Journal of Agricultural and Resource Economics*, Vol. 53, iss. 4, pp. 467-483.

[61] Wang, Z., (2009), "(Mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation, " *Journal of Political Economy*, Vol. 117, iss. 6, pp. 987-1030.

Chapter 3

Asymmetric cost transmission and market power: An examination of the Swedish retail gasoline market

Ritvana Rrukaj[a]*, and Frode Steen[b]

[a]Department of Business and Management Science, NHH Norwegian School of Economics, 5045 Bergen, Norway

[b]Department of Economics, NHH Norwegian School of Economics, 5045 Bergen, Norway

Abstract

Estimating non-linear autoregressive distributed lag models, using a panel of detailed daily station-level data from the Swedish market, we establish short-run cost pass-through for both average pump prices and volume-adjusted gasoline prices. The disequilibrium error takes longer to be fully corrected when we test for the existence of volume-adjusted price asymmetry than in the case of pump price asymmetry, suggesting that oil companies are more focused on pricing on days and at stations with larger sales. Estimating models across regions, different station service levels, and local competition levels, we find that rural areas and gasoline stations less exposed to local competition are able to impose larger and more prolonged asymmetry on retail gasoline prices. This supports several studies concluding that asymmetry in cost pass-through is positively correlated with the degree of market power. The same pattern is uncovered when we estimate models for varying service levels. Full-service stations have a higher and more prolonged asymmetry in pricing than automated self-service stations. The former group of stations has, in several studies, been found to have higher market power than the latter.

Keywords: Gasoline markets, asymmetric short- and long-run cost pass-through, market power, volume-adjusted prices, station heterogeneity, local competition, asymmetric dynamic multipliers

*We thank David P. Byrne from the University of Melbourne for very useful insights and suggestions.

3.1 Introduction

Many studies have aimed to examine the behavior of firms in the gasoline market. One important strand of literature has focused on Edgeworth price cycles (Edgeworth, 1925; Maskin and Tirole, 1988) and saw tooth-like pricing patterns.[1] Several studies have focused on studying the presence of coordination and coordinated effects.[2] Others have attempted to investigate whether retail gasoline markets have prices where cost changes are asymmetrically reflected in consumer prices.[3] This pattern, often referred to as "rocket and feathers" (Bacon, 1991), has been observed and analyzed in several markets and across many products (Peltzman, 2000).[4] In response to wholesale price changes, consumer prices rise quickly, and the costs are passed-through immediately, while they tend to fall at a much slower pace when the wholesale price decreases. This asymmetrical pattern is well documented in a large number of empirical studies of gasoline markets.[5]

Some studies also look at how this asymmetry is related to the competitive situation in the market. For instance, Borenstein et al. (1997) attribute the asymmetry to tacit collusion. Likewise, Verlinda (2008) attributes the asymmetry to market power and provides evidence that a high service level on the station level (implying higher market power) increases the pass-through asymmetry. Another strand of literature focuses more directly on how market characteristics, such as spatial competition, differences in station amenities, and the effects of brand identity and contractual forms, influence the degree of competition in retail gasoline markets. In a recent study, Nguyen-Ones and Steen (2021) examine the combined effect of several factors, using the same dataset as we do here, and show that these factors have a non-trivial impact on the level of competition. They also provide an overview of this literature.

However, fewer studies relate the order of price-cost asymmetry to market characteristics as those outlined above. In this study, we exploit a particularly rich dataset to uncover to which extent asymmetric cost pass-through is related to such market characteristics. An

[1] See, for example, Noel (2016) and Eckert (2013) for comprehensive surveys on pricing, and Edgeworth cycles and competitive pricing games. Others have treated cycles as an outcome of intertemporal price discrimination (e.g., Conlisk et al., 1984), and some authors have argued that complexity through cycles can be used to soften price competition (e.g., Carlin, 2009; Ellison and Wolitzky, 2012).

[2] See, for example, Byrne and De Roos (2019) on coordination and collusion; Foros and Steen (2013a), Foros et al. (2021) on coordination and price cycles; and among others Hastings (2004), Slade (1987), Netz and Taylor (2002), and Nguyen-Ones and Steen (2021) on local competition.

[3] See, for example, Borenstein et al. (1997), Asplund et al. (2000), Johnson (2002), Verlinda (2008), Apergis and Vouzavalis (2018), and Byrne (2019).

[4] Peltzman (2000) analyzes as many as 77 different consumer goods and 165 producer goods (food, agriculture, metal products, textiles). He found asymmetric price responses for several of these.

[5] See eg. Borenstein et al., 1997; Deltas, 2008; Asplund et al., 2000; Eckert, 2002; Verlinda, 2008; Balmaceda and Soruco, 2008; Hofstetter and Tovar, 2010; Lewis, 2011, Lewis and Noel 2011, Apergis and Vouzavalis, 2018 and Byrne, 2019.

observation from the "rocket and feather" literature is that lack of competition will allow for short-run asymmetry in pass-through rates. However, the literature is in agreement regarding long-run pass-through, where some authors even impose this in their models (see, e.g., Lewis, 2004; Apergis and Vouzavalis, 2018; Byrne, 2019).

To the extent that the degree of competition is linked to short- and long-run asymmetry, we anticipate that the longer the retailers can impose asymmetry from cost changes in the wholesale prices, the higher is their market power. Byrne (2019) also posits that rural stations exposed to lower competition exhibit longer asymmetric cost pass-through than stations in cities exposed to more competition from neighboring gasoline stations.

Likewise, differences in competition due to station heterogeneity have been shown to create differences in prices and market power. Eckert and West (2005) find evidence that station characteristics affect sellers' price setting, and suggest the presence of imperfect competition. Haucap et al. (2017) conclude that prices are positively related to station service levels, and Shepard (1991) finds that stations charge a full-service markup. Estimating a structural conjectural variation model, Nguyen-Ones and Steen (2021) find that manned service stations in Sweden have a higher market power (markups). This suggests the presence of a more persistent short-run pricing asymmetry for manned service stations.

We analyze the extent to which the Rotterdam spot market prices are passed on to consumer prices for gasoline in Sweden and how quickly this happens. We estimate short- and long-run effects using dynamic time-series models for different regions and allow for station heterogeneity both with regards to station service-level and differences in local competition.

Furthermore, as far as we know, earlier studies have not been able to use volume-adjusted prices in their analyses of pass-through asymmetry because they require detailed data on quantities that are seldom available. Hence, recent and abundant empirical literature has used (average) pump prices or recommended (quoted) prices to show that asymmetric price adjustments exist.

Our dataset includes daily observations on transaction prices and quantities for 147 gasoline stations in Sweden for 2012. As can be seen in many national gasoline markets, the Swedish downstream market is concentrated, with four companies sharing 99% of the market in our sample period. Statoil Fuel & Retail AB has a market share of 34.9% in gasoline volume, St1 Energy AB has 22.6%, OK-Q8 has 27.9%, and Preem AB has 14.2% (SPBI, 2013).

In line with previous studies, the dataset allows us to not only analyze asymmetry for average pump prices but also use volume-weighted prices in our models. A company should be more concerned with obtaining "profitable" asymmetry at stations with higher sales than at stations with lower sales. The volume-adjusted prices will account for this. Hence, we expect to find clearer pass-through asymmetry by applying volume-adjusted prices.

We apply a non-linear autoregressive distributed lag (NARDL) model developed by Shin et al. (2013) to estimate these dynamic effects. Similar to the empirical process followed by

Apergis and Vouzavalis (2018), we have tested the presence of short- and long-run asymmetries by deriving both the positive and negative partial sum decompositions of the dependent variable. This methodology is a one-step estimation model and can test the presence of cointegration between variables through a bounds-testing process irrespective of whether these variables are I(0) or I(1). Moreover, the asymmetric dynamic multipliers obtained from the model would allow us to observe the long- and short-run responses of the gasoline prices to positive and negative oil innovations.

Using the NARDL models, we find clear evidence of short-run price asymmetry both using average pump prices and volume-adjusted prices for the Swedish national market. Prices increase more after input price increases, then they fall after input price reductions. Interestingly, and as expected, the volume-adjusted price asymmetry takes longer to correct the disequilibrium error than the average pump price asymmetry does. This suggests that companies are indeed more focused on the pricing at larger gasoline stations since the profits from asymmetry will be higher from pricing "correctly" for these than for smaller stations.

In the following sections, we examine the presence of asymmetric pass-through for Rotterdam spot prices on volume-adjusted prices across six different geographical regions, between service and self-service stations, and between stations with a distance to the nearest competitor of more than 3 km and within 3 km.

Turning first to heterogeneity across regions, we find a larger pricing asymmetry in rural markets than in cities, which is consistent with what Byrne (2019) finds from Ontario, Canada. Related to this finding, our estimates of cost pass-through for gasoline stations that have less local competition measured from the distance to the nearest competing station also exhibit more persistent short-run asymmetry. These results are consistent with the empirical gasoline literature, which show that markets that enjoy some market power have an incentive to prolong the pass-through of cost changes to consumer prices (Verlinda, 2008; Deltas, 2008; Byrne, 2019).

However, Stockholm is an exception in our analysis, contradicting the above findings. Somewhat surprisingly, the patterns found in Stockholm resemble those found in rural markets. This is in line with what Foros and Steen (2013b) reported in studying the Swedish gasoline market. Despite being the largest city in Sweden, Stockholm has price patterns, both in terms of dynamics and levels, that are closer to what one observes in rural areas rather than what is observed in other large Swedish cities.

When it comes to heterogeneity at service levels, pricing asymmetry in terms of the timing and the amplitude of the adjustment is found to be less pronounced in the self-service stations than in the service gasoline stations. This difference is likely due to the market power that the service stations exhibit. Using mostly the same dataset, Nguyen-Ones and Steen (2021) estimate a significant difference in the markup between self-service and service stations, where the self-service stations have 29% lower markups.[6]

[6]In their Table 6, p.22, model (4) they find a significant benchmark markup of 0.37 for all stations. For

Finally, we also perform "back of the envelope" calculations to show the effects of the asymmetric pricing behavior. We find that as much as 1,136,982 thousand SEK is the estimated overcharge in the average price model. This increases by 49,008 thousand SEK, when we account for the effect of one extra day and volume-adjusted prices. Hence, this suggests that the model using average prices underestimates the full overcharge due to asymmetry by 4.1%. The overcharge amounts to only 2.7% of all gasoline spending in 2012, obviously due to the high exogenous taxes and the Rotterdam wholesale price representing on average 93% of the retail price. Focusing on the margin, this overcharge represents 40–60% of the daily average margin. Thus, we also calculate the total gross margin in SEK for 2012 to compare it to these overcharges. The total margin in 2012 is around 3 billion SEK. Hence, the 2.7% increase in total expenditure for gasoline implies a doubling of the gross margin, suggesting that for the oil companies, asymmetric pricing is a major way of increasing profits. In markets such as transparent oligopolies, where prices are mostly determined by exogenous costs and overall demand is typically inelastic, companies have strong incentives to impose asymmetric pricing.

Section 3.2 presents the literature review, Section 3.3 describes our dataset and the Swedish market, Section 3.4 presents the econometric approach, and Section 3.5 presents and discusses the results. Section 3.6 discusses potential policy implications. Finally, Section 3.7 offers concluding remarks while proposing future research directions.

3.2 Price dynamics and cost pass-through in gasoline markets

3.2.1 The empirical literature

This section briefly summarizes some of the research on price dynamics and cost pass-through in the gasoline market. The studies considered typically used international wholesale gasoline spot prices and average consumer level pump prices for gasoline. Most of the studies confirm that the gasoline price dynamics resemble the "rockets and feathers" phenomenon. This phenomenon applies across countries, data periods, and data-use frequency. Some studies use other dynamic modeling approaches but eventually arrive at the same conclusions.

The vast majority of studies use error correction models (ECM) to identify possible asymmetries in short-run consumer price responses due to changes in wholesale prices (for a detailed survey on ECM as an econometric specification in gasoline literature, see Grasso and Manera, 2007). Most studies find asymmetry in the price responses to costs. Borenstein et al. (1997) and Borenstein and Shepard (1993) found asymmetric price responses. Borenstein and Shepard (1993) found empirical support for their model of asymmetric price responses and implicit cooperation, and therefore argue that collusion drives "rockets and feathers." Borenstein et al. (1997) used the semi-monthly index prices from regional stations in the US from 1986 to 1992. They again found that prices increase faster with an increase in spot

the self-service stations the markup estimate is only 0.26 (Nguyen-Ones and Steen, 2021).

prices than the rate at which they decrease with reductions in spot prices.

Asplund et al. (2000) investigated the Swedish market, where they analyzed indicative daily pump prices from Shell for the period 1980–1996. Applying an error correction model showed that there is only a gradual adjustment of consumer prices to the cost of gasoline in the short run but a one-to-one relationship in the long run. They also found asymmetry in price responses. The pump prices obtained an immediate significant price change of 0.70 SEK/liter after an increase in the Rotterdam spot price of one SEK. On the contrary, the corresponding reduction in the Rotterdam spot price led to a price change of only 0.35 SEK/liter.[7]

Johnson (2002) analyzed 15 gasoline markets in his study (1996–1998, daily index prices in the US). He calculated average development in the price after a positive and negative shock in the spot price and differences between spot and consumer price. Based on his data, Johnson (2002) found that negative shocks are absorbed far slower in consumer price than positive shocks. Verlinda (2008) analyzed the US market for the period 2000 to 2003. Utilizing weekly station data and in line with other studies, he found asymmetric price responses. Lewis (2011) analyzed the US market but used somewhat longer weekly data (2000–2007) and observed both the station and area levels, also finding asymmetry.

Balmaceda and Soruco (2008) analyzed the petrol market in Chile based on the weekly data of 44 petrol stations in Santiago from 2001 to 2004. They also used an error correction model and demonstrated asymmetry in pricing. Station prices reacted more quickly to an increase in the spot price (6.2% increase over increase in purchase price) than to a decrease in the spot price (the reduction in the station price is 10.2% lower than the reduction in the cost of gasoline). Faber (2015) analyzed Dutch daily station data for the period 2003–2005. He found price symmetry, but only for 38% of the stations in the dataset.

A more recent study, Apergis and Vouzavalis (2018) apply the NARDL model as we use here. They address pricing asymmetry utilizing weekly data that span from January 2009 to July 2016 in the US, UK, Spain, Italy, and Greece. They find mixed results across countries for short- and long-run asymmetries. Asymmetric price responses due to wholesale price changes are found only in Italy and Spain. Byrne (2019) analyzes data from Ontario and estimates gasoline pricing asymmetry across different regional areas. Using weekly data, he finds symmetric pass-through of positive and negative cost shocks in urban markets. However, in rural markets, a short-run asymmetry is observed. 45% of positive shocks and 21% of negative shocks are passed through to prices at the same week, and the rest is corrected fully in the upcoming week.

Some studies do not find such asymmetries, but they are in the minority. For example, Bachmeier and Griffin (2003) found, in the US gasoline market, that the results depend on the

[7] Asplund et al. (2000) include both spot prices and tax in their error correction model. They can thus also comment on how the guide price changes as a result of changes in taxes and fees. They find that tax change has an immediate effect on the retail price of gasoline, in the sense that the short-term effect of tax on the recommended price is very close to one both in the short-run and in the long-run.

estimation method and data frequency. In particular, they found that when they use daily data, there is less asymmetry and they used standard two-step estimation of error correction models (see Engle and Granger, 1987). They also found less asymmetry than Borenstein et al. (1997) when using the same estimation method as Borenstein et al. (1997). Unlike Borenstein et al. (1997), Bachmeier and Griffin (2003) had access to a higher-frequency dataset.

The literature unambiguously concludes that all changes in spot prices are absorbed in their entirety in the long-run, which means that one can expect full pass-on (cost transfer) at consumer prices in the long-run. Some studies have even imposed this on their models (see, e.g., Lewis, 2011). The supportive argument is that one should theoretically expect full pass-through, and to the extent that this is not observed, there are biases in the models (Lewis, 2004; Verlinda, 2008).

Although most studies have used error correction models, several studies have looked at price dynamics for gasoline prices where other methods were used. Here too, one typically finds asymmetry in price responses when wholesale costs are measured through spot price change. Noel (2007) used a Markov (switching) model to analyze price cycles across 19 Canadian cities during the period 1989–1999, a phenomenon that has also been observed in many markets, including, for example, the Norwegian market (See Foros and Steen, 2013). Noel (2007) found that the price cycles are less asymmetrical, shorter, and larger (higher variance/amplitude) when more competition is measured as a greater presence of independent gasoline stations in the market. He attributes this to the literature and models around asymmetric price dynamics.

Later, Lewis and Noel (2011) used a similar Markov model and show that passing on costs is even faster in markets with clear Edgeworth cycles. They argue here that cycles are more important for passing on costs than market characteristics. Eckert (2003) found less evidence of asymmetry when he used error correction models to show that cycle movements dominate the asymmetric price response in the Canadian gasoline market from 1989 to 1994. Douglas and Herrera (2010) have a slightly different approach to price dynamics and pricing systems where they analyzed the daily price patterns for nine petrol stations in Philadelphia, Pennsylvania, USA. They extended an earlier study of the same market by Davis and Hamilton (2004), who estimated hazard models that predicted price changes and analyzed pricing systems for the same market. However, Douglas and Herrera (2004) had instead estimated an autoregressive probability model. The model predicts the probability of price change at a time t due to the historical distribution of price changes, including previous price changes and historical distance between the cost of goods and prices. They typically also found asymmetric price responses.

3.2.2 Sources of asymmetric cost pass-through in gasoline

In gasoline literature, asymmetric cost pass-through has two strands of explanations: search costs and collusive behavior. Consumers' search intensity in gasoline markets is considered to reflect competition in the market. More specifically, this search intensity explains firms' asymmetric price responses to oil price fluctuations (Tappata, 2009). Tappata (2009) argues that when the input prices for gasoline are high in the current period, consumers search very little because they do not expect a significant variation in prices. However, when the input prices are low in the current period, consumers expect that prices will exhibit higher variations in the next period, therefore they increase their search. Moreover, Tappata (2009) develops an oligopolistic model with competitive firms. Consumers have partial information and they endogenously choose to search less when gasoline input prices are anticipated to be high in the market setting. The "rocket and feathers" pattern has also been confirmed empirically by Lewis (2011) and Yang and Ye (2008). Lewis (2011) examines retail price dispersion using station-level data from Southern California, USA. He relates the price dispersion to models of consumer search. Similarly, Yang and Ye (2008) provide a search-based theoretical and empirical analysis of asymmetrical pricing.

Douglas and Herrera (2010) test whether the pricing dynamics can be attributed to uninformed customers, uninformed manufacturers, or strategic pricing, and they concluded that the results do not support that the price patterns are driven by menu costs or information access/collection (search costs). Instead, they believe in finding support for strategic pricing to explain the price dynamics for petrol stations. Therefore, as per their observations, larger price increases due to cost shocks are passed on to customers immediately, while negative cost price shocks take longer to be passed on to customers.

Verlinda (2008) expands the model of Borenstein et al. (1996) by introducing differentiation, where the differentiation takes the form of geographical distances. Increased geographical distances can be seen as increased product differentiation, which reduces competition as the players have more coordination options than just price. This opportunity, in turn, enables them to set prices asymmetrically. Consistent with these findings, Byrne (2019) and Deltas (2008) also empirically documented that increased local market power allows for an increased degree of asymmetric pricing.

Indeed, several authors argue that collusion may be behind the asymmetry in fuel markets' price responses. Clarke and Houde (2013) analyzed a price cartel from 2005 to 2006 that was revealed by the competition authorities in Canada; this cartel collaborated by using a typical asymmetric pricing pattern that followed the dynamics of wholesale price changes. Rather than collaborating around a fixed margin, they collaborated through an asymmetric pricing pattern where they typically increased prices more after an increase in wholesale price. At the same time, when there was a decline in wholesale price, there was a less decrease in prices, thereby increasing their profits. At the same time, they collaborated to delay price increases to increase market shares from competitors outside the cartel. Borenstein et

al. (1996) argue more in the direction of implicit collusion, stating that the collaboration on pricing collapses faster when the product cost increases since competitors' potential penalties are lower. Balmaceda and Soruco (2008) report that the results of asymmetric pricing responses due to the markets' transparency and unique design in Chile cannot be explained based on search costs but rather on implicit cooperation.

In sum, we can say that the literature has not unambiguously clarified which factors influence and drive such asymmetric price responses (Lewis and Noel, 2011). Although previous studies have investigated pricing asymmetry, they focused only on a country or city-level analysis. In what is the closest article to ours, Byrne (2019) addresses pricing asymmetry in different regional markets: urban and rural. In this context, rural markets are found to exhibit more asymmetry in gasoline price responses due to oil price fluctuations. The gasoline price associated with a specific date and market is calculated as the average price across all stations on a given date and at a given market. However, for the station-level analysis, the average price of a station during a week is used.

While our approach has several similarities to the literature, some important features differ. First, due to a rich dataset, we incorporated both daily station prices and volume-adjusted prices. We can analyze both these and compare the results to a corresponding analysis of average prices. To the extent that firms are engaging in pass-through asymmetry, we anticipate the effect to be more pronounced for volume-weighted prices. To the extent that firms are asymmetrically passing through costs to consumer prices, they should be more focused on obtaining this at the stations where and on days when sales are at their highest.

Furthermore, we investigate pricing asymmetry across different station characteristics, mirroring heterogenous competition levels. In particular, we differentiate between service and self-service gasoline stations across regions and stations with heterogeneous spatial competition.

3.3 A first look at the market: Data description

The dataset consists of daily data on Rotterdam spot prices, average transaction prices, and gasoline volumes sold for 147 branded stations across 6 different geographical areas in Sweden. This dataset was obtained from the Swedish Competition Authority (SCA). The period covered is from January 1, 2012 to December 31, 2012. Four major companies dominated the Swedish gasoline market during our sample period: Statoil Fuel & Retail AB (operating the brands Statoil and Jet), Preem AB, St1 Energy AB (operating the brands St1 and Shell), and OK-Q8 AB. In 2012, they controlled more than 99% of the market: Statoil Fuel & Retail AB holding 34.86% of volumes, OK-Q8 holding 27.93%, ST1 Energy AB holding 22.62%, and Preem AB holding 14.22%, leaving only 0.36% of the sales to others.[8] Of these brands, Jet and St1 only operated self-serviced stations. While Statoil, Preem, OK-Q8, and Shell only represented full-service stations. We require that a station must have

[8]https://www.konkurrensverket.se/globalassets/publikationer/uppdragsforskning/forskrap2013-5.pdf

demand and price observations available for each day during 2012 to be considered. Forty three stations of Preem did not meet this requirement. Therefore, Preem has been removed from the analysis.

In addition, we have information on the distance of each station to their nearest competitor. For the analysis that follows, an essential variable is the daily Rotterdam spot prices. Throughout this paper, "input prices" will refer to these Rotterdam spot prices unless otherwise indicated. The volume-adjusted prices are computed at the national, regional, (self)-service station, and station levels. This price is a function of three components:

1. The station's daily volume of gasoline sold.
2. The station's daily average pump price.
3. Total daily gasoline sold nationally, regionally, across brands, and stations.

For example, the daily volume-adjusted price in Stockholm is computed as the sum of revenue of all Stockholm stations divided by the corresponding total daily gasoline sales. Table 1 lists the descriptive statistics of the variables used in our analysis.

The data was collected to be used in a study of the Swedish gasoline market. The stations were picked by the SCA and the oil companies, to be representative of the six different regions in Sweden (Foros and Steen, 2013b). The regions are "larger cities" (Stockholm, Gothenburg and Malmö—the three largest cities in Sweden), "larger cities" (cities with a population between approximately 30,000 and 80,000),"the E6 highway"[9] and "rural areas" (population below 10,000).

Table 1: Summary statistics at the national level

Variables	**Mean**	**Std. Dev**	**Min**	**Max**
Av. pump prices (SEK/liter)	14.760	.440	13.908	15.730
Vol-adj. prices (SEK/liter sold)	14.727	.435	13.870	15.709
Volume (liter)	818.000	87.690	464.000	1,210.000
Input prices (SEK/liter)	5.390	.359	4.800	6.151
Service (SEK/liter sold)	14.842	.439	14.011	15.819
Self-Service (SEK/liter sold)	14.609	.432	17.701	15.579
≤ 3 (SEK/liter sold)	14.692	.433	13.806	15.672
>3 (SEK/liter sold)	14.748	.438	13.900	15.731

Note: The sample period is January 1, 2012, to December 31, 2012. N=366..

In the sample, national pump prices averaged 14.760 SEK/liter (including taxes), volume-adjusted prices averaged 14.727 SEK/liter sold, while input prices were approximately 9.3

[9]The E6 highway is a part of the international E-road network. It is defined as a separate geographical region as customers who frequently purchase from stations along the highway mostly are highway commuters.

SEK/liter lower. Higher volume-adjusted prices are observed in Stockholm and in rural areas and lower in Gothenburg. On average, prices in the service market differ from the automated market prices by approximately 0.2 SEK/liter sold. A smaller difference in prices is observed between stations with the distance to their nearest competitor within 3 km and more than 3 km. The daily average observed pump prices, volume-adjusted prices, and input prices from January 1, 2012 to December 31, 2012 have been plotted in Figure 1.

Figure 1: Gasoline and input price patterns for average and weighted pump prices in Sweden (2012)

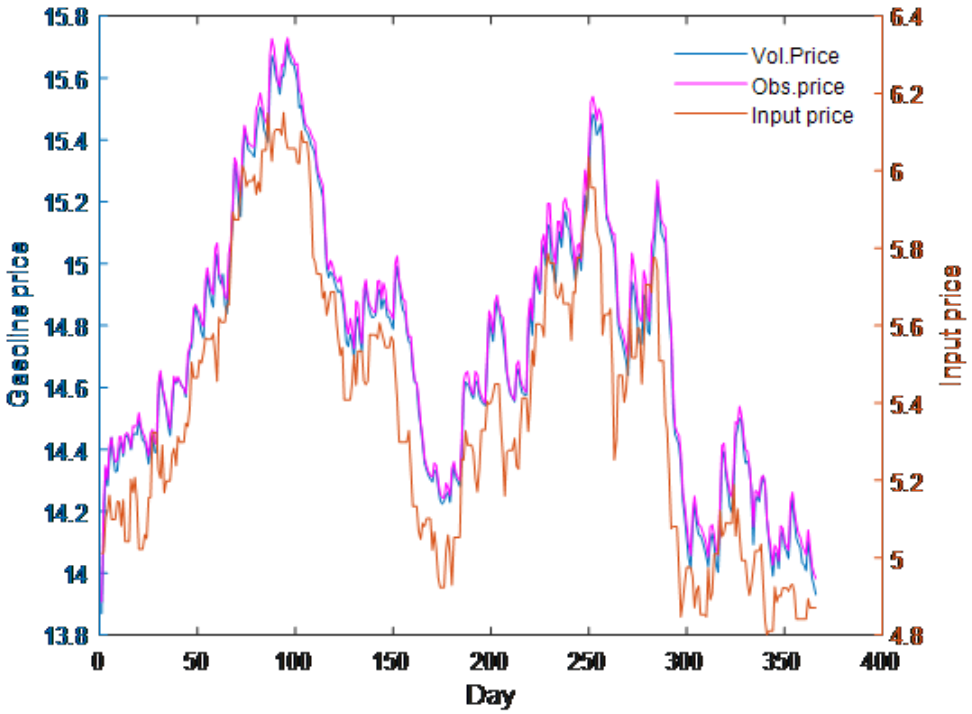

As can be seen, nationally, the daily fluctuations in the pump prices and volume-adjusted prices closely follow the cost movements. The vertical axes to the left contain the gasoline price values, while the vertical axes to the right contain the input price values. The volume-adjusted prices seem to be lower than the observed pump prices, but both exhibit a similar seasonal pattern as the input prices. Following the time series from left to right, we can see that the vertical values' length is relatively compressed when prices rise and somewhat stretched when the prices fall, implying rigidity in prices. Nevertheless, the compression is less pronounced than the stretching, indicating an asymmetric behavior.

3.4 Empirical strategy and methods used

To the extent that variables are non-stationary, the ARDL and NARDL models constitute a very suitable modeling tool since they accommodate both I(0) and I(1) variables. The alter-

native to ECM modeling is to look at the first differences of the variables, as non-stationary I(1) then become stationary, and both asymptotic properties and standard statistical tests can be used. However, one loses all long-run information from the level-form of the variables, and one must be content with estimating only short-run asymmetry. In this study, we want to evaluate both short- and long-run effects of input prices on gasoline prices and cannot just look at first-differenced variables. Therefore, the NARDL model has been used. This estimation method precisely assumes that the dependent variables' response to increases (+) and decreases (-) of the independent variable is asymmetric. Moreover, the NARDL model, as an extension of the ARDL model proposed by Pesaran and Shin (1998) and Pesaran et al. (2001), can entrain a mix of I(O) and I(1) regressors, and the long-run terms are specified. However, unlike the ARDL model, the NARDL model accommodates the decomposing of the regressor into its positive and negative changes. Hence input prices (IP_t) is decomposed into its positive and negative partial sums for
Increases:

$$IP_t^+ = \sum_{j=1}^{t} \Delta IP_j^+ = \sum_{j=1}^{t} max(\Delta IP_j, 0)$$

and decreases in (IP_t)

$$IP_t^- = \sum_{j=1}^{t} \Delta IP_j^- = \sum_{j=1}^{t} min(\Delta IP_j, 0)$$

where the partial sum of positive change in IP is denoted by IP_j^+ and the partial sum of negative change in IP, denoted by IP_j^-. Both IP_j^+ and IP_j^- are included as separate regressors in the model.
The non-linear model for our study takes the following form:

$$\Delta VP_t = \beta_0 + \sum_{i=1}^{p-1} \gamma_i \Delta VP_t - i + \underbrace{\sum_{i=0}^{q} \delta_i^+ \Delta IP_{t-i}^+ \sum_{i=0}^{q} \delta_i^- \Delta IP_{t-i}^-}_{\text{Short-run}} + \theta VP_{t-1} + \underbrace{\varphi^+ IP_{t-1}^+ + \varphi^- IP_{t-1}^-}_{\text{Long-run}} + v_t \tag{1}$$

$$\Delta OP_t = \beta_0 + \sum_{i=1}^{p-1} \gamma_i \Delta OP_t - i + \underbrace{\sum_{i=0}^{q} \delta_i^+ \Delta IP_{t-i}^+ \sum_{i=0}^{q} \delta_i^- \Delta IP_{t-i}^-}_{\text{Short-run}} + \theta OP_{t-1} + \underbrace{\varphi^+ IP_{t-1}^+ + \varphi^- IP_{t-1}^-}_{\text{Long-run}} + v_t \tag{2}$$

where, VP_t is the daily volume-adjusted gasoline price, OP_t is the daily observed pump price, IP_t is the daily Rotterdam spot price, and v_t is an i.i.d. process with zero mean and constant variance. The symbol delta (Δ) denotes first differences, while p, and q are the lag orders of the dependent and independent variables, respectively.

The optimal number of lags for p, and q is obtained using the Akaike information criteria. We then test for a cointegrating relationship between the variables VP_t OP_t , IP_t^+, and IP_t^- using the following NARDL bounds-tests (Banerjee et al., 1998; Pesaran et al., 2001) for an asymmetric long-run cointegration:

$$\begin{aligned} &\text{F-test:} && H_0 : \theta = \varphi^+ = \varphi^- = 0, \text{ and } H_A : \theta \neq \varphi^+ \neq \varphi^- \neq 0 \\ &\text{t-test:} && H_0 : \theta = 0, \text{ and } H_A : \theta < 0 \end{aligned}$$

The two tests are denoted as F_{PSS} and t_{BDM} (Shin et al., 2013). The critical value bounds for both depend on the number of regressors, k. In the situation of long run asymmetry, where we have VP_t (OP_t), IP_t^+, and IP_t^- , the value of k ranges between 1 and 2. If we reject H_0 (of no cointegration), we conclude that the variables are cointegrated in the presence of asymmetry. In this context, this means that one would expect that the gasoline price and the input price levels follow one-to-one in the long-run and thus be cointegrated, while they may differ from each other, and partly strongly in the short-run.

Next, we calculate the NARDL long-run level asymmetric coefficients following Shin et al. (2013). The long-run asymmetric effects of IP_t on VP_t (OP_t) are calculated by dividing the negative of the coefficient of IP_{t-1}^+ (φ^+) by the coefficient of $VP_{t-1}(\theta) : L_{IP}^+ = \frac{-\varphi^+}{\theta}$; and, likewise, by dividing the negative of the coefficient of $IP_{t-1}^-(\varphi^-)$ by the coefficient of $VP_{t-1}(\theta) : L_{IP}^- = \frac{-\varphi^-}{\theta}$.

Then, if a long-run relationship exists (bounds test), we proceed to test using a Wald test whether the difference in the long-run asymmetric coefficients is statistically significant:

$$H_0 : \frac{-\varphi^+}{\theta} = \frac{-\varphi^-}{\theta} \qquad\qquad H_A : \frac{-\varphi^+}{\theta} \neq \frac{-\varphi^-}{\theta}$$

If we reject H_0, it means we have long-run asymmetry. In other words, the magnitude of the change in VP_t (OP_t) when IP_t increases is not the same as when IP_t decreases.

The next step is to test the following hypotheses for short-run asymmetry, again by employing a Wald test.

$$H_0 : \sum_{i=0}^{q} \delta_i^+ = \sum_{i=0}^{q} \delta_i^- \qquad\qquad H_A : \sum_{i=0}^{q} \delta_i^+ \neq \sum_{i=0}^{q} \delta_i^-$$

Again, if the Wald test for the equality of the sum of (+) and (-) lags of IP is rejected, we conclude that the impact of IP_t on VP_t (OP_t) is asymmetric.
If there are no long- and short-run asymmetries, then model (1) is reduced to the traditional ECM model. In case only short- or long-run asymmetry is detected, then the model (1) is transformed to the NARDL with short-run asymmetry (model 3) and the NARDL with long-run asymmetry (model 4) below:

$$\Delta VP_t = \beta_0 + \sum_{i=1}^{p-1} \gamma_i \Delta VP_t - i + \sum_{i=0}^{q} \delta_i^+ \Delta IP_{t-i}^+ \sum_{i=0}^{q} \delta_i^- \Delta IP_{t-i}^- + \theta VP_{t-1} + \varphi IP_{t-1} + v_t \quad (3)$$

$$\Delta VP_t = \beta_0 + \sum_{i=1}^{p-1} \gamma_i \Delta VP_t - i + \sum_{i=0}^{q} \delta_i \Delta IP_{t-i} + \theta VP_{t-1} + \varphi^+ IP_{t-1}^+ + \varphi^- IP_{t-1}^- + \upsilon_t \quad (4)$$

In the same way, we have constructed both models for ΔOP_t. Finally, we have calculated the asymmetric dynamic multipliers in order to see how VP_t (OP_t) adjusts to its new long-run equilibrium following a positive (+) or negative (-) change in IP_t.
The cumulative dynamic multiplier effects of IP_t^+, and IP_t^- on VP_t are evaluated as:

$$m_h^+ = \sum_{j=0}^{h} \frac{\partial VP_{t+j}}{\partial IP_t^+}, m_h^- = \sum_{j=0}^{h} \frac{\partial VP_{t+j}}{\partial IP_t^-} \text{ for } h = 0, 1, 2, 3, ...,$$

In the same way, the asymmetric dynamic multipliers for OP_t are calculated. If $h \to \infty$,then $m_h^+ \to \frac{-\varphi^+}{\theta}$ and $m_h^- \to \frac{-\varphi^-}{\theta}$. Thus, the asymmetric dynamic multipliers will show us the duration of the temporary disequilibria after an input price IP_t change, and also the magnitude of an increase (decrease) in the gasoline price (VP_t, OP_t) due to the positive (negative) change.
In sum, the NARDL framework provides information on the existence of asymmetry, whether it exists in the short- or long-run, and also provides a measure of the accumulated effect of the asymmetry.

3.5 Empirical results

Models (1), (3), and (4) were used first for national-level data. Here, we aim to see whether the relationship between the daily input prices for gasoline and volume-adjusted retail prices differs from the relationship between the input prices and daily average pump prices typically estimated in the empirical gasoline literature.

3.5.1 Price asymmetry in Sweden using national prices: Average and volume-adjusted prices

We have calculated the daily volume-adjusted prices and average gasoline pump prices nationally. The NARDL model was then used to examine the long- and short-run price reactions due to input price changes.

We start out by uncovering the stationarity properties of our price variables. We apply the Augmented Dickey-Fuller (1979) test on all our variables, and we found the presence of a unit root at the 1% significance level. All variables become stationary after applying the first differences. Table 1A in the Appendix shows the results of the Augmented Dickey-Fuller test. As described in the methodology section above, we started with a general specification of the NARDL model, and through appropriate testing, end with one final model. Table 2 below presents our preferred NARDL model findings for the national-level analysis. Tables 2A and 3A in the Appendix show all the four NARDL models estimated for the two retail prices. Column 2 in Table 2 shows the results using the volume-adjusted price, and the results for the average pump prices are shown in column 4.

Table 2: Estimation results in Sweden as suggested by the NARDL models

Method:	**NARDL with**	**LR Symmetry**	**and SR Asymmetry**
Variables		**Variables**	
Vol.-adj. prices	Coefficients	Pump prices	Coefficients
Constant	2.669 ***(0.215)	Constant	3.029 ***(0.226)
VP_{t-1}	-0.325 ***(0.025)	OP_{t-1}	0.371 ***(0.026)
IP_{t-1}	0.392 ***(0.030)	IP_{t-1}	0.451 ***(0.032)
ΔVP_{t-1}	0.347 ***(0.043)	ΔOP_{t-1}	0.285 ***(0.042)
ΔVP_{t-2}	-0.238 ***(0.043)	ΔOP_{t-2}	-0.185 ***(0.041)
ΔVP_{t-3}	0.147 ***(0.041)	ΔOP_{t-3}	0.127 ***(0.039)
ΔIP_t^+	0.238 ***(0.070)	ΔIP_t^+	0.249 ***(0.073)
ΔIP_{t-2}^+	-0.191 ***(0.072)	ΔIP_{t-2}^+	-0.310 ***(0.075)
ΔIP_t^-	-0.148 **(0.075)	ΔIP_t^+	-0.226 ***(0.071)
ΔIP_{t-1}^-	-0.212 ***(0.079)	ΔIP_{t-1}^-	-0.244 ***(0.084)
LR_{IP}	1.205 ***(0.023)	LR_{IP}	1.225 ***(0.021)
λ	-0.325 ***(0.025)	φ	-0.371 ***(0.026)
R^2	0.485	R^2	0.491
t_{BDM}	-13.633 ***	t_{BDM}	-13.397 ***
F_{PSS}	85.188 ***	F_{PSS}	92.717 ***
IP_{SR}	0.407 ***(0.151)	IP_{SR}	0.409 ***(0.158)

Note: Standard errors in in parentheses, and $^*p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

The results suggest a potential long-run relationship between input prices and retail, volume-adjusted, and pump prices. Therefore, as a next step, we have tested our variables for cointegration through a bounds-testing process. The cointegration tests tBDM and FPSS are 12.433 and 52.559, -13.397 and 92.717, respectively, for both cases, and the null hypothesis for no cointegration is rejected at a 1% significance level. Thus, we have evidence of cointegration between volume-adjusted prices and input prices when long- and short-run asymmetry is taken into account. We find a long-run relationship between the average pump prices and input prices for gasoline when only short-run asymmetry is considered. The positive and negative long-run coefficients are 1.187 and 1.193, respectively, and are highly significant. They indicate that in the long-run, a 1% price increase (decrease) of input prices causes a 1.187% (1.193%) increase (decrease) in the volume-adjusted prices. The long-run coefficient in the case of observed prices (column four) (LR_{IP}) is 1.225 and is highly significant, implying that a 1% price increase (decrease) of input prices results in the rise (fall) of pump prices by 1.225% in the long run, ceteris paribus. The error correction λ has the correct negative sign and is statistically significant in both cases, suggesting a somewhat higher adjustment speed for the average pump prices. The Wald tests for short run (IP_{SR}) asymmetry in columns 2 and 4 implies that the short-run asymmetry responses

are present when we consider both volume-adjusted and average prices. However, in column 2, we find some weak evidence on a 10%-level and also for the long-run (IP_{LR}) asymmetry for the volume-weighted retail prices (see column 5, Table A2: IP_{LR}=-0.007).[10] Figure 2 below presents the asymmetry from the dynamic multipliers for both cases, volume-adjusted prices, and pump prices. Specifically, Figure 2 highlights the differences in the asymmetry between the volume-adjusted and average pump prices.

Figure 2: The asymmetry + C.I. for NARDL

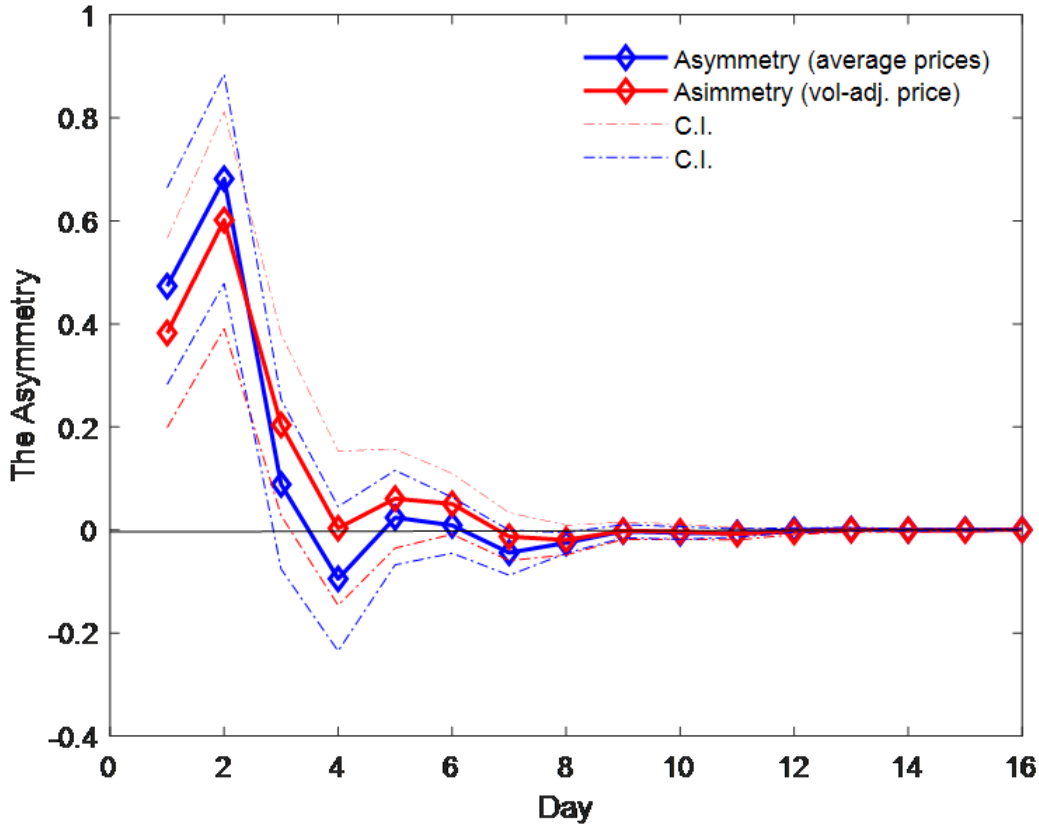

The solid red and blue lines in the middle are the asymmetry plots for volume-adjusted and pump prices, respectively. They reflect the difference between the dynamic multipliers of positive and negative changes in IP. The asymmetry lines lie within the upper and the lower bounds of the 95% confidence interval. Thus, if the zero line falls within the confidence interval boundary, there is no asymmetry because it will not be statistically significant.

Figure 2 shows that in the short- run, an increase due to a positive change in input prices is more significant than the magnitude of decrease due to a negative change. There is a cumulative positive and statistically significant effect of input price on gasoline prices from

[10]The models pass the tests of stability and serial correlation. Besides, the NARDL model with long- and short-run asymmetry also passes the normality test.

the first four days, indicating asymmetry in favor of positive cost innovations. This difference mirrors the observed short-run asymmetry that we already determined from the Wald tests where IP_{SR} for volume-weighted prices, and the average pump prices are nearly similar (0.407 vs. 0.409). Thus, the asymmetry line with its confidence interval moves away quite distinctly from the zero lines in the early stages compared to the later stages. Similarly, the thick blue line shows that the magnitude of the increase of the observed prices due to a positive change is larger in the short-run. Another important feature of the NARDL estimates (graph in Figure 2) is the time it takes before the asymmetry in the short-run disappears, and we are back to a symmetric equilibrium. Both dynamic multipliers show strong temporal pricing asymmetry, but the adjustment speed toward long-run equilibrium is relatively slower for the volume-adjusted prices than the average prices. Actually, while the average pump price returns to a long-run equilibrium after two days, it takes three days for the same to happen with the volume-weighted price (see Table A2). Hence, in line with the adjustment speed parameters, it takes volume-weighted prices longer to return to equilibrium. This result suggests that to the extent the Swedish firms manage to impose asymmetric pass-through costs, they are more concerned with determining prices on stations or days with higher sales. Estimating asymmetry based on average prices (as has been common in the literature) underestimates the degree of asymmetry. In the following, we will use the volume-adjusted prices as a dependent variable when we analyze price asymmetry heterogeneity across regions, service levels, and spatial competition levels.

3.5.2 Regional heterogeneity: Price asymmetry across regions

As described in the data section, we have available data for six Swedish regions that are anticipated to differ regarding supply (e.g., number and density of stations) and demand (driving pattern, income, and consumer density). Table 3 displays the station distribution across regions and Figure 3 the volume-adjusted prices.

Table 3: Station distribution by brands across regions

Regions/Brands	**Statoil**	**OK-Q8**	**Shell**	**Jet**	**St1**	**Total**
Stockholm	5	5	3	4	5	22
Gothenburg	5	2	4	5	3	20
Malmo	5	5	4	4	5	23
E6	4	4	5	5	5	23
Smaller cities	6	4	4	6	6	26
Rural areas	12	4	2	4	12	33
Total	37	24	22	28	36	147

Figure 3: Volume-adjusted prices across regions in Sweden in 2012

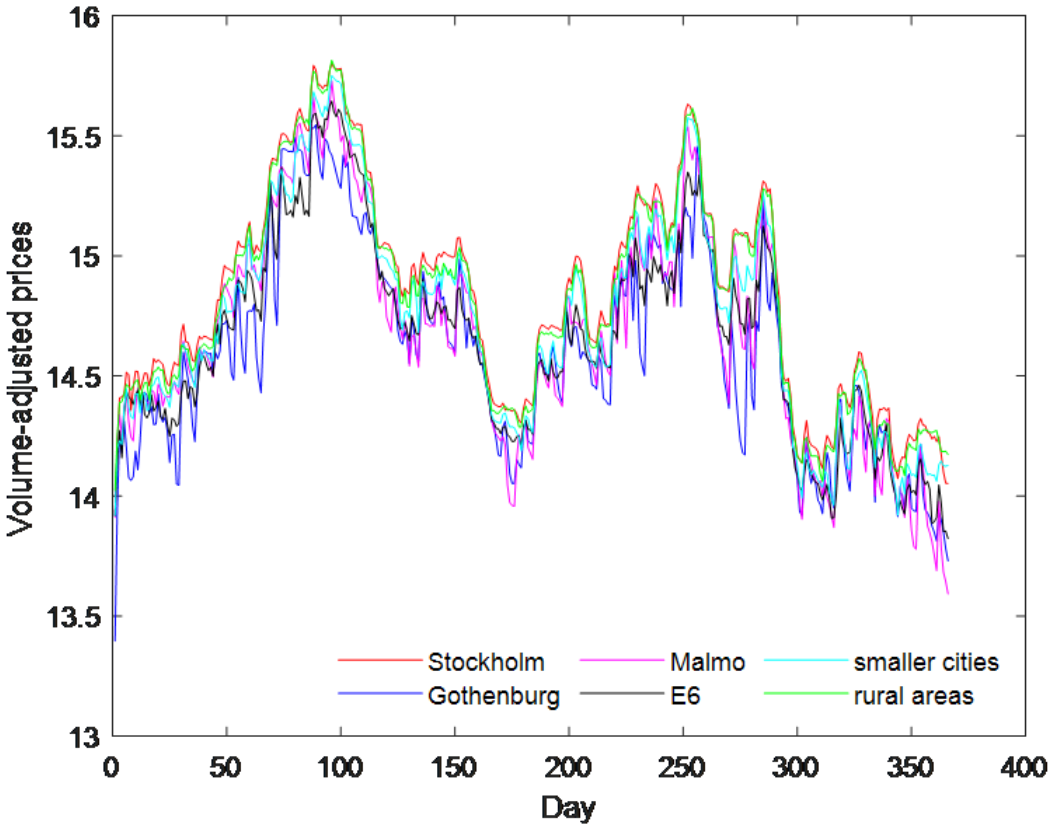

We found the highest prices in rural areas and in Stockholm. The latter is surprising, given Stockholm's status as both the largest city and the capital of Sweden. The former is less surprising as the stations in rural areas are much less exposed to competition from nearby stations. This is also in line with what Foros and Steen (2013b) observed. Further, smaller cities exhibit higher prices. The lowest prices are found in Gothenburg and Malmö, which are number two and three cities in Sweden, respectively. Table 4 shows more detailed descriptive statistics, showing much the same picture. Regarding quantities, the largest volumes were found in Stockholm and along the E6 highway. Table 5 presents the results of asymmetric responses of volume-adjusted prices to input price changes across regions.[11] The evidence of a long-run relationship between input and retail prices is found for cities, highways, and rural areas. The Wald tests document that there is an asymmetric behavior in gasoline prices' short-run responses to input prices across all regions.

We also find some evidence of long-run asymmetry in Malmö, the E6 highway, and rural areas, but as above, the numbers are very close to zero, ranging from -0.02 to 0.007 (see Table

[11]In Table 5, we only include our preferred models, the underlying 3x6 NARDL models are not included here but the results of the NARDL models with LR and SR asymmetry across regions is presented in Table A4 in the Appendix.

Table 4: Descriptive statistics regional average weighted gasoline prices and average station quantities

(a) Average station prices (b) Average station quantities

Variables **Regions**	**Mean**	**Std. Dev.**	**Min**	**Max**	**Mean**	**Std. Dev.**	**Min**	**Max**	**Obs.**
Stockholm	14.841	.445	13.929	15.802	7165	3730.4	21.540	23 701	8052
Gothenburg	14.600	.428	13.396	15.546	6586	4448.2	564.840	28 242	7320
Malmö	14.666	.463	13.593	15.73	5629	3299.4	36.120	20 568	8418
E6	14.654	.422	13.823	15.645	6790	4475.3	26.650	29 834	8418
Smaller c.	14.741	.448	13.919	15.751	6018	3343.9	54.520	21 609	9516
Rural areas	14.817	.446	14.024	15.815	3022	2078.0	96.530	17 505	12 078

Note: This table shows descriptive statistics on a sample of the Swedish gasoline market during 2012. Panel A presents the statistics of average station prices (SEK/liter sold). In Panel B, the average quantity related to stations is provided. Statistics are provided separately for six different geographical region samples. The number of observations (N) is offered by day.

A4). We find negative and highly significant adjustment speed in all areas. The adjustment speed in Malmö is the highest (-0.361), whereas rural regions and Stockholm appear to have lower adjustment speeds (-0.283 to -0.286). The difference between rural areas and Malmö is not very surprising and is in line with the expectation that in rural areas with potentially less competition, asymmetry can last longer. The low adjustment speed for Stockholm is more surprising but is in line with the fact that the price dynamics and price levels in rural areas and Stockholm are very similar (see also Foros and Steen, 2013b). Stockholm also stands out because the accumulated short-run effect is negative (-0.308) and, in common with the smaller cities and the rural areas, lower in magnitude than the others.

The short-run asymmetry as measured through IP_{SR} is varying, but is now clearly the lowest for rural areas (0.236) versus Stockholm where IP_{SR} is much the same as for the E6 highway and smaller cities (between 0.41 and 0.51). Again, the other large cities Gothenburg and Malmö, stand out with IP_{SR} around 0.84–0.85. Figure 4 contains the corresponding dynamic multipliers for each area. The graphs mirror to a large extent what we see in Table 5. In Gothenburg, gasoline prices are back on to long-run symmetry already after one day. Malmö, smaller cities, and on the E6 highway, gasoline prices take three to four days to reach the long-run equilibrium. In contrast, this process takes six days for both Stockholm and rural areas, and the dynamics magnitude is much lower. Additionally, the amplitude of the dynamics differs. In Gothenburg and Malmö, the firms can impose much higher asymmetric pass-through, double as high as what we saw for the national prices above. IP_{SR} is estimated to be in the range of 0.84 to 0.85, whereas the national average (weighted prices) was 0.41. Again, the amplitude of the pass-through dynamics in Figure 4 reveals a

much lower amplitude for rural areas, smaller cities, and Stockholm.

The finding of a cumulative positive and statistically significant effect of input price changes on volume-adjusted prices, lasting somewhat longer in Stockholm and rural areas, might suggest less competition in Stockholm than other leading cities. Stockholm takes almost a whole week to return to a symmetric equilibrium as compared to, in the extreme, one day in Gothenburg. From Table 4, we can see that this also corresponds to the price levels observed. Stockholm has the highest average price of all six regions (14.841), whereas Gothenburg has the lowest (14.60).

Table 5: Estimation results across local markets

Regions:	**Stockholm**	**Gothenburg**	**Malmö**	**E6**	**smaller c**	**rural areas**
Method:	LR Symmetry&	LR Symmetry &	LR Symmetry &	LR Symmetry &	LR Symmetry &	LR &
NARDL with	SR Asymmetry	SR Asymmetry	SR Asymmetry	SR Asymmetry	SR Asymmetry	SR Asymmetry
Variables	Coefficients	Coefficients	Coefficients	Coefficients	Coefficients	Coefficients
Constant	2.331***	2.653***	2.697***	2.468***	2.493***	2.307***
	(0.154)	(0.310)	(0.274)	(0.408)	(0.181)	(0.149)
VP_{t-1}	-0.285***	-0.315***	-0.345***	-0.299***	-0.309***	-0.281***
	(0.018)	(0.034)	(0.032)	(0.024)	(0.021)	(0.017)
IP_{t-1}	0.350***	0.358***	0.435***	0.353***	0.381***	0.344***
	(0.022)	(0.039)	(0.041)	(0.028)	(0.026)	(0.021)
ΔVP_{t-1}	0.361***	0.264***	0.283***	0.098**	0.345***	0.349***
	(0.040)	(0.044)	(0.046)	(0.043)	(0.043)	(0.042)
ΔVP_{t-2}	-0.164***	-0.214***	-0.191***	-0.109***	-0.151***	-0.116**
	(0.043)	(0.045)	(0.046)	(0.043)	(0.043)	(0.043)
ΔVP_{t-3}	0.138***		0.126***		0.154***	0.173***
	(0.039)		(0.045)		(0.041)	(0.039)
ΔVP_{t-7}		0.112***			0.106***	
		(0.042)			(0.039)	
ΔVP_{t-17}						0.071**
						(0.035)
ΔIP_t^+	0.142***	0.263**	0.448***	0.229***	0.211***	0.114**
	(0.057)	(0.133)	(0.122)	(0.088)	(0.074)	(0.055)
ΔIP_{t-2}^+		-0.410***	-0.273**		-0.292***	-0.121**
		(0.138)	(0.121)		(0.076)	(0.056)
ΔIP_{t-4}^+		-0.223*				
		(0.135)				
ΔIP_{t-6}^+	-0.164**	0.486***				
	(0.055)	(0.140)				
ΔIP_{t-9}^+						0.090*
						(0.052)
ΔIP_{t-13}^+						0.113**
						(0.053)
ΔIP_{t-15}^+						-0.101**
						(0.052)
ΔIP_t^-	-0.155***		-0.355***	-0.285***	-0.173***	-0.142***
	(0.055)		(0.119)	(0.086)	(0.073)	(0.053)
ΔIP_{t-1}^-		-0.432***	-0.322***			
		(0.140)	(0.131)			
ΔIP_{t-6}^-		-0.295**			-0.136**	
		(0.136)			(0.071)	
ΔIP_{t-6}^-					-0.182***	
					(0.071)	
LR_{IP}	1.228***	1.136***	1.260***	1.179***	1.233***	1.223***
	(0.021)	(0.047)	(0.038)	(0.032)	(0.025)	(0.021)
λ	-0.284***	-0.306***	-0.345***	-0.335***	-0.310***	-0.271***
	(0.017)	(0.034)	(0.033)	(0.028)	(0.021)	(0.025)
R^2	0.550	0.391	0.369	0.343	0.499	0.574
t_{BDM}	-15.873***	-9.364***	-11.116***	-12.914***	-14.680***	-16.857***
F_{PSS}	128.387***	44.423***	58.351***	81.495***	110.287***	133.961***
IP_{SR}	0.424***	0.844***	0.853***	0.514***	0.409***	0.236*
	(0.103)	(0.349)	(0.257)	(0.138)	(0.171)	(0.132)

Note: Standard errors in in parentheses, and $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Figure 4: Asymmetric dynamic multipliers across regions

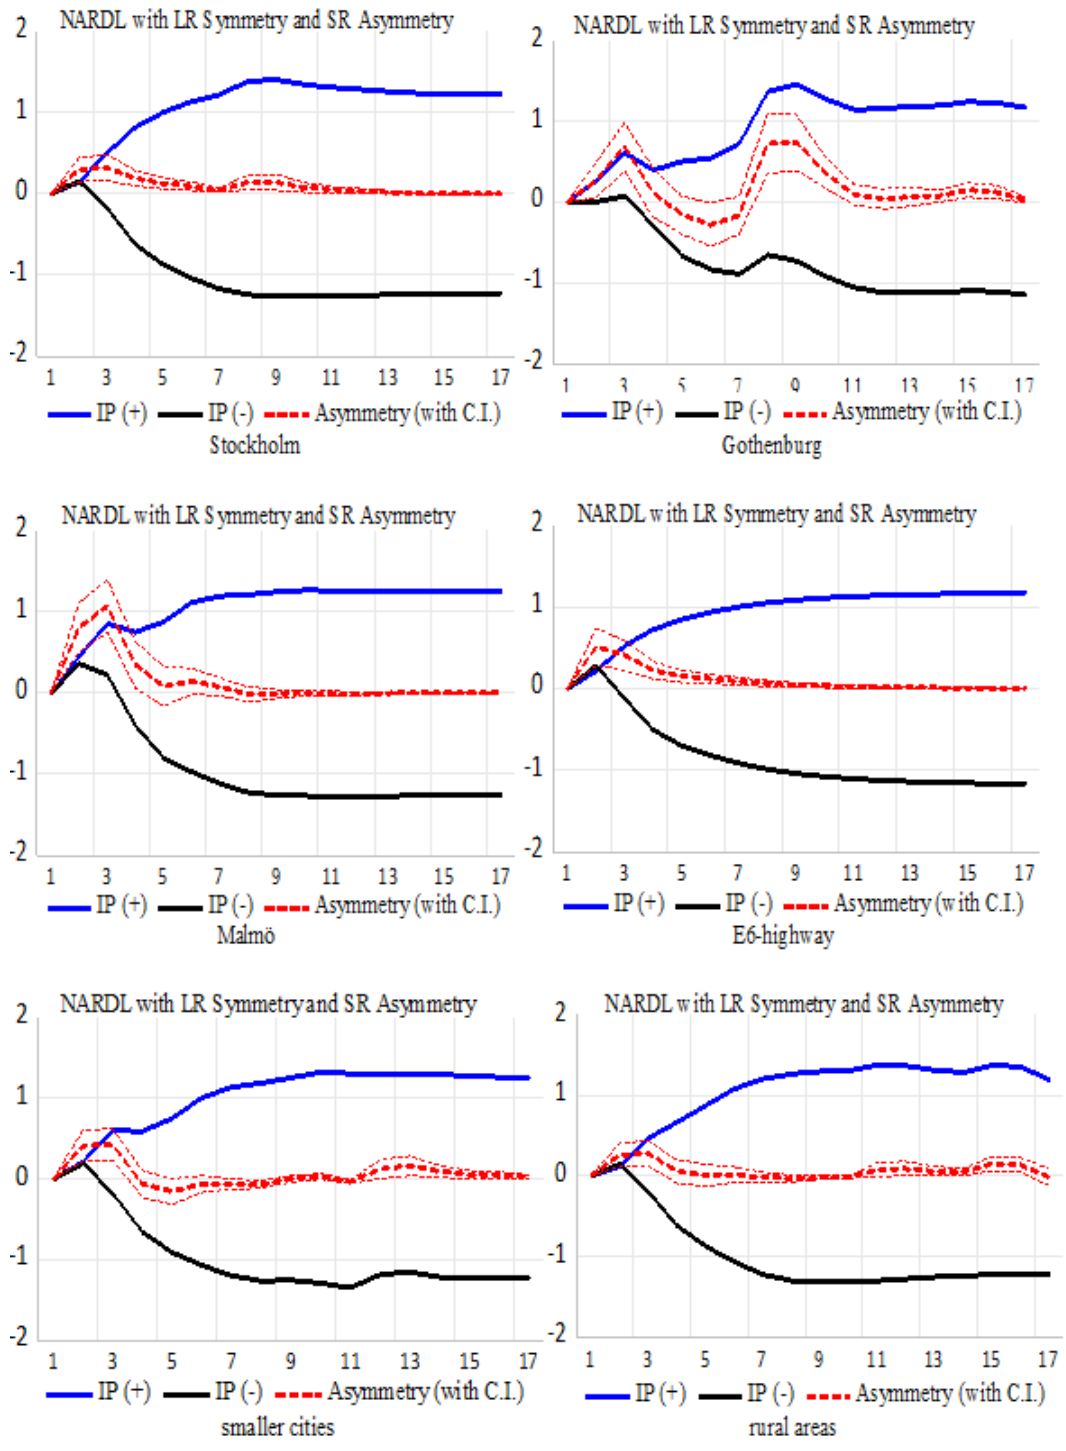

One explanation for Stockholm could have been the higher station volumes, but also Gothenburg has high average station volumes (refer to Table 4). A natural question is also whether the Swedish market exhibits Edgeworth-like cycles and saw-tooth patterns. This was explicitly analyzed by Foros and Steen (2013b), using the same dataset as we do here,

and they found no sign of this. Indeed, they found very little evidence of price cycling. However, to the extent that they uncovered (very moderate) price cycles, these were most pronounced in markets that have the lowest prices, that is, Gothenburg and Malmö, though they did not resemble saw-tooth-like patterns.

The only study besides us that has estimated asymmetry in pass-through is Byrne (2019). Byrne (2019), using data from Canada, found more asymmetry in rural markets than in cities. If we disregard the Stockholm result, we find evidence of the same pattern, estimating asymmetry across different markets.

Consistent with what we also found at the national-level analysis. A temporal cumulative positive and statistically significant effect of input price innovations on gasoline prices in the short-run is evident in six different regional markets, though close to zero for Gothenburg, indicating the existence of asymmetry in favor of positive cost changes. In other words, retailers are quick to raise and slow to reduce prices when an input change occurs.

3.5.3 Station heterogeneity: Service and self-service gasoline stations

From what we saw above, and in line with Byrne (2019), market power seems to be associated with a more asymmetric cost pass-through. Using the same data, while also including the Preem stations, Nguyen-Ones and Steen (2021) find evidence of higher markups for service stations as compared to self-service stations. Hence, in our sample, we have chosen to calculate two price series: one representing the daily volume-adjusted gasoline prices for service stations and the other representing price series for the automated stations. We then test for the presence of asymmetric behavior of volume-adjusted prices to input price deviations for these two station groups. To our knowledge, we are the first to provide estimates for the input-volume-adjusted price relationship from these two station types separately. Table 6 shows the results. The variables are cointegrated in both markets, and the adjustment speeds λ are statistically significant and have the correct sign. We observe a slower speed of adjustment in the service market than in the self-service market. The subsequent step is to implement the Wald tests for potential long- or short-run asymmetries. The Wald tests accept long-run symmetry with respect to input price, indicating the pass-through from input prices to the volume-adjusted price of gasoline in the long-run. The short-run symmetry is rejected at a 1% significance level, which signifies asymmetric pass-through in the short-run. The long-run coefficient LR_{IP} is statistically significant at a 1% significance level and has a positive causal effect on VP. In Figure 5, the temporal asymmetric behavior can be observed.

IP_{SR} is statistically significant for both gasoline station groups, but in line with the higher adjustment speed, somewhat higher for the self-service stations. However, we find that short-run asymmetry, in favor of positive responses, results in both markets. This indicates that similar to the national and regional-level analysis, retailers are quick to increase and slow to decrease gasoline prices when an input change happens.

Interestingly, we find some support for differences correlated with the degree of market power. We find that manned service retailers adjust the asymmetry first after four days, whereas, at the automated self-service stations, the asymmetric cost pass-through only lasts for three days (see Table 6). Several studies associate high service levels with higher prices or more market power. Eckert and West (2005) find evidence that station characteristics affect sellers' price setting and suggest the presence of imperfect competition. Haucap et al. (2017) conclude that prices are positively related to station service levels, and Shepard (1991) finds that stations charge a full-service markup. Nguyen-Ones and Steen (2021) confirm higher market power for manned service stations in Sweden. From our asymmetry result, and in line with this, we find that the oil companies are better able to extract market power through asymmetric pass-through of costs for a longer duration for these stations.

Table 6: Estimation results across self-service and service gasoline stations

Market:	**Service stations**	**Self-service stations**
Method:	NARDL with SR	NARDL with SR
Variables	Coefficients	Coefficients
Constant	2.806*** (0.211)	3.028*** (0.042)
VP_{t-1}	-0.339*** (0.024)	-0.373*** (0.028)
IP_{t-1}	0.413*** (0.029)	0.449*** (0.034)
ΔVP_{t-1}	0.372*** (0.043)	0.183*** (0.042)
ΔVP_{t-2}	-0.239*** (0.043)	-0.132*** (0.040)
ΔVP_{t-3}	0.159*** (0.041)	-
ΔIP_t^+	0.168*** (0.067)	0.339*** (0.083)
ΔIP_{t-2}^+	-0.221*** (0.071)	-0.307*** (0.084)
ΔIP_t^-	-0.135** (0.067)	-0.259*** (0.081)
ΔIP_{t-1}^-	-0.266*** (0.077)	-0.245*** (0.095)
LR_{IP}	1.217*** (0.021)	1.201*** (0.023)
λ	-0.339*** (0.024)	-0.372*** (0.029)
R^2	0.505	0.429
t_{BDM}	-13.825***	-13.116***
F_{PSS}	96.939***	87.029***
IP_{SR}	0.348*** [0.018]	0.537*** [0.003]

*Note: Standard errors in (), p-value in [], and *p<0.10, **p<0.05, ***p<0.01.*

Figure 5: Asymmetric dynamic multipliers for service and self-service gasoline stations

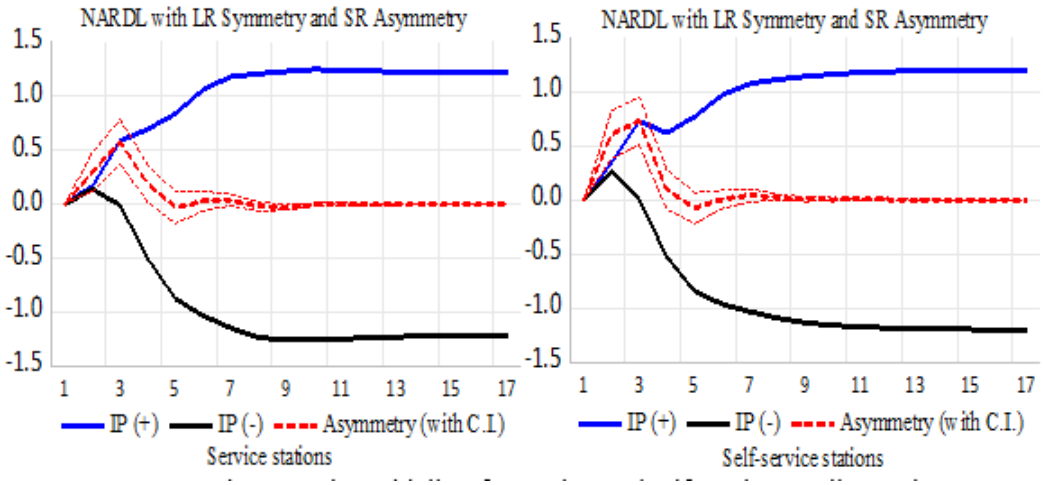

3.5.4 Results for stations with a different degree of local competition

As discussed above, there exists a large literature base discussing the effects of local competition on retail gasoline markets (see, e.g., Hastings, 2004; Slade, 1987; Netz and Taylor, 2002; Nguyen-Ones and Steen, 2021). Nguyen-Ones and Steen (2021), estimating a Bresnahan-Lau (1982) structural model for the Swedish market, find markup results suggesting that differentiation in terms of location has a significant impact on a stations' level of market power in Sweden. The increased density of stations reduces each sellers' markup and resulting in a positive effect on the local competition. These findings are in line with those of Barron et al. (2004), Barron et al. (2008), and Clemenz and Gugler (2006), who show that higher station density tends to lower average prices suggesting that a higher number of sellers raises local competition.

To scrutinize this, we analyze how a distance to the nearest competitor affects pricing asymmetry. We calculated the daily volume-adjusted prices for all stations with their closest competitor within a span of 3 km and volume-weighted prices for those that have their closest competitor more than 3 km away.

Table 7 contains the estimates of these two cases. The volume-adjusted prices and input prices are cointegrated, suggesting a long-run relationship between these two variables. We find only short-run asymmetry for both relationships, for both those having a close competitor and those that do not. Thus, we have used the NARDL model, which accounts only for short-run asymmetric responses.

Table 7: Estimation results of stations based on the distance to their nearest competitor

Distance:	**Stations** $\leq$ **3km**	**Stations** $>$ **3km**
Method:	NARDL with LR& SR	NARDL with SR
Variables	Coefficients	Coefficients
Constant	2.913*** (0.233)	2.458*** (0.174)
VP_{t-1}	-0.356*** (0.027)	-0.302*** (0.020)
IP_{t-1}	0.429*** (0.033)	0.376*** (0.025)
ΔVP_{t-1}	0.295*** (0.043)	0.216*** (0.041)
ΔVP_{t-2}	-0.194*** (0.042)	-0.094*** (0.023)
ΔVP_{t-3}	0.106 *** (0.041)	0.104*** (0.041)
ΔIP_{t}^{+}	0.287*** (0.081)	0.178*** (0.067)
ΔIP_{t-2}^{+}	-0.284*** (0.082)	-
ΔIP_{t}^{-}	-0.209*** (0.076)	-0.179*** (0.064)
ΔIP_{t-1}^{-}	-0.247*** (0.088)	-
LR_{IP}	1.205*** (0.023)	1.247*** (0.024)
λ	-0.356*** (0.026)	-0.302*** (0.019)
R^2	0.463	0.454
t_{BDM}	-13.215***	-14.916***
F_{PSS}	87.542***	113.583***
IP_{SR}	0.461*** [0.007]	0.358*** [0.000]

*Note: Standard errors in (), p-value in [], and *p<0.10, **p<0.05, ***p<0.01.*

We find that if the input price increases (decreases) by 1%, the VP increases (decreases) by about 1.205% and 1.247%.[12] Hence, we observe almost a one-to-one relationship. The Wald tests fail to reject long-run symmetry with respect to input price, indicating a symmetric pass-through from input prices to the retail price of gasoline in the long-run.[13]

In terms of the short-run, the Wald tests decisively reject the null of additive short-run symmetry with respect to input price. This pattern of asymmetry determines the shape of the dynamic multipliers presented in Figure 6. At a 1% confidence interval, we reject the null hypothesis of summative symmetric adjustment, and P_{SR} is equal to 0.461 and 0.358. This indicates that the positive impact of IP^{+} on VP is significantly greater than the positive impact of IP^{-} in the short-run, and mostly so for the gasoline stations with closer local competitors.

Columns (1)–(2) show the negative and statistically significant adjustment speed λ for both markets, whereas anticipated, the adjustment speed is somewhat slower for stations

[12]The long-run level coefficient for both cases is LR_{IP} and the error correction forms are EC= VP-(1.205*IP_{t-1}) and EC= VP-(1.247*IP_{t-1}), respectively.

[13]Both NARDL models passed the test of stability and serial correlation.

having lower local competition. Further, the error correction term provides evidence that being in a market more than 3 km away from the nearest competitor can significantly slow the speed of pass-through asymmetry. The dynamic multipliers in markets with a distance between stations less or equal to 3 km has a rapid return to long-run symmetry after just three days. Stations with more than 3 km to their closest competitor are able to impose short-run asymmetric cost pass-through for as much as seven days (see Table 7).

Figure 6: Asymmetric dynamic multipliers for NARDL

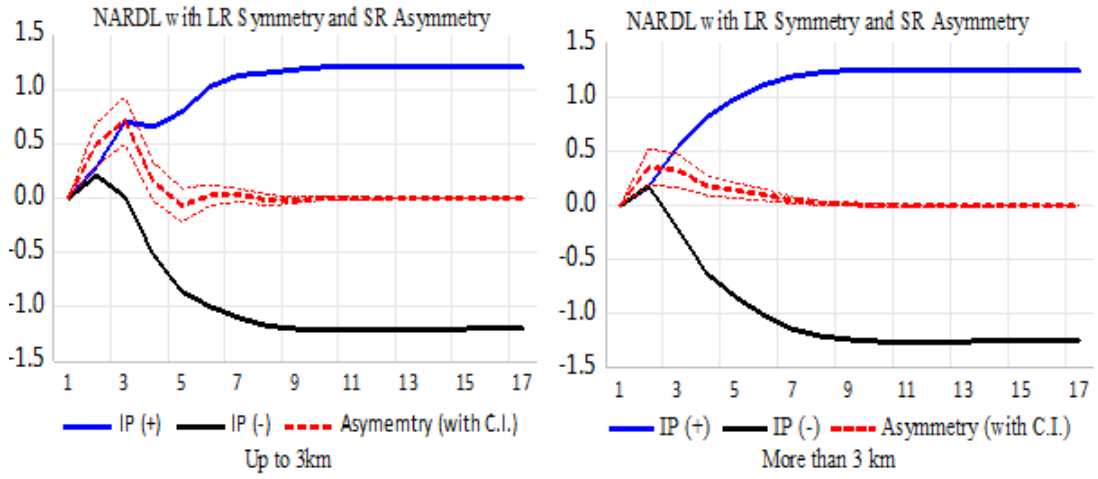

This can be seen clearly from Figure 6, which shows the dynamic multiplier effects. Since the average distance to competing gasoline stations correlates with regions in the sense that, in rural areas, one typically finds a longer distance between competitors as compared to in cities. Interestingly enough, the longest period of short-run asymmetry in cost pass-through is found for rural regions (6 days) and stations with no competitors nearby (7 days). As such, our results are in line with what Byrne (2019) found for the Canadian market.

3.6 Economic implications of asymmetric pass-through of oil prices to gasoline prices

Using the NARDL models, we find clear evidence of short-run price asymmetry both using average pump prices and volume-adjusted prices for the Swedish national market. Prices increase more after input price increases compared to their decrease when input price reduces. Figure 7 shows the difference in asymmetry between the model using the average prices and the model using the volume-weighted prices. The volume-adjusted price asymmetry takes an extra day to correct the disequilibrium error as compared to what we find for the average pump price asymmetry. Hence, this indicates that companies are indeed more focused on pricing at larger gasoline stations.

Figure 7: The Asymmetry + C.I. for NARDL in Sweden

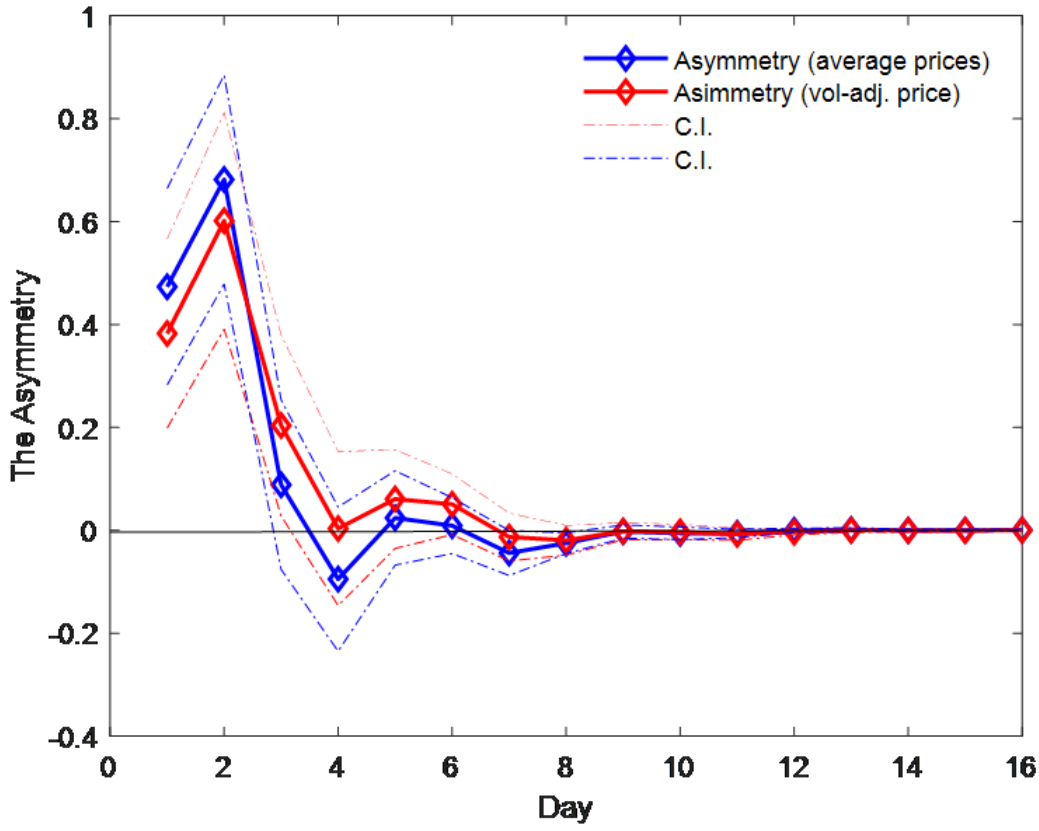

Indeed this also impacts how we interpret the effects of asymmetric pricing behavior. If we consider only average prices, we underestimate the effect of the asymmetry. In our case, for Sweden aggregated, we find one more day of asymmetry. This extra day accounts for a significant part of the overcharge created by the asymmetric pricing behavior. In Table 8, we do a "back-of-the-envelope" calculation of the effects of asymmetry. Based on our estimates from using average prices and the model using volume-weighted prices (see Figure 7), we find an approximately average overcharge for the 2(3) days predicted by the models.[14] Note that the overcharge, the prices, and the margins will differ across these models. Calculating these overcharges, we find that as much as 1,136,982 thousand SEK is the estimated overcharge

[14]Average price and margin increase due to asymmetry on Day 1 (Day 2). For instance, this overcharge can be calculated as asymmetry on Day 1 (Day 2) divided by the daily average price: 0.48/14.76= 3.3% (0.68/14.76= 4.6%), and by the daily average margin 0.48/1.0478=45.8% (0.68/1.0478=64.9%). Then, to calculate the cost of the asymmetric price responses for Day 1 (Day 2) in Sweden, we have used the daily gasoline consumption in 2012 (see Sweden Gasoline consumption—data, chart/TheGlobalEconomy.com) and multiplied it with the asymmetry value: 0.48 (0.68)*8167972= 470,475,187 (666,506,515) SEK. Finally, the annual cost of asymmetry in SEK in Sweden is calculated as the daily cost times 120, the number of days when the gasoline price increases in our sample. Similarly, the calculations are made for the model based on volume-weighted prices, but the only difference is the presence of Day 3 asymmetry and using volume-weighted average price (14.73) and margin (1.0238).

for the average price model.

Table 8: "Back-of-the-envelope" calculations

Models	**Day 1**	**Day 2**	**Day 3**	
1. The average price model				**Total effect SEK**
Price increase due to asymmetry (%)	3,3 %	4,6 %		
Margin increase due to asymmetry (%)	45,8 %	64,9 %		
Annual cost of asymmetry in SEK in Sweden	470 475 187	666 506 515		**1 136 981 702**
2. The vol-weighted price model	**Day 1**	**Day 2**	**Day 3**	**Total effect SEK**
Price increase due to asymmetry (%)	2,6 %	4,2 %	1,4 %	
Margin increase due to asymmetry (%)	38,1%	66,4%	19,5%	
Annual cost of asymmetry in SEK in Sweden	382 261 090	607 697 117	196 031 328	**1 185 989 534**
Difference in overcharge in SEK				**- 49 007 832**
Difference in overcharge in %				**-4,13%**

This increases by 49,008 thousand SEK when we account for the effect of one extra day, and volume-adjusted prices.[15] Hence, it suggests that the model using average prices underestimates the full overcharge due to an asymmetry of 4.1%. In fact, the overcharge amounts to 2.7% of the total expenditure on gasoline in Sweden in 2012. Table 8 also shows, as compared to the margin, this overcharge represents 40–60% of the daily average margin. This is obviously due to the significant taxes that, together with the Rotterdam price, leave the companies with a very low gross margin per liter. Thus, we also calculate the total gross margin for 2012 to compare it to these overcharges. The total margin is around 3 billion SEK; therefore, there is an overcharge of around 1 billion SEK. Hence, the 2.7% increase in total expenditure for gasoline implies a doubling of the gross margin, suggesting that for the oil companies, asymmetric pricing presents a major avenue for increasing profits. Since the last part of the retail price is determined by taxes and the Rotterdam wholesale price, this increase in the margin is not very noticeable for consumer expenditures, who end up paying less than 3% extra for their gasoline.

3.7 Summary and discussion

Given the rich dataset that we had access to from the Swedish retail gasoline market, we have been able to examine whether there is a significant difference between asymmetric average pump price responses and volume-adjusted price responses to input price variations. We observe that when using volume-adjusted prices, the disequilibrium error takes more time to be fully corrected to the long-run equilibrium, suggesting that the oil companies are more concerned with the price setting on days and stations with higher volumes.

[15]Calculated as the difference in the annual cost due to asymmetry in pricing between the model estimated using average prices (SEK 1,236,981,702) and the model using volume-weighted prices (SEK 1,185,989,534).

In the rest of our empirical analysis, we continue using only volume-adjusted prices as the dependent variable. We analyze how asymmetry is related to the degree of competition by estimating models across regions, and considering service-level heterogeneity and the local competition.

The results for our regional models are overall consistent with what Byrne (2019) found. Asymmetric price responses are more pronounced in rural markets than in urban markets, although Stockholm seems to be an exception. When controlling for local competition pressure by dividing stations into those being exposed to nearby competitors within 3 km, and those stations where this is not the case, we find a similar result. The stations not exposed to as much local competition exhibit a similar pattern to what we found for the rural region. Here, the most prolonged short-run asymmetry was concluded among all estimated models.

Additionally, we investigated the presence of asymmetric price responses for service and self-service stations. The results indicate that the pricing asymmetry is larger and more prolonged for service stations. This result, that is, stations with higher service levels have larger pricing asymmetry, supports the results found by Verlinda (2008). It is also in line with the findings of Byrne (2019), that is, evidence suggesting that market power increases cost pass-through asymmetry. Several studies have concluded that a higher service level is associated with higher prices and market power, suggesting that the prolonged asymmetry found here can be attributed to differences in competition level (see, e.g., Eckert and West, 2005; Haucap et al., 2017; Shepard, 1991; Nguyen-Ones and Steen, 2021).

Overall, our results are largely consistent with the existing literature that relates market power with the presence of asymmetric cost pass-through to the retail price of gasoline. Moreover, the analysis confirms that gasoline prices in Sweden during the sample period display the "rocket and feathers" pattern by responding more rapidly to the input price increases than decreases. When input price increased, retailers were able to quickly pass-through the input price increases to volume-adjusted price. Conversely, when input prices decreased, the retailers delayed the pass-through.

As we saw in the section above, the economic impact on retailer's gross margins is substantial, still only affecting average prices and expenditure marginally. In transparent oligopolies, where taxes and exogenous costs make up a significant part of the retail price and overall demand typically is inelastic, incentives to impose asymmetric pricing are very high.

The literature seems to conclude that market power leads to larger pricing asymmetry (see, e.g., Byrne, 2019; Verlinda, 2008; Deltas, 2008). Indeed, we support their findings. The volume-adjusted price converges with its long-run equilibrium position through a slower adjustment speed in rural areas, the service market, and stations with no competitor within 3 km.

For all models, and in line with the literature, we find a return to long-run pricing asymmetry. Another interesting feature of using daily data is that for all models, we find

that the return to long-run asymmetry is always achieved within one week, both for average pump prices and for using volume-weighted prices. Several studies using more aggregated data find longer adjustment periods, probably due to the lower data frequency but also due to longer datasets. Our dataset, though representing real transaction data, is after all, only 366 days long.

We have aggregated across stations using different selection criteria; regions, service level, and local competition level. Since we have access to a detailed panel of transaction prices and volumes per day for as many as 147 gasoline stations, the next interesting step would be to estimate dynamic models for each station and relate the results to variables measuring the particularities of the stations' demographics with regards to local demand, competition level, and station heterogeneity. This would potentially take us further in the direction of what Byrne (2019) did while estimating price dynamics across stations and comparing the results across local markets in Canada.

References

[1] Apergis, N., and Vouzavalis, G., (2018), "Asymmetric pass through of oil prices to gasoline prices: Evidence from a new country sample, " *Energy Policy*, Vol.114, pp. 519-528.

[2] Asplund, M., Eriksson, R., and Friberg, R., (2000), "Price adjustments by a gasoline retail chain, " *Scandinavian Journal of Economics*, Vol.102, iss.1, pp. 101-121.

[3] Bachmeier, L., and Griffin, J. M., (2003), "New evidence on asymmetric gasoline price responses, " *Review of Economics and Statistics*, Vol.85, iss.3, pp. 772-776.

[4] Bacon, R. W., (1991), "Rockets and feathers: The asymmetric speed of adjustment of UK retail gasoline prices to cost changes, " *Energy Economics*, Vol.13, iss.2, pp. 211-218.

[5] Balmaceda, F., and Soruco, P., (2008), "Asymmetric dynamic pricing in a local gasoline retail market, " *Journal of Industrial Economics*, Vol.56, iss.3, pp. 629-653.

[6] Banerjee, A., Dolado, J., and Mestre, R., (1998), "Error-correction mechanism tests for cointegration in a single-equation framework, " *Journal of time series analysis*, Vol.19, iss.3, pp. 267-283.

[7] Barron, J., Taylor, B., and Umbeck, J., (2004), "Number of sellers, average prices, and price dispersion, " *International Journal of Industrial Organization*, Vol.22, iss.8, pp. 1041-1066.

[8] Barron, J., Umbeck, J., and Waddell, G., (2008), "Consumer and competitor reactions: Evidence from a field experiment, " *International Journal of Industrial Organization*, Vol.26, iss.2, pp. 517-531.

[9] Borenstein, S., Cameron, C., and Gilbert, R., (1997), "Do gasoline prices respond asymmetrically to crude oil price changes?, " *Quarterly Journal of Economics*, Vol.112, iss.1, pp. 305-339.

[10] Borenstein, S., and Shepard, A., (1993), "Dynamic pricing in retail gasoline markets, " *National Bureau of Economic Research.*

[11] Bresnahan, T., (1982) , "The oligopoly solution concept is identified, " *Economic Letters*, Vol.10, iss.1-2, pp. 87-92.

[12] Byrne, D. P., (2019), "Gasoline Pricing in the Country and the City, " *Review of Industrial Organization*, Vol.55, iss.2, pp. 209-235.

[13] Byrne, D. P., and De Roos, N., (2019), "Learning to coordinate: A study in retail gasoline, " *American Economic Review*, Vol.109, iss.2, pp. 591-619.

[14] Carlin, B. I., (2009), "Strategic price complexity in retail financial markets, " *Journal of Financial Economics*, Vol.91, iss.3, pp. 278-287.

[15] Clemenz, G., and Gugler, K., (2006), "Locational choice and price competition: some empirical results for the Austrian retail gasoline market, " *Empirical Economics*, Vol.31, iss.2, pp. 291-312.

[16] Conlisk, J., Gerstner E., and Sobel J., (1984), "Cyclic Pricing by a Durable Goods Monopolist, " *Quarterly Journal of Economics*, Vol.99, iss.3, pp. 489-505.

[17] Clark, R., and Houde, J. F., (2013), "Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case, " *American Economic Journal: Microeconomics*, Vol.5, iss.3, pp. 97-123.

[18] Deltas, G., (2008), "Retail gasoline price dynamics and local market power, " *Journal of Industrial Economics*, Vol.56, iss.3, pp. 613-628.

[19] Douglas, C., and Herrera, A. M., (2010), "Why are gasoline prices sticky? A test of alternative models of price adjustment, " *Journal of Applied Econometrics*, Vol.25, iss.6, pp. 903-928.

[20] Eckert, A., (2013), "Empirical studies of gasoline retailing: A guide to the literature, " *Journal of Economic Surveys*, Vol.27, iss.1, pp. 140-166.

[21] Eckert, A., (2003), "Retail price cycles and the presence of small firms, " *International Journal of Industrial Organization*, Vol.21, iss.2, pp. 151-170.

[22] Eckert, A., (2002), "Retail price cycles and response asymmetry, " *Canadian Journal of Economics*, Vol.35, iss.1, pp. 52-77.

[23] Eckert, A., and West, D., (2005), "Price uniformity and competition in a retail gasoline market, " *Journal of Economic Behavior & Organization*, Vol.56, iss.2, pp. 219-237.

[24] Edgeworth, F. Y., (1925), "The Pure Theory of Monopoly, " *Papers Relating to Political Economy*, Vol.1, pp.111-142.

[25] Ellison, G., and Wolitzky A., (2012), "A search cost model of obfuscation, " *The RAND Journal of Economics*, Vol.43, iss.3, pp. 417-441.

[26] Engle, R. F., and Granger, C. WJ., (1987), "Co-integration and error correction: representation, estimation, and testing, " *Econometrica: journal of the Econometric Society*, pp. 251-276.

[27] Foros, Ø., Nguyen-Ones, M., and Steen, F., (2021), "The Effects of a Day Off from Retail Price Competition: Evidence on Consumer Behavior and Firm Performance in Gasoline Retailing, " *International Journal of Economics of Business*.

[28] Foros, Ø., and Nguyen-Ones, M., (2020), "Coordinate to obfuscate? The role of prior announcements of recommended prices, " *Economics Letters*, Vol.198.

[29] Foros, Ø., Nguyen-Ones, M., and Steen, F., (2013a), "Vertical control and price cycles in gasoline retailing?, " *The Scandinavian Journal of Economics*, Vol.115, iss.3, pp. 640-661.

[30] Foros, Ø., and Steen, F., (2013b), "Retail pricing, vertical control and competition in the Swedish gasoline market, " *retail-pricing-vertical-control-and-competition-in-the-swedish-gasoline-market.pdf*.

[31] Granger, C. WJ., and Newbold, P. E., (1974), "Spurious regressions in econometrics. Baltagi, Badi H. A , " *Companion of Theoretical Econometrics*, pp. 557-561.

[32] Grasso, M., Manera, M., (2007), "Asymmetric error correction models for the oil-gasoline price relationship, " *Energy Policy*, Vol.35, iss.1, pp. 156-177.

[33] Haucap, J., Heimeshoff, U., and Siekmann, M., (2017), "Fuel Prices and Station Heterogeneity on Retail Gasoline Markets, " *The Energy Journal*, Vol.38, iss.6, pp. 81-103.

[34] Hofstetter, M., and Tovar, J, (2010), "Common Knowledge Reference Price and Asymmetric Price Adjustments, " *Review of Industrial Organization*, Vol.37, iss.2, pp. 141-159.

[35] Johnson, R. N., (2002), "Search costs, lags and prices at the pump, " *Review of Industrial Organization*, Vol.20, iss.1, pp. 33-50.

[36] Hastings, J., (2004), "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California, " *American Economic Review*, Vol.94, iss.1, pp. 317-328.

[37] Lau, L. J., (1982), "On identifying the degree of competitiveness from industry price and output data, " *Economic Letters*, Vol.10, iss.1-2, pp. 93-99.

[38] Lewis, M. S., and Marvel, H. P., (2011), "When do consumers search?, " *Journal of Industrial Economics*, Vol.59, iss.3, pp. 457-483.

[39] Lewis, M. S., and Noel, M. D., (2011), "The speed of gasoline price response in markets with and without Edgeworth cycles, " *Review of Economics and Statistics*, Vol.93, iss.2, pp. 672-682.

[40] Lewis, M. S., (2004), "Asymmetric price adjustment and consumer search: An examination of the retail gasoline market.

[41] Lewis, M. S., (2011), "Asymmetric price adjustment and consumer search: An examination of the retail gasoline market, " *Journal of Economics & Management Strategy*, Vol.20, iss.2, pp. 409-449.

[42] Maskin, E., and Tirole, J., (1988), "A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles., " *Econometrica: Journal of the Econometric Society* , pp. 571-599.

[43] Nguyen-Ones, M., and Steen, F., (2021), "Market Power in Retail Gasoline Markets, " *Centre for Economic Policy Research.*

[44] Pesaran, M. H., Shin, Y., and Smith, R. J., (2001), "Bounds testing approaches to the analysis of level relationships, " *Journal of applied econometrics*, Vol.16, iss.3, pp. 289-326.

[45] Shin, Y., Yu, B., and Greenwood-Nimmo, M., (2014), "Modelling asymmetric cointegration and dynamic multipliers in a non-linear ARDL framework, " *Book,Springer*, pp. 281-314.

[46] Swedish Petroleum and Biofuel Institute (2013), "SPBI Branschfakta 2013 (Facts of the Industry, " *Stockholm: Swedish Petroleum and Biofuel Institute.*

[47] Netz, J. S., and Taylor, B. A., (2002), "Maximum or Minimum Differentiation? Location Patterns of Retail Outlets, " *The Review of Economics and Statistics*, Vol.84, iss.1, pp. 162-175.

[48] Noel, M. D., (2016), "Retail Gasoline Markets, " *In: Basker, E. ed. Handbook on the Economics of Retailing and Distribution. Cheltenham: Edward Elgar Publishing*, pp. 392-412.

[49] Noel, M. D., (2007), "Edgeworth price cycles, cost-based pricing and sticky pricing in retail gasoline markets, " *Review of Economics and Statistics*, Vol.89, iss.2, pp. 324-334.

[50] Peltzman, S., (2000), "Prices rise faster than they fall, " *Journal of Political Economy*, Vol.108, iss.3, pp. 466-502.

[51] Shepard, A., (1991), "Price Discrimination and Retail Configuration, " *Journal of Political Economy*, Vol.99, iss.1, pp. 30-53.

[52] Slade, M., (1987), "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion, " *The Journal of Industrial Economics*, Vol.35, iss.4, pp. 499-516.

[53] Tappata, M., (2009), "Rockets and feathers: Understand asymmetric pricing, " *RAND Journal of Economics*, Vol.40, iss.4, pp. 466-502.

[54] Verlinda, J. A., (2008), "Do rockets rise faster and feathers fall slower in an atmosphere of local market power? Evidence from the retail gasoline market, " *Journal of Industrial Economics*, Vol.56, iss.3, pp. 581-612.

[55] Yang, H., and Ye, L., (2008), "Search with learning: Understanding asymmetric price adjustments, " *RAND Journal of Economics*, Vol.39, iss.2, pp. 547-564.

3.8 Appendix

Table A1. Augmented Dickey-Fuller test for unit root

Variables	**Test Statistic**	**1% Critical Value**	**5% Critical Value**	**10% Critical Value**
VP	-1.634	-3.451	-2.875	-2.570
OP	-1.692	-3.451	-2.875	-2.570
IP	-1.528	-3.451	-2.875	-2.570
Variables	**Test Statistic**	**1% Critical Value**	**5% Critical Value**	**10% Critical value**
ΔVP	-14.295	-3.451	-2.875	-2.570
ΔOP	-15.409	-3.451	-2.875	-2.570
ΔIP	-18.936	-3.451	-2.875	-2.570

Table A2. Estimation results in Sweden using volume-adjusted prices, daily dataset

Method: ARDL		**NARDL with LR**		**NARDL with LR & SR**	**NARDL with SR**
Variables	Coefficients	Coefficients	Variables	Coefficients	Coefficients
Constant	2.694*** (0.251)	3.376*** (0.298)	Constant	4.689*** (0.359)	2.669*** (0.215)
VP_{t-1}	-0.326*** (0.029)	-0.236*** (0.021)	VP_{t-1}	-0.328*** (0.025)	-0.325*** (0.025)
IP_{t-1}	0.391*** (0.035)		IP_{t-1}		0.392*** (0.030)
IP^{+}_{t-1}		0.281*** (0.025)	IP^{+}_{t-1}	0.389*** (0.030)	
IP^{-}_{t-1}		0.282*** (0.025)	IP^{-}_{t-1}	0.391*** (0.030)	
ΔVP_{t-1}	0.370*** (0.045)	0.388*** (0.045)	ΔVP_{t-1}	0.343*** (0.043)	0.347*** (0.043)
ΔVP_{t-2}	-0.235*** (0.044)	-0.291*** (0.049)	ΔVP_{t-2}	-0.238*** (0.042)	-0.238*** (0.043)
ΔVP_{t-3}	0.126*** (0.043)	0.168*** (0.048)	ΔVP_{t-3}	0.143*** (0.040)	0.147*** (0.041)
ΔVP_{t-4}		-0.105** (0.045)	ΔIP^{+}_{t}	0.234*** (0.070)	0.238*** (0.070)
ΔVP_{t-7}	0.117*** (0.038)		ΔIP^{+}_{t-2}	-0.190*** (0.072)	-0.191*** (0.072)
ΔIP_{t}		-0.122*** (0.046)	ΔIP^{-}_{t}	-0.548*** (0.075)	-0.148** (0.069)
ΔIP_{t-1}	-0.110** (-0.055)		ΔIP^{-}_{t-1}	-0.221*** (0.079)	-0.212*** (0.079)
ΔIP_{t-2}	-0.124** (0.052)				
LR_{IP}		1.191*** (0.039)	LR_{IP}		1.205*** (0.023)
LR^{+}_{IP}		1.191*** (0.039)	LR^{+}_{IP}	1.193*** (0.024)	-
LR^{-}_{IP}		1.198*** (0.036)	LR^{-}_{IP}	1.187*** (0.026)	-
λ	-0.320*** (0.029)	-0.322*** (0.028)	φ	-0.326*** (0.026)	-0.325*** (0.025)
R^2	0.467	0.414	R^2	0.489	0.485
t_{BDM}	-13.513***	-13.058***	t_{BDM}	-13.433***	-13.633***
F_{PSS}	68.197***	44.804***	F_{PSS}	57.976***	85.188***
LR_{IP}	-	-0.061 (0.169)	LR_{IP}	-0.007* (0.004)	-
IP_{SR}	-	-	IP_{SR}	0.813*** (0.159)	0.407*** (0.151)

*Note: Standard errors in in parentheses, and *p<0.10, **p<0.05, ***p<0.01.*

Table A3. Estimation results in Sweden using **average pump prices**, daily dataset

Method:	**ARDL with LR & SR Symmetry**	**NARDL with LR Asymmetry & SR Symmetry**		**NARDL with LR & SR Asymmetry**	**NARDL with LR Symmetry & SR Asymmetry**
Variables	Coefficients	Coefficients	Variables	Coefficients	Coefficients
Constant	3.021*** (0.268)	4.021*** (0.353)	Constant	5.297*** (0.381)	3.029*** (0.226)
VP_{t-1}	-0.369*** (0.032)	-0.280*** (0.025)	VP_{t-1}	-0.371*** (0.027)	-0.371*** (0.026)
IP_{t-1}	0.450*** (0.039)		IP_{t-1}		0.451*** (0.032)
IP^{+}_{t-1}		0.337*** (0.030)	IP^{+}_{t-1}	0.449*** (0.033)	
IP^{-}_{t-1}		0.338*** (0.030)	IP^{-}_{t-1}	0.450*** (0.032)	
ΔVP_{t-1}	0.317*** (0.044)	0.318*** (0.047)	ΔVP_{t-1}	0.283*** (0.042)	0.285*** (0.042)
ΔVP_{t-2}	-0.189*** (0.043)	-0.193*** (0.046)	ΔVP_{t-2}	-0.186**** (0.041)	-0.186*** (0.041)
ΔVP_{t-3}	0.124*** (0.041)	0.120*** (0.044)	ΔVP_{t-3}	0.126*** (0.039)	0.128*** (0.040)
ΔIP_{t}		-0.191*** (0.050)	ΔIP^{+}_{t}	0.247*** (0.073)	0.249*** (0.073)
ΔIP_{t-1}	-0.119** (-0.059)		ΔIP^{+}_{t-2}	–0.309*** (0.075)	-0.310*** (0.075)
ΔIP_{t-2}	-0.176*** (0.055)	-0.107** (0.055)	ΔIP^{-}_{t}	-0.680*** (0.078)	-0.226*** (0.071)
ΔIP^{-}_{t-1}			ΔIP^{-}_{t-1}	-0.246*** (0.084)	-0.244*** (0.085)
LR_{IP}	1.219*** (0.022)		LR_{IP}		1.218*** (0.021)
LR^{+}_{IP}		1.203*** (0.035)	LR^{+}_{IP}	1.211*** (0.024)	-
LR^{-}_{IP}		1.206*** (0.032)	LR^{-}_{IP}	1.213*** (0.022)	-
λ	-0.369*** (0.032)	-0.368*** (0.030)	φ	-0.367*** (0.027)	-0.371*** (0.026)
R^2	0.441	0.377	R^2	0.491	0.491
t_{BDM}	-13.923***	-13.213***	t_{BDM}	-13.875***	-14.113***
F_{PSS}	68.197***	44.348***	F_{PSS}	65.746***	98.552***
LR_{IP}	-	-0.003 (0.005)	LR_{IP}	-0.002 (0.004)	-
IP_{SR}	-	-	IP_{SR}	0.865*** (0.169)	0.409*** (0.158)

*Note: Standard errors in in parentheses, and *p<0.10, **p<0.05, ***p<0.01.*

Table A4. NARDL models with LR & SR asymmetry across regions

Regions:	**Stockholm**	**Gothenburg**	**Malmö**	**E6**	**smaller cities**	**rural areas**
Method:	NARDL with LR & SR asymmetry	NARDL with LR & SR asymmetry	NARDL with LR & SR asymmetry	NARDL with LR & SR asymmetry	NARDL with LR & SR asymmetry	NARDL with LR & SR asymmetry
Variables	Coefficients	Coefficients	Coefficients	Coefficients	Coefficients	Coefficients
Constant	4.091*** (0.258)	4.515*** (0.484)	5.156*** (0.465)	4.863*** (0.408)	4.402*** (0.301)	4.156*** (0.248)
VP_{t-1}	-0.285*** (0.017)	-0.319*** (0.034)	-0.362*** (0.032)	-0.342*** (0.028)	-0.309*** (0.021)	-0.290*** (0.017)
IP^{+}_{t-1}	0.352*** (0.022)	0.351*** (0.039)		0.392*** (0.033)	0.379*** (0.027)	0.359*** (0.022)
IP^{-}_{t-1}	0.352*** (0.022)	0.354*** (0.039)		0.395 *** (0.033)	0.380*** (0.026)	0.357*** (0.021)
ΔVP_{t-1}	0.362*** (0.041)	0.265*** (0.045)		0.105** (0.044)	0.345*** (0.043)	0.341*** (0.041)
ΔVP_{t-2}	-0.164*** (0.043)	-0.214*** (0.045)		-0.111*** (0.042)	-0.151*** (0.043)	-0.120** (0.043)
ΔVP_{t-3}	0.139*** (0.040)			0.392*** (0.033)	0.153*** (0.041)	0.166*** (0.039)
ΔVP_{t-7}		0.109*** (0.043)			0.105*** (0.039)	
ΔIP^{+}_{t}	0.143** (0.057)	0.246* (0.140)	0.436*** (0.120)	0.205*** (0.087)	0.210** (0.074)	0.137*** (0.056)
ΔIP^{+}_{t-2}		-0.412*** (0.139)	-0.282*** (0.120)	-0.152** (0.089)	-0.291** (0.076)	-0.124** (0.056)
ΔIP^{+}_{t-4}	0.127** (0.055)	0.491*** (0.140)				
ΔIP^{+}_{t-6}		-0.223* (0.136)				
ΔIP^{+}_{t-9}						0.105** (0.053)
ΔIP^{+}_{t-13}						0.096* (0.052)
ΔIP^{-}_{t}	-0.504*** (0.061)	-0.340** (0.146)	-0.825*** (0.124)	-0.662*** (0.092)	-0.554*** (0.079)	-0.485*** (0.058)
ΔIP^{-}_{t-1}		-0.446*** (0.141)	-0.369*** (0.130)	-0.211** (0.097)		
ΔIP^{-}_{t-6}		-0.307** (0.137)			-0.136** (0.072)	
ΔIP^{-}_{t-8}						0.099** (0.053)
ΔIP^{-}_{t-10}					-0.183*** (0.071)	
ΔIP^{-}_{t-16}						0.101** (0.051)
ΔIP^{-}_{t-18}						0.113** (0.050)
LR^{+}_{IP}	1.234*** (0.024)	1.100*** (0.053)	1.200*** (0.040)	1.146*** (0.030)	1.229*** (0.030)	1.237*** (0.025)
LR^{-}_{IP}Z	1.233*** (0.023)	1.111*** (0.049)	1.218*** (0.037)	1.156*** (0.029)	1.230*** (0.028)	1.231*** (0.022)
λ	-0.286*** (0.018)	-0.295*** (0.044)	-0.361*** (0.032)	-0.332*** (0.029)	-0.312*** (0.026)	-0.283*** (0.025)
R^2	0.550	0.394	0.385	0.362	0.499	0.582
t_{BDM}	-15.636***	-9.184***	-11.116***	-11.914***	-14.923***	-16.769***
F_{PSS}	85.492***	30.198***	42.528***	49.142***	73.347***	94.612***
LR_{IP}	0.002 (0.004)	-0.011 (0.008)	-0.018*** (0.006)	-0.010** (0.005)	-0.001 (0.005)	0.007* (0.004)
IP_{SR}	0.774*** (0.107)	1.195*** (0.384)	1.349*** (0.264)	0.926*** (0.197)	0.791*** (0.173)	0.385*** (0.150)

*Note: Standard errors in in parentheses, and *p<0.10, **p<0.05, ***p<0.01.*

Chapter 4

Risk Management in Wholesale Electricity Markets: A Signal Processing Approach*

Benjamin P. Fram[a], Ritvana Rrukaj[a†], Leif K. Sandal [a]

[a]Department of Business and Management Science, NHH Norwegian School of Economics, 5045 Bergen, Norway

Abstract

Wholesale electricity markets exhibit high price volatility; therefore, market participants routinely hedge price risks through the use of financial derivatives, such as by acquiring futures contracts that allow them to buy or sell electricity for a pre-determined price over a set time period. Considering the perspective of a market participant, this study applies techniques from the field of signal processing to map medium-term relationships between load and prices in the New York State wholesale electricity market (NYISO). We first sort and filter raw market data from the year 2006 to 2018 into annual, seasonal (90-day), and weekly moving averages. We then develop three nonlinear models that relate the moving averages of load and prices while accounting for seasonality in the data. As a next step, we extract the model from the filtered data and employ the Box–Jenkins methodology for the time series analysis to model the resulting noisy residuals. We conclude by performing out-of-sample forecasts of the seasonal moving average model using data from 2019 and 2020. The forecast results indicate that our modelling approach is useful in predicting the relationship between seasonal load and prices in the NYISO and that it may assist market participants to better assess market risks and adjust their hedging positions.

Keywords: Electricity markets, signal processing, risk management, mathematical modelling, nonlinear data analysis

*A six-page version of the paper has been published on the IISE Annual conference proceedings page. A shorter version of this paper has been submitted to the IEEE journal, Transactions on Power Systems.

†We thank conference participants at the 41st International Symposium on Forecasting (ISF) 2021, 2nd International Workshop on Forecasting for Social Good, and IISE Annual conference 2021 for helpful comments.

4.1 Introduction

In the mid-1990s, electricity markets across Europe, North America, and Australia underwent a dramatic period of transformation. The electric power industry, which had historically been dominated by vertically integrated and state-regulated monopolies, was restructured into an open-market paradigm with free entry and auctions where market participants were required to bid competitively to buy and sell power. Though this shift in market design has generally led to increased competition and the entry of new market participants, restructuring has also created new challenges for market participants and regulators (Monitoring Analytics, 2020).

One notable challenge is that restructured power markets can exhibit extreme volatility. The supply and demand for power must always be equal and even minor fluctuations in the weather, fuel costs, or plant availability can cause dramatic price spikes. This price volatility was brought into sharp focus during the California energy crisis of the early 2000s and the Polar vortex during the winter of 2014.[1] Though the former was partially due to market manipulation and the latter was due to extreme weather, both resulted in prolonged periods of extreme price volatility that caused billions of dollars in excess charges and catastrophic losses for some industry participants. To protect themselves against price volatility, electricity market participants commonly engage in financial hedging strategies.[2] Previous studies have highlighted that hedging practices are crucial for protecting against spikes in the price of the commodities that generate electricity in the case of generators and, in the case of retailers, in the price of the electricity itself. For example, Consolidated Edison (ConEd) notes the following in its 2019 Annual Report:

"To reduce the volatility of its customers' electric energy costs, the company has contracts to purchase electric energy and enters into derivative transactions to hedge the costs of a portion of its expected purchases under these contracts and through the NYISO's wholesale electricity market." (Consolidated Edison 2019 Annual Report – Page 22).

"Con Edison's subsidiaries hedge market price fluctuations associated with physical purchases and sales of electricity, natural gas, steam and, to a lesser extent, refined fuels by using derivative instruments including futures, forwards, basis swaps, options, transmission congestion contracts and financial transmission rights contracts" (Consolidated Edison 2019 Annual Report – Page 160).

However, there are costs and risks associated with the purchase of futures contracts, forward agreements, and other derivatives; as a consequence, inefficient hedging strategies can result in significant financial losses. Because of the potential negative impact that such

[1] See Stoft (2002) and Monitoring Analytics (2014).

[2] From Klemola and Sihvonen (2015): "Electricity companies are among the most active users of financial derivatives. This is mostly because highly volatile electricity prices require rigorous hedging, but also because increased competition within the industry forces companies to find new sources of profit, such as arbitrage and speculation."

"hedging losses" can impose on the public in the form of higher rates, some regulators require market participants to file reports outlining their hedging strategies and they closely monitor those programs.[3] In addition, many companies, including publicly traded companies such as ConEd, routinely report on the overall financial performance of and risks associated with their hedging programs. Put simply, hedging programs are integral to the electricity markets, and the participants in those markets are strongly motivated to implement programs that are not only effective but financially prudent and efficient.

In this study, we develop a modelling approach that can potentially be utilized by power market participants to fine-tune their hedging strategies. We develop this approach using data from the New York Independent System Operator (NYISO), which is the wholesale power market clearinghouse for the state of New York. These data are of high-resolution, well organized, and are freely available online. However, this modeling approach could likely be applied to many other electricity markets. In order to take advantage of the fact that medium-term demand for power and medium-term prices for power are highly synchronous, we use a series of non-linear models to bi-directionally map these two time series to each other. This mapping enables anyone with an accurate forecast of medium-term load to be able to estimate future medium-term prices, and vice versa.

A *signal* is defined as "any sequence of numerical data that varies with respect to an underlying independent variable, mostly time," and signal processing is the subfield of applied engineering that focuses on the analysis of signal behavior (Nepal, 2015). Signal processing developed into a formal field of study in the 1940s in order to construct and analyze communication systems and has since then been widely used in physics and engineering. In recent years, however, signal processing has begun to receive more attention in the field of finance where it can be used to analyze economic variables such as derivative prices.

Signal processing can be broadly broken down into three main modeling components: filtration, waveform identification, and noise analysis (Nepal, 2015). Filtration involves performing mathematical transformations on a signal in an effort to distill its random or "noisy" component. This smoothens the signal and makes it easier to carry out the next step, namely, fitting a waveform function that approximates the signal's persistent, low frequency periodicity. In the case quasi-periodic signals, it is common for practitioners to use aperiodic functions, a flexible technique that consists of summing a series of trigonometric functions that collectively model the amplitude and frequency of the signal in question. Finally, the leftover signal noise, that is, the model residuals or "error terms" that are not captured by the filtration and waveform fitting steps, is analyzed to ensure that it is stationary white noise and that it does not contain any additional systemic information.

This study employs these three modeling techniques to model the medium-run relationship between load and prices in the NYISO. We begin by sorting thirteen years of hourly real-time NYISO market data into four distinct time periods that reflect the consumer de-

[3]See Washington State Docket UG-132019.

mand for power over the course of the day: morning, midday, evening, and night. We then consolidate these hourly data by calculating daily observations of average load and load-weighted average prices for each of these time periods.

We then employ a series of filtrations on these daily observations. First, we calculate a 365-day moving average (MA) of the daily load and price observations. We then create a new data series where we divide the daily observations by the annual MA on the same date and subtract one. This new filtered data series expresses the daily NYISO load and prices as percentage deviation from the annual MA. As a final step, we calculate the 90-day seasonal MA and 7-day weekly MA of the percentage deviation data.

After creating these annual, seasonal, and weekly moving averages of the NYSIO load and prices, we develop three signal functions that express the load as a function of purely seasonal expressions, prices, and interaction terms between time and prices. As a final modelling step, we subtract the fitted waveform functions from the actual seasonal MA observations to create model residuals, that is, the model noise that is left over from the filtration and modeling steps. We model these noisy residuals using the Box–Jenkins methodology for time series analysis via autoregressive regressions that capture any remaining seasonality left in the data. We choose the seasonal model for the noise analysis as the model yielded the highest R-squared values and because it most closely resembles the typical duration of power futures contracts.

To validate the predictive capability of this signal-processing approach, we perform simple out-of-sample forecasts of the seasonal model using 2019 and 2020 NYISO market data as testing data. We find that the out-of-sample forecasts yield model residuals that closely resemble the training data residuals. This indicates that the seasonal model yields relatively accurate predictions of the relationship between seasonal load and prices across all NYISO zones and time periods. We do note, however, that some forecasts are more accurate for some NYISO zones than others, suggesting that some market zones may require more aggressive hedging strategies than others.

While there are similarities between the standard econometric analysis and the signal processing approach we follow in this study, we believe that there are several key differences. First, while it is common to mathematically transform raw economic data before carrying out an econometric analysis, our approach employs heavier data filtration than is common in most econometric studies. For instance, in the econometric analysis of financial time series, it is common to transform raw price data into percentage returns. In contrast, we immediately transform NYISO market data into annual moving averages and then apply several additional filtrations thereafter in order to eliminate more structures (quasi-periodicities and trends) from the signals before we study the residuals. Second, we perform additional analysis on our waveform model residuals, which is uncommon in most econometric studies. Though some econometric studies provide descriptive statistics about model residuals, few actually go so far as to build additional models that describe residual behavior. This part of the

analysis, that is, analyzing the noise left over from the filtration and waveform modelling steps, is standard practice in signal processing. By carrying out this additional step, we were able to find left over explanatory information in our model residuals despite having initially high model R-squared values. Finally, while it is the intent of most econometric pursuits to identify causality in a given economic system, causality is not the primary focus of our analysis. Undoubtedly, short-term fluctuations in power demand cause changes in electricity prices and the two variables are highly correlated across time. However, it is our intent to accurately model and predict this correlational relationship in order to exploit it for the purpose of risk management in electricity markets. In our waveform functions that describe the basic structural behavior of load, we allow for prices modifying amplitudes and frequencies.

By inspecting the estimated parameter values of our models, we find that they can reveal a highly detailed picture of the dynamics between load and prices across the NYISO market areas and time periods. The vast majority of these parameter values are consistent with the known physical characteristics of the NYISO market landscape, and the waveform functions yield high R-squared values—two results that strengthen the validity of our analysis.

We believe that this paper makes several unique contributions to applied economic research. First, we use innovative filtration methods on the raw NYISO load data: we have not seen any previous studies express load in terms of percentage deviations from the annual mean. Second, while most previous research has focused solely on power market load or price forecasting, we actually focused on forecasting the relationship between the two. Third, we take advantage of the highly synchronous moving averages of load and prices by using prices to modulate the amplitude and frequency of load in a nonlinear model setting. Finally, due to the extensive time period and zone based sorting of the NYISO market data, we are able to present a very detailed analysis of the price and load dynamics across the NYISO and offer insights about price elasticity in the market. This analysis spans thirteen years of the NYSIO market operation and, due to the nonlinear nature of our models, reveals many extra details and dynamics that are not easily visible using standard linear regressions.

We also believe that the modeling paradigm and market analysis presented in this paper could be of interest to electricity market participants and other risk management practitioners. Knowing that many load serving entities (LSEs) already employ sophisticated hedging strategies in practice, we do not suggest any new hedging strategies in this paper. The signal processing approach we have developed can, however, be used to independently estimate the medium-term average price of power given that an LSE has a reasonably accurate medium term load forecast. Since we have developed functions that relate load, time, and price, an improved forecast of average prices will likely yield improved forecast of load. Most LSEs are capable of accurately forecasting short to medium term load and we therefore believe that the former is more useful to market participants. With independent, data-driven estimates of medium-term power prices, LSEs could potentially make better-informed decisions that

can lead to lower hedging costs. Such an estimate can be compared to the current price of power in the futures markets and potentially help utilities better determine whether power is selling at a premium, discount, or fair market value.

In addition, by using information about the demand for electricity, we believe that our modeling approach potentially yields an enhanced methodology for assessing price risk in wholesale power markets. Rather than simply using past information about prices to predict future values, we use past information about both load and prices together. Because we are using moving averages in addition to periodic terms and flexible interaction terms in the waveform functions, our models can capture long-run, persistent trends and medium-term trends alike.

This paper is organized in the following manner. Section 4.2 presents a review of relevant literature. Section 4.3 provides a brief overview of electricity markets and the physical characteristics that influence their dynamics, with a special focus on the NYISO. We also briefly discuss futures contracts for electricity. Section 4.4 provides a thorough overview of the NYISO market data used in the analysis and discusses their statistical properties. Section 4.5 presents the processes used to sort and filter the raw market data and three waveform-like functions used to model three different moving averages of the filtered data. Section 4.6 presents the results of the parameter estimation of models and discusses the market dynamics that they reveal about the NYSIO. In Section 4.7, we analyze the noisy residuals of the 90-day seasonal model using the Box–Jenkins time series methodology. Section 4.8 presents the results of our out-of-sample forecasts of the seasonal model and offers some insights for risk management in the NYISO based on these results. Finally, Section 4.9 presents our concluding remarks and potential paths for future research.

4.2 Literature review

During the past decade, load and price forecasting in electricity markets has been an active field of research (Aggarwal et al., 2009). Researchers have concentrated not only on models but also on data pre-processing techniques and other mathematical methods that increase forecasting accuracy. This research has focused on forecasting over variety of time scales, such as; short term which may span several hours to several days; medium-term, which may span a week, several months or even several years, and; long-term which may span several years to a decade. For the short-term, three distinct data-driven methods have been used to predict load or price time series: Regression or casual models (Engle et al., 1992) ARIMA configurations (Christiansen, 1971; Taylor and McSharry, 2007; Taylor, 2011) and Artificial intelligence models, including boosting machines (Lloyd, 2014), Artificial Neural Network (ANN) models (Charytoniuk and Chen, 2000; Han et al., 2018; Taylor and Buizza, 2006), and Generalized Additive Models (GAM) (Goude et al., 2013). ARIMA and exponential smoothing models are based on the initial assumption that predicted variable are strongly stationary, which may limit their application in the case of load or price prediction.

Moreover, these methods are preferred by some researchers when dealing with linear data. The so-called similar-day approach is simple but has low accuracy when large variations exist in load time series. Exponential smoothing methods (ESM) perform better than the standard moving average method, but it is difficult to determine the parameter of smoothing factors. GAM is a flexible nonlinear modeling approach and provides better estimation accuracy than conventional linear models. It works best, however, for trends that are steady and systematic. For longer forecasting horizons, a limited number of studies exist.

Medium-term electricity load forecasting (MTLF), however, is an area of research that requires more attention considering its vital applications in the operation and planning of power systems at the generation, transmission, and marketing levels. In the study by Ghiassi and Zimbra (2002), the authors develop an MTLF model in Taiwan by developing a univariate, dynamic multi-layered artificial neural network that uses a Fourier series to capture the non-linearity of load. Their annual and seasonal model forecasts produce mean average percentage error (MAPE) values below 1%.

Similar attempts have been made by Juen and Chang (2004) using the SVM-based model and by Chen et al. (2008) applying the manifold learning methodology-locally linear embedding (LLE). In Juen and Chang (2004) article, the authors succeeded in forecasting next month's daily peak load using load data from 1997 to 1998 provided by EUNITE competition2001, while the forecast accuracy in Chen et al. (2008) is demonstrated by numerical results using historical price data taken from the Eastern U.S. electric power markets. In the studies by Lucia and Schwartz (2002) and Benth and Schmeck (2014), stochastic factors affecting electricity prices are presented. In the article by Lucia and Schwartz (2002), an Ornstein-Uhlenbeck processis used to model the stochastic component of spot prices in the Nordic Power Exchange.

Moreover, a simple deterministic function with one trigonometric expression is utilize to capture the seasonal component of spot prices. Later, Benth and Schmeck (2014) examine the spiking behavior in daily spot electricityprices of the European Energy Exchange (EEX). They develop a deterministic sinusoidal function to capture seasonality in these raw price data. One set of Fourier expressions captures annual seasonality, and another captures semi-annual seasonality. Benth and Schmeck (2014) and Klüppelberg et al. (2010) use deterministic functions with trigonometric expressions to model seasonality in power prices and then proceed to filter the deseasonalized data before analyzing the resulting stochastic component.

Features such as weather are also incorporated into load forecasting models. Chang et al. 2010 study proposes a modified fuzzy neural network considering variables such as temperature, wind, rainfall, air pressure, rainy time, relative humidity, and day light time. The model provides monthly electricity load forecasting in Taiwan. However, features such as weather and climate data are recommended to be avoided in medium-term load forecasting due to uncertain information (Amjady, 2007). However, none of aforementioned models can

address the need of moreflexible model that gives an accurate medium-term forecast horizon.

4.3 The NYISO and Power Market Fundamentals

The NYISO is a federally regulated, non-profit entity that is tasked with two primary functions. The first is to facilitate the sale of bulk electric power in the state of New York by acting as a market clearinghouse. The second is to safely operate the bulk power transmission grid via centralized dispatch, that is, act as a sort of "air traffic controller" for the power grid.

These two functions are deeply connected in the everyday operations of the NYISO. After collecting bids and offers to buy and sell power, the NYISO aggregates these quantities into market demand and supply curves and then clears the market by solving a unit dispatch problem that maximizes economic welfare while respecting the physical limitations of the grid to ensure system reliability. During real-time power delivery, the NYISO constantly balances the power flow on the grid by operating a real-time spot market and directing generators to adjust their power output based on the current market conditions.

4.3.1 NYISO as the grid operator

The physical power grid is comprised primarily of power plants that generate electricity and end users such as private residences, commercial buildings, and industrial facilities. Individual end users often lie within a service territory where a utility or LSE is responsible for purchasing bulk power from the high-voltage grid and delivering it directly to the end user. The demand for power is often times called "load" and we will henceforth use these two terms interchangeably.

Generators and end users are connected to each other via a network of power lines that transmit electricity produced by generators to the end users. These power lines operate at different voltage levels, that is, at different electrical pressures. In the state of New York, the NYISO is responsible for operating power flowing on high-voltage "highway" lines, also called the "transmission network," that have voltages between 115 kV to 765 kV. Power lines that operate at voltages below 115 kV are part of the "distribution" network that directly connect to end users and are operated by local utilities.

In addition to voltage levels, power lines differ in the amount of electricity they can safely transmit. Electrical flow produces heat as a byproduct, and each line has a power rating that dictates how much electricity can flow over the line without causing it to melt. If a line becomes too hot and melts, it could cause a power surge that leads to a cascade of subsequent line failures on the grid and a system blackout.

The high-voltage grid uses alternating current (AC) power flow. In AC power flow, electrons oscillate back and forth at a specific frequency, measured in hertz (Hz). In order for power to flow in an AC setting, all components of the power grid, including generators, must be synchronized at precisely the same frequency, which is 60 Hz in the case of the NYISO.

In practical terms, this means that the demand for electricity must always be precisely equal to the amount that is supplied, that is, power must always be perfectly balanced. If the frequency is too high or low, it may cause electrical equipment such as generators to spin too quickly or too slowly and power blackouts can occur as a result.

As the grid operator, the NYISO is responsible for dispatching generators such that the frequency on the grid is maintained at 60 Hz and that high-voltage line thermal limits are not violated. This requires that the NYISO constantly maintain grid power balance 24 hours a day, every day of the year, by directing the power output of each generator connected to the grid.

4.3.2 NYISO as a market clearinghouse

In addition to operating the physical power grid for system reliability, the NYISO conducts a series of auctions in which all NYISO market participants trade electricity. There are two primary auctions: a "day ahead" market, where participants submit bids to purchase and sell electricity 36 hours before actual power delivery; and a real-time or "balancing" market, where adjustments to the day ahead market take place to reflect the actual physical demand for power. Combined, these two auctions comprise what is commonly referred to as the "two-settlement" system (Stoft, 2002).

In the day-ahead market, all entities that wish to purchase electricity submit hourly bid curves and all entities that wish to sell electricity submit hourly offer curves to the NYISO. These curves must be submitted by 05:00 on the day before the actual power delivery day (i.e., twenty hours before the dispatch day begins).[4] For each hour of the day, each participant submits a curve that is comprised of up to ten increments of price and quantity pairs that indicate how much power the given participant would be willing to sell or purchase at a given price. For each hour of the day, the NYISO then horizontally aggregates all offer curves in a total market supply curve and horizontally aggregates all bid curves into a market demand curve. After calculating the aggregate supply and demand curves, the NYISO clears the market by solving a constrained optimization problem with the objective of maximizing total system economic welfare while taking into account the technological constraints of the power system such as generating capacity and ramping rates of units, thermal line limits, and so on. In the NYISO, this market-clearing mechanism is known as security constrained unit commitment. Finally, the NYISO calculates the day-ahead hourly market prices, and those who submitted cleared demand bids pay the cleared market sellers at the specified day-ahead market prices.

During real-time power delivery, however, any number of factors may cause demand or supply to deviate from the cleared day-ahead market quantities. Unexpected weather conditions, forced plant outages, and excess wind or solar production can all affect the real-time power balance. For this reason, the NYISO operates a real-time or "balancing"

[4] See Craan (2019) for more details.

market where generators are instructed to deviate from their day-ahead market positions to accommodate these changes in power flow.[5] When load deviates from the day-ahead market solution, the demand curve will shift along the supply curve and the marginal cost of power will consequently change. When demand changes suddenly and generating units must quickly ramp up production to meet load, real-time price spikes can occur, adding a major source of risk to market participants. We discuss price volatility more thoroughly in Section 4.2. Because power lines are limited in the amount of power they can transmit, the NYISO employs a zonal pricing scheme to manage congestion in the transmission network. Zonal pricing is a form of locational pricing that will cause prices to differ when transmission lines reach their thermal limits. For example, if the demand for electricity in Zone K exceeds the amount that can be transmitted over the lines that connect it with Zone J, a local generator in Zone K will need to be turned on or ramped up in order to maintain the power balance in Zone K. Since prices are set equal to the cleared offer of the marginal generator, the price in Zone K will increase to reflect the new marginal generator balancing the local power supply. If the marginal generator remains unchanged in Zone J, so too will the price in Zone J. Consequently, Zones J and K will exhibit different prices. The NYISO is comprised of eleven such zones, depicted in Figure 1 below. This figure was directly taken from the Federal Energy Regulatory Commission website on ISOs in the US.[6] When a zone

Figure 1: Zonal map of the NYISO

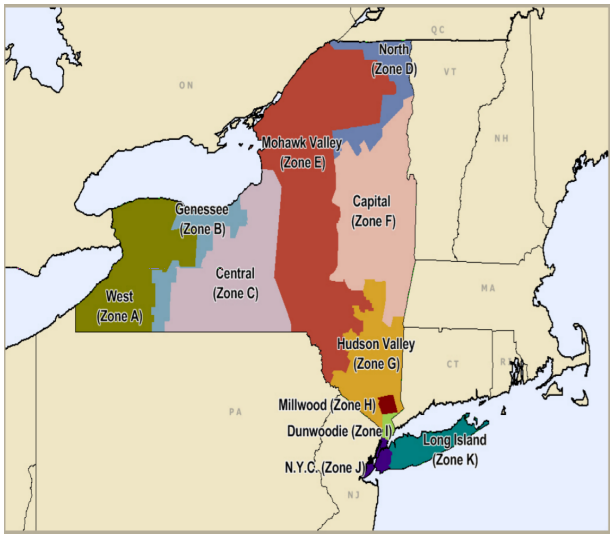

becomes isolated from neighboring zones due to congestion, it may exhibit dramatically different market dynamics than when lines are not congested. For example, before the line connecting Zones J and K becomes congested, the marginal unit in both zones may be a coal

[5]Because generators have already been paid for their output in the day- ahead market, they should be indifferent to increasing or decreasing their power output in real time in a perfectly competitive market. If they are ramped down, they will be forced to buy back power that they had sold in the day- ahead market at a lower price than they sold it for, thus earning a profit. If they are ramped up, they will sell their power at a higher price since they have a monotonically increasing supply curve. See Stoft (2002) for more details.

[6]https://www.ferc.gov/market-assessments/mkt-electric/new-york/2013/10-2013-elec-ny-archive.pdf

plant located in Zone J. After the line becomes congested, however, Zone K may need to rely on a gas-fired generator that is more flexible but more expensive and subject to volatility in the real-time natural gas market. As a result of congestion in the line connecting it to Zone J, price volatility in Zone K might therefore increase.

4.3.3 Unique power market characteristics

The fact that the short-run demand for power in wholesale electricity markets is extremely inelastic is well known; it is a consequence of the fact that end users only pay for their aggregate power consumption via a monthly bill and are largely ignorant to price changes in real time (Stoft, 2002; Monitoring Analytics, 2020).[7] The vast majority of end users consume power when they need it and do not consume less power even when prices drastically spike in the wholesale markets. As a result of this pricing disconnect, the demand for power largely acts as a purely exogenous market force, that is, it is not responsive to real-time prices whatsoever. ISOs simply direct generators to increase or decrease the supply of power as a response to movements in load and are largely powerless to control load under normal circumstances. If one were to graph power prices as a function of demand, it would appear that demand increases as prices increase, implying that the demand for power increases as it becomes more expensive. This result is contrary to standard economic theory and makes little sense without understanding that short-run demand acts independently of market price signals and suppliers must respond to movements in demand at the direction of the NYISO. As a consequence of this highly independent nature of power demand, short-run prices "respond" to movements in demand but not the other way around.[8]

For a variety of reasons, there are administrative constraints on energy offers in the NYISO. Perhaps the most apparent of these constraints is a price cap on both day-ahead and real-time market energy offers as a measure to mitigate the use of market power under tight market conditions. Though there are many details behind this rule, in practice, it effectively dictates that under normal circumstances, the price of electricity cannot exceed $1,000/MWh in the NYISO. Under extraordinary circumstances, for example, when extreme weather causes severe scarcity in natural gas markets, generators may exceed this offer cap if they can document that their actual incurred cost will exceed this value. Instances where prices exceed the offer cap, however, are rare. As an example, in Zone I, there were only 41 hours where prices exceeded $1,000/MWh out of a total of 120,840 hourly price observations that span 13 years.

[7] Since the inception of restructured power markets in the late 1990s, this extreme demand elasticity has persisted as one of the largest challenges of wholesale power market design. Most US markets have even introduced special "demand response" programs in an attempt to increase demand-price elasticity with very limited success (Monitoring Analytics, 2020).

[8] There is an ongoing debate among researchers about the long-run price elasticity of demand for power, that is, the price elasticity of demand that spans several years or longer. Most research has empirically confirmed highly price inelastic short-run demand for power but elasticity estimates are mixed for longer time periods. See Bernstein and Griffin (2006), Deryugina et al. (2017), and Feehan (2018) for a more in-depth discussion.

Finally, prices for power in the NYISO can and do take on negative values. From a power balance perspective, negative prices indicate that there is an oversupply of power on the grid and consumers will be paid for each unit of electricity they use and power producers must pay for each unit they produce. Although it is possible that negative prices can result from sudden congestion in a transmission line, the more likely culprit in several zones are Federal production tax credit (PTC) programs for wind generators. Under these programs, wind generators receive a fixed dollar amount for every MWh of power produced, regardless of the current spot market price. For instance, if a wind generator receives $10 for every MWh they produce, even if the spot price of electricity is $0/MWh, they will receive $10 for 1 MWh output. If spot prices are -$9/MWh, they will still earn a $1/MWh profit. Consequently, wind generators will often offer to sell power for negative prices so long as the absolute value of the negative offer is lower than the value of the PTC. In market areas where the wind-generating capacity is high, prices can and do often become negative during the middle of the night when demand is low and wind speeds are typically the highest as the demand curve will intersect a negative portion of the supply curve.

4.3.4 Supply, demand, and price dynamics

The NYISO is responsible for ensuring that there is adequate generating capacity to meet system peak load and a capacity reserve margin of extra capacity available in the event of an emergency during extreme weather events. The supply of power, that is, the generating capacity of power plants, is generally fixed in the short- and medium-run. This is because new, largescale power plants are generally expensive and take several years to build. Even older, less-efficient units that retire may announce their planned retirement several years in advance so that the NYISO can allow them to shut down without causing the available generating capacity to fall below the minimum capacity reserve requirements.

Coal and nuclear units typically have long-term contracts for fuel and are seldom at risk of fuel shortages since they can store fuel on-site. Run-of-river hydro plants or hydro plants with pumped storage capabilities are also generally immune to water shortages and can run when needed. Gas and oil-fired units also typically have long-term contracts for fuel but have less on-site storage capacity and may experience problems obtaining fuel during extreme weather conditions. Gas units may be especially prone to spot market risks as gas must be delivered in real time to the unit. Renewable power generation from wind and solar in the NYISO is limited but steadily growing. Neither solar nor wind power is upwards dispatchable and, despite having zero fuel costs, are beholden to real-time weather conditions, that is, sunshine and wind speeds. Flexible generators such as gas and hydro units are often used to balance the sudden fluctuations caused by wind and solar resources. End-user demand for electricity systematically varies over the course of a day.

Demand is typically low during the night when consumers are asleep but rises precipitously during the morning hours when people wake up and begin their day. Demand further

increases during the middle of the day when people go to work; large commercial buildings utilize lighting and heating, ventilation, and air conditioning (HVAC) systems, people use electronic devices, factories begin production plans, and so on. Finally, as people return home from work in the evening, demand tends to increase even further as households cook meals, run large electronic appliances such as washing machines and dishwashers, and turn on HVAC systems and lighting. Figures 2 and 3 show the hourly demand and price for electricity in Zone J (New York City) on January 1, 2018 and represent typical daily load and price patterns in the NYISO.

Figure 2: Hourly load in Zone J on January 1, 2018

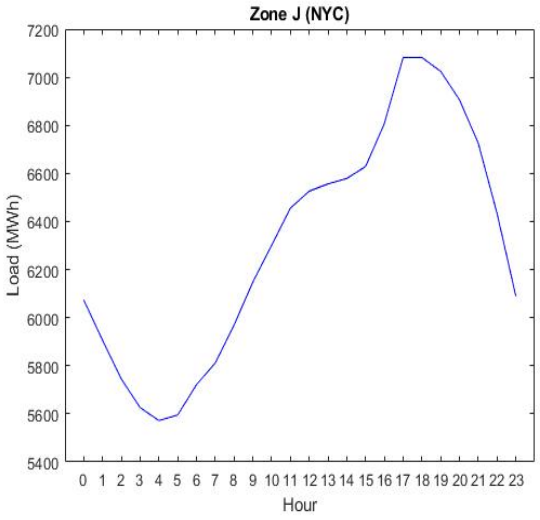

Figure 3: Hourly real-time power prices in Zone J on January 1, 2018

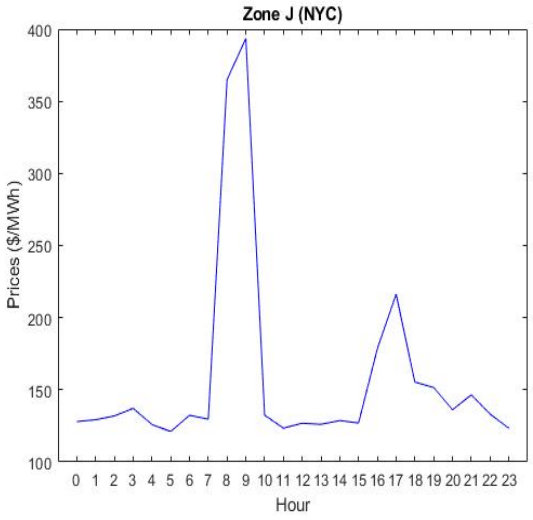

Figure 2 clearly shows that power demand is lowest at night, begins to rise rapidly around 07:00, and then plateaus at an elevated level from about 11:00 to 15:00 as people are at work and school. The load then rapidly increases again around 15:00 until it reaches its peak from 17:00 to 18:00 and begins to drop off as people end their days. In Figure 3, we see that prices are generally stable throughout the day but rise precipitously during the morning and evening ramp-up periods when demand is rising rapidly.

In addition to daily patterns, the demand for electricity exhibits several additional time-dependent patterns. Demand for power is typically higher during the weekdays and both lower and with less-pronounced peaks during the weekends. There are also strong seasonal patterns in load; load is generally lowest during the spring and autumn when weather is mild and cooling and heating needs are low, but tends to increase during the winter as more power is used for heating. In the NYISO, demand increases sharply during the summer when hot weather leads to a precipitous rise in air conditioning. Though these annual seasonal patterns occur rather cyclically, the intensity of these weather patterns may vary from year to year, causing peak demand in one year to be higher than in another. Finally, in the very long run, the demand for power may be affected by population growth or shrinkage or gradual

changes in electricity use such as energy efficiency, the emergence of electric vehicles, and other technological, demographic, and behavioral factors. Figures 4 and 5 show the hourly load in Zone J over the course of 2018 and during a five-year period from 2013 to 2018. In

Figure 4: Hourly load in Zone J in the year 2018

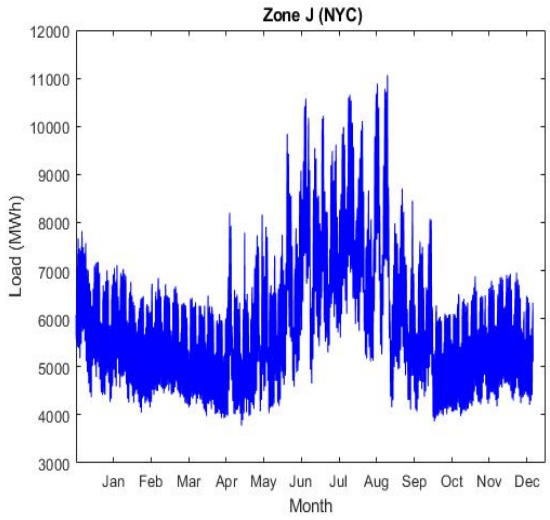

Figure 5: Hourly load in Zone J from 2014 to 2018

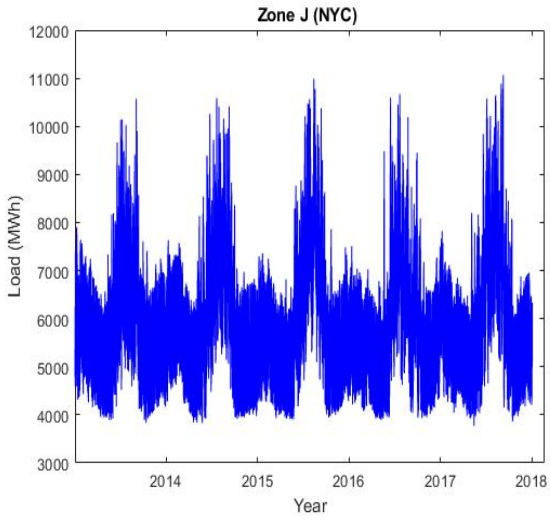

Figure 4, we clearly see an annual seasonal pattern where load is elevated slightly during the winter, is lower during the spring and autumn, and is significantly elevated during the summer. Figure 5 shows how this annual seasonal pattern repeats itself year after year. We can also see the yearly variations of this seasonal pattern in Figure 5, as the summers of 2015 and 2016 seem to have been cooler than the summers of 2014 and 2017. Despite these differences in seasonal peaks, the overall demand for power seems to exhibit a stable long-term trend, indicating that load was not growing or shrinking dramatically during this time period. Although power prices generally follow the same seasonal pattern as load, prices are far more volatile which in turn makes these seasonal patterns more difficult to see. Figure 6 lists the power prices for Zone J during the course of 2018.

Figure 6: Hourly real-time power prices in Zone J in the year 2018

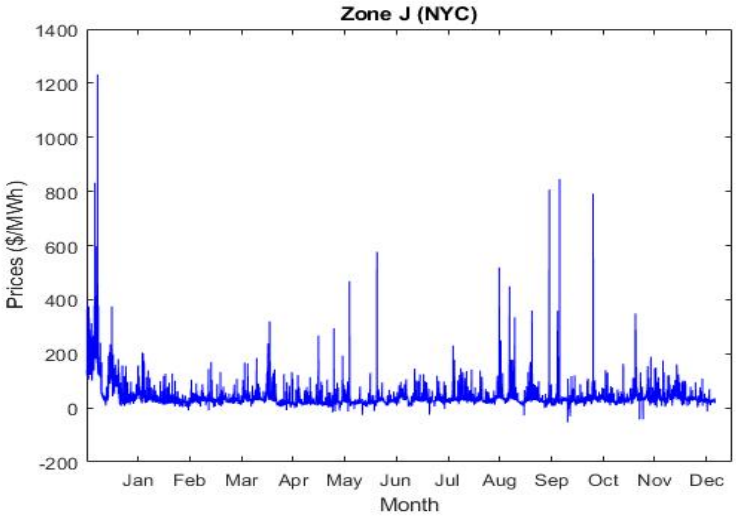

From Figure 6, it is difficult to see any clear seasonal patterns in prices, though they do appear elevated during the winter. We will see next that the NYISO power futures curve more clearly depicts seasonal price patterns as it measures expected average monthly prices as opposed to raw spot prices.

4.3.5 Risk management in the NYISO: Futures contracts

Many market participants wish to shield themselves from price risk by hedging because electricity prices can be extraordinarily volatile. Though there are several methods and products available that allow participants to hedge, one popular means is by obtaining futures contracts for power. A *futures contract* is a financial derivative where two parties agree to a pre-determined price, commonly called a "strike" price, for goods to be delivered on some future date. At the time of delivery, the buyer of the futures contract will pay the supplier the strike price regardless of the current prevailing spot market price and is therefore protected from spot market-price movements.

Futures contracts are typically standardized derivative products traded openly on exchanges. NYISO power futures are traded on several centralized exchanges such as the CME Group (formed from the mergers of the Chicago Mercantile Exchange, Chicago Board of Trade, and New York Mercantile Exchange) and the Intercontinental Exchange (ICE). In contrast to an internal bilateral transaction where two market participants directly negotiate specific contract terms, futures contract details such as strike price, quantity of deliverable power, duration of delivery, settlement rules, among others, constitute publicly available information for all market participants. Indeed, even members of the general public can easily access the details of power futures contracts traded on the CME or ICE.[9]

[9]Power futures contract information can be accessed directly on the CME Group website. https://www.cmegroup.com/trading/energy/electricity/nyiso-zone-a-5-mw-peak-calendar-month-day-ahead-lbmp-swap-futures-contract-specifications.html

Another important characteristic of futures contracts for power is whether they are for "peak" or "off-peak" power delivery. Peak power contracts cover all weekdays (Monday to Friday) from 07:00 to 22:00 when the demand for power is typically the highest. Peak contracts traded on the CME and ICE exclude North American Electric Reliability Corporation (NARUC) holidays such as July 4th and Thanksgiving when schools and businesses are closed and power demand will resemble a weekend day. Off-peak contracts cover periods of lower demand for power, that is, weekday nights (from 23:00 to 06:00), weekends (all day Saturday and Sunday), and NARUC holidays.

NYISO power futures are available for both day-ahead and real-time prices where the former are traded primarily on the CME and the latter on the ICE. Though it is possible to purchase futures contracts with annual or seasonal durations, the most commonly traded futures have a monthly duration. This means that the contract holder is entitled to a fixed amount of power at the contract strike price over an entire month.[10]

Futures contracts are traded for power in each zone because the NYISO is a zonal market where prices are settled within each zone. Figure 7 shows the monthly futures curve for power in six NYISO zones traded on the CME as of October 2019. These contract prices reflect the market expectations of average monthly power prices and are traded for up to five years for future delivery.

Figure 7: Forward curve for monthly peak day-ahead power futures in select NYISO zones

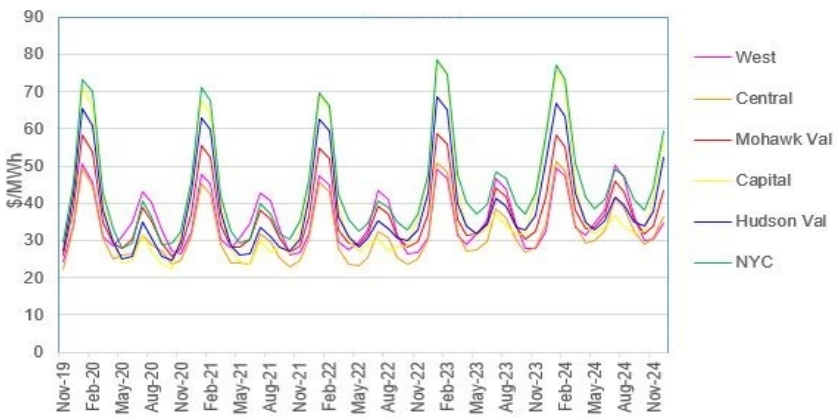

Although the prices differ between zones, there is a clear pattern of significantly elevated prices during the winter, slightly elevated prices during the summer, and low prices during the "shoulder months" of spring and autumn across all zones, which mirrors the patterns seen clearly in the load data. We will discuss why prices differ systematically between the NYISO zones in the following section.

[10]Contracts are sold in units of capacity but translate to units of power. For example, if an LSE holds a peak contract for 5 MW, they will pay the strike price for 5 MWh for each of the 16 peak hours during a weekday. Thus, the 5 MW futures contract actually translates to 80 MWh of power delivery over the course of a normal weekday.

4.3.6 The NYISO market landscape

New York possesses several geographic and demographic characteristics that have determined the location and fuel type of power plants within its borders, resulting in the evolution of the state's high-voltage power grid. These physical power grid attributes are not trivial and actively influence the NYISO's market dynamics.

The majority of New York State's population lives in or is in close proximity to New York City. In 2010, the US Census counted a total population of 19.3 million people living in the state of New York, 13.0 million of whom live in the dozen counties that comprise the New York City metropolitan area. The rest of the state is far more sparsely populated and the majority of the remaining 6.3 million inhabitants live in clusters around the northern cities of Buffalo, Rochester, Syracuse, and the state capital, Albany. Consequently, Zones G to K of the NYISO contain over two thirds of the state's population and account for the vast majority of power demand in the NYISO.

Although New York's population is heavily concentrated in the southeastern part of the state, the so-called "upstate" regions, represented by Zones A to E, contain a comparatively large share of the state's power generation capacity. Figure 8, below, was directly taken from the NYISO's 2019 "Power Trend Report," and it depicts the generation fuel mix in different regions of New York and has a simplified depiction of the high-voltage transmission network that spans the state.[11]

Figure 8: High-voltage transmission network and generation fuel mix in the NYISO

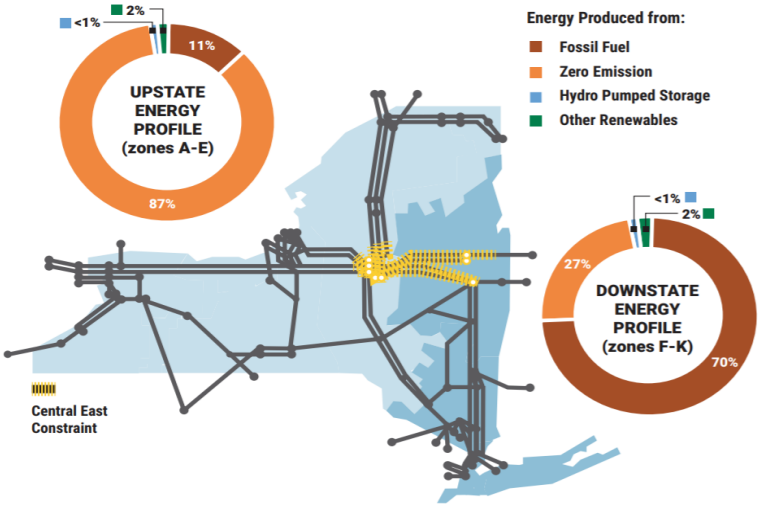

As Figure 8 shows, 87% of the installed generating capacity in the upstate region comes from zero-emission power plants, that is, hydro, wind, or nuclear generating resources. Most of the hydro, wind, and nuclear power plants in New York are situated in the northern part of the state around the Niagara region and in close proximity to Lake Ontario. The Robert

[11]The NYISO 2019 Power Trends Report can be accessed directly here: https://www.nyiso.com/documents/20142/2223020/2019-Power-Trends-Report.pdf/0e8d65ee-820c-a718-452c-6c59b2d4818b?t=1556800999122

Moses Niagara hydro plant, the largest in the state, is located in Zone A and other large nuclear and hydro plants hug the shores of Lake Ontario. In addition, the upstate region contains relatively few generators that burn coal, oil, or natural gas to produce electricity. An important economic characteristic of hydro and nuclear plants is that they have short run marginal costs that are effectively equal to zero. They therefore almost always offer to sell power to the market at lower prices than plants that burn fossil fuels.

Figure 8 also shows that the "downstate" zones, that is, Zones F to K, have a power generation fuel mix that is significantly different than the upstate region. Although the large Indian Point nuclear facility is located approximately 50 km from New York City, approximately 70% of the generating capacity downstate is comprised of conventional thermal power plants that burn natural gas and/or oil. Because they must pay for fuel, thermal plants have positive short-run marginal costs and almost always offer power to the market at higher prices than hydro or nuclear plants. During the winter and summer months, when the demand for electricity is very high, capacity limits on the major natural gas pipelines that feed into the New York City metro area may cause temporary fuel shortages for gas-fired power plants. This in turn forces these units to offer their power output at higher marginal prices to the market and regularly causes elevated prices and price volatility in the downstate zones (Patton et al., 2019).

The disparate locations of load and generation in the NYISO has also had consequences for the development of the high-voltage transmission network. Due to long distances and the heavily forested terrain of the Adirondack Mountains, the capacity of the high-voltage grid that connects the upstate and downstate regions is limited, frequently experiencing congestion. Figure 8 highlights the Central East Constraint, the major transmission bottleneck in the NYISO. When demand for power in the downstate zones becomes higher than these high-voltage lines are able to safely transmit, the NYISO effectively cleaves into two different sub-markets. During this situation, prices in the upstate zones typically decrease as lower power demand and cheaper marginal generating resources form a new market equilibrium. Conversely, prices in the downstate zones typically rise as higher demand must be met with more expensive marginal generating resources such as gas-fired generators.

Finally, the NYISO is connected to four adjacent balancing authorities: PJM Interconnection to the south, the New England ISO (ISO-NE) to the east, the Ontario Independent Electricity System Operator (IESO) to the northwest, and the Hydro-Quebec service territory to the northeast. PJM, ISO-NE, and IESO are wholesale power markets whereas Hydro-Quebec is a regulated, government-owned utility that is responsible for power generation and transmission in the Canadian province of Quebec. The high-voltage interconnectors between the NYISO and these four adjacent control areas allow market participants to import and/or export power out of the NYISO. As the IESO and Hydro-Quebec control areas are dominated by relatively inexpensive nuclear and hydro generators, producers in these markets regularly export power into the NYISO in an attempt to receive higher market prices

for their electricity output. Zone A, which constitutes the Buffalo metropolitan area and the westernmost portion of the state, is the NYISO zone directly connected to the IESO. Zone D, which constitutes the far northeastern portion of the state, is the NYISO zone directly connected to Hydro-Quebec. Because of their direct connection to these adjacent control areas, Zones A and D are frequently exposed to large influxes of cheap power from Canada that can cause transmission lines within these zones to become constrained. These internal zonal constraints in turn lead to periods of high price volatility in these two zones, which does not regularly occur in other upstate zones (Patton et al., 2019).

4.4 Data

From the NYISO website, we downloaded all available real-time hourly load and price data for all 11 zones.[12] The earliest available data for both series begins on January 1, 2006 and we proceeded to download all hourly data from this start date up to and including November 30, 2019. We used all data from 2006 to 2018 to estimate our models and all data from 2019 exclusively for out-of-sample forecasting. For each zone, we have 122,712 hourly observations of real-time prices and 122,712 observations of real-time load. Across all zones, this gives us a total of 1,349,832 observations each for real-time load and prices.

As discussed in Section 2, there are distinct intra-daily patterns of consumer electricity use, a fact reflected by the presence of futures contracts for peak and off-peak hours of power demand. We were able to obtain a large number of load and price observations. We then sorted these raw hourly data into four separate time periods; morning, midday, evening, and night. From a data analytics perspective, we believe that by splitting the days into four rather than two distinct time periods we are able to uncover a richer picture of the dynamic relationship between load and prices during different parts of the day.
The hours that constitute each time period are presented in Table 1. Although we deviate

Table 1: Hours contained within the four distinct time periods

Time Period	**Hours Beginning**
Morning	07:00, 08:00,09:00
Midday	10:00, 11:00, 12:00, 13:00, 14:00, 15:00
Evening	16:00, 17:00, 18:00,19:00, 20:00
Night	21:00, 22:00, 23:00, 00:00,01:00 02:00, 03:00,04.00, 05:00, 06:00

slightly from the standard definition of peak and off-peak as defined in the power futures contracts traded on the CME, we are effectively splitting the peak time period of HB 07:00 to HB 22:00 into three distinct time periods: morning, midday, and evening. The night-time period is almost identical to the hourly definition of off-peak except with the added hour of

[12]In the NYISO real-time market, load and prices are measured in 5-minute increments. All market quantities and prices, however, are ultimately settled on an hourly basis. The NYISO facilitates these hourly settlements by calculating the time-weighted/integrated (TWI) prices that are hourly averages of the 5-minute increments that fall within each hour. We therefore downloaded all of the load and price data labeled "TWI" on the NYISO website.

HB 21:00. We chose the hours contained in each time period based on the visual inspection of the average daily load curve discussed in Section 2.

At the same date and hour across all eleven NYISO zones, there were nine hourly load observations with zero or missing values. These missing or zero observations generally appear in the early years of the data and appear to be related to daylight savings time. Table 2 shows the date and hour for each of these nine missing values.

Table 2: Missing values in NYISO load and price data

Date	**Hours Beginning**
March, 4, 2006	05:00, 06:00,07:00
October, 29, 2006	01:00
November, 4, 2007	01:00
November, 2, 2008	01:00
November, 1, 2009	01:00

To address this issue, we estimated these missing observations by taking the simple average of the two preceding and two succeeding averages of load. In the instance of March 4, 2006 when there were multiple missing observations in a row, we first calculated the load value at 5:00 by taking the average of hours beginning 3:00, 4:00, 8:00, and 9:00. The load value at 6:00 was then calculated using the average of 4:00 using the average of our load estimates at 5:00 , 8:00, and 9:00. Finally, we calculated the load value at 7:00 using the average of our load estimates at 5:00 and 6:00, 8:00 and 9:00.

In addition, there were twelve negative observations for hourly load in Zone H. Most of these were concentrated in the hours beginning 00:00 to 06:00 on March 14, 2014. In contrast to prices or net power balance, we have no reason to believe that load can be a negative quantity as it has a lower bound of zero. Although we are suspicious that these negative load values indicate some sort of data entry error, they constitute a minuscule number of observations and we simply retain them in our dataset. The remaining 1,349,820 hourly load observations across all zones in the NYISO were positive.

4.4.1 Descriptive statistics: Load

Tables 3–6 present descriptive statistics of real-time hourly zonal load in the NYISO from 2006 to 2018. First, across all zones, the time varying load patterns are overwhelmingly consistent with the known patterns of consumer demand for electricity discussed in Section 2.

Table 3: Descriptive statistics of **load**. Time period: **Morning**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	1831	1167	1930	634	950	1389	1143	325	698	6170	2442
Median	1842	1175	1926	642	939	1382	1134	321	687	6083	2382
Mode	1842	1176	1896	468.	928	1415	1073	327	724	5191	2338
St. Dev	197	148	227.05	134.66	132.92	171.30	149.25	58.35	114.85	1033.06	413.38
Kurtosis	-0.06	0.42	-0.12	-1.36	-0.11	0.52	1.22	2.91	1.43	0.76	2.06
Skewness	-0.22	0.02	0.03	0.08	0.21	0.24	0.59	0.43	0.81	0.75	1.05
Minimum	1179	753	1305	289	579	844	547	-83	220	3199	499
Maximum	2649	1900	2907	963	1425	2222	2008	1141	1309	10769	4821
Observations	14244	14244	14244	14244	14244	14244	14244	14244	14244	14244	14244

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 4: Descriptive statistics of **load**. Time period: **Midday**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	1937	1255	2017	634	985	1461	1265	352	773	6837	2776
Median	1921	1228	1993	630	969	1419	1205	339	730	6572	2548
Mode	1942	1219	1863	494	900	1395	1224	301	740	6600	2462
St.Dev	216.21	181.41	233.69	131.11	131.82	207.94	233.96	73.16	154.06	1171.78	671.29
Kurtosis	0.91	2.35	0.87	-1.46	-0.03	1.75	2.43	1.36	1.98	0.95	1.71
Skewness	0.59	1.20	0.67	0.06	0.43	1.18	1.49	0.97	1.41	1.07	1.43
Minimum	1119	832	1426	291	534	962	649	135	293	3801	622
Maximum	2921	2198.5	3153.4	1219	1558.6	2402.9	2495.1	729.5	1530.3	11442.9	5935
Observations	28488	28488	28488	28488	28488	28488	28488	28488	28488	28488	28488

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 5: Descriptive statistics of **load**. Time period: **Evening**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	1957	1279	2054	642	1008	1525	1344	383	808	6825	2972
Median	1941	1258	2030	644	996	1498	1299	375	774	6572	2814
Mode	1931	1208	2112	734	880	1542	1184	339	768	6104	2766
St.Dev	212.12	178.91	241.39	133.35	138.16	213.95	234.95	77.86	149.31	1089.66	631.03
Kurtosis	0.29	1.41	-0.02	-1.36	-0.41	0.88	1.55	0.52	1.97	1.39	1.55
Skewness	0.40	0.89	0.45	0.08	0.29	0.79	1.15	0.53	1.32	1.21	1.25
Minimum	1159	824	1406	331	576	946	670	-172	404	4157	713
Maximum	2908	2188	3098	964	1540.1	2425	2495	726	1544	11456	5915
Observations	23740	23740	23740	23740	23740	23740	23740	23740	23740	23740	23740

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 6: Descriptive statistics of **load**. Time period: **Night**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	1629	993	1661	596	784	1176	1005	277	603	5315	2149
Median	1606	970	1636	603	767	1151	972	269	575	5089	2035
Mode	1524	946	1726	432	661	1106	879	229	536	5073	1743
St.Dev	195.29	151.38	229.99	132.39	141.80	198.43	185.13	68.14	127.48	1014.98	481.74
Kurtosis	0.53	1.44	0.19	-1.36	-0.07	0.93	2.29	1.56	2.29	1.53	3.04
Skewness	0.68	0.97	0.63	0.08	0.58	0.78	1.20	0.59	1.28	1.19	1.43
Minimum	1025	668	1173	267	300	702	445	-342	235	2859	447
Maximum	2750	2038	2872	990	1452	2306	2258	862	1412	10659	7089
Observations	47480	47480	47480	47480	47480	474800	47480	47480	47480	47480	47480

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

The mean and median load values for all zones are lowest during the night when people are asleep. Average load then increases sharply during the morning when people wake up, and climbs higher still during the midday hours when people are at work and school. With the exception of Zone J, which corresponds to New York City, mean and median load values are highest during the evening when people return home from work and increase their use of large appliances, lighting, and HVAC systems. The mean and median values for load in Zone J are highest during the midday time period. At the same time, Zones G, H, and K, corresponding to the suburban areas outside of New York City, exhibit the largest increases in mean and median load from midday to evening. Taken together, these statistics seem to reveal a pattern where commuters from the suburbs enter New York City for work during the day and then return home to the suburbs during the evenings.

As expected, Zones J and K, where Zone K corresponds with the heavily populated Long Island, consistently exhibit the highest mean and median load values across all time periods. Zones D, H, and I exhibit the lowest mean and median load values across all time periods. This result is not surprising due to the fact that Zone D corresponds to sparsely populated territories in the Adirondacks, and Zones H and I are geographically small zones that correspond with the Northern suburbs of New York City.

Across almost all zones, the standard deviation and variance values for load are higher during the midday and evening and lower during the morning and night. In Zones D and E,

the exceptions to this trend, the standard deviation and variance of load values were almost identical across all time periods. In terms of absolute magnitude, Zones J and K show the largest standard deviation and variance values, where Zone J shows standard deviations of over one thousand in each time period. Zone H exhibits the lowest standard deviations in each time period.

In contrast to the time period trends in mean and standard deviation, the overall time period trends for load ranges are weaker. In eight out of eleven zones, however, range values are lowest during the morning, indicating that load during this time period is generally more uniform than in other time periods.

In addition, the midday and night time periods show the highest range values in nine out of the eleven zones, indicating potentially more load value uncertainty during these time periods. Across all time periods, the relative magnitude of range values closely follows the patterns seen in the mean, median, and standard deviations. Zone J has by far the largest range values across all time periods; the range for each time period in this zone exceeds 1,000 MWh. Zone K shows the second highest ranges across all time periods. Zones D, E, and H have the smallest ranges across all time periods.

We next examine the skewness and kurtosis of the load. Skewness is the statistical measurement of symmetry in a given distribution. A skewness of zero indicates that a distribution is perfectly symmetrical. In the context of this discussion, a positive skewness indicates that there is an asymmetrically higher probability of large load values and lower probability of small load values. A negative skewness indicates an asymmetrically higher probability of small load values. Kurtosis is the statistical measure of the tails of a distribution, that is, the likelihood of outcomes that lie far away from the mean. In finance, kurtosis is frequently called "tail risk" and serves as an important measure of the likelihood of "extreme" events that lie very far away from the mean of the distribution. Distributions that exhibit kurtosis higher than the normal distribution are called "leptokurtic" or sometimes more informally "fat-tailed" to indicate the higher likelihood of events taking place in the extreme regions of the distribution. In risk management terms, an economic variable that possesses a distribution with large kurtosis values is inherently riskier and a risk-averse market participant would seek to mitigate this risk through hedging.

In Tables 4 and 6, we see that almost all zones exhibit skewness that is positive or close to zero. Across all zones, midday load generally exhibits the highest skewness and morning load generally exhibits the smallest skewness. The downstate zones exhibit significantly larger positive skewness than the upstate zones across all time periods with the exception of Zone B, which shows high positive skewness during midday, evening, and night. The midday, evening, and night time loads in Zones G, I, and J have a positive skewness larger than one, and the load in Zone K has a positive skewness greater than one across all time periods. Though they share some similar trends with the skewness values, the kurtosis values for load reveal additional market dynamics. The time-varying trend of kurtosis across all zones is

generally the same as for skewness, that is, highest during the middays and lowest during the mornings. Notable exceptions to this trend are Zones J and K where the kurtosis is highest during the night. All of the downstate zones (F through K) exhibit leptokurtosis with Zones G, I, and K showing kurtosis values larger than one for every time period. With kurtosis values approximately equal to three, Zones H and K exhibit the highest absolute kurtosis values during the morning and night, respectively. However, there is no one clear trend that dominates the upstate zones' kurtosis values. While Zone B shows significant leptokurtosis much like the downstate zones, Zones A, C, and E all generally show low kurtosis values outside of the midday time period. Zone D exhibits negative kurtosis across all time periods, which indicates a thin-tailed distribution where load values that are extremely far from the mean are less likely to occur than in a standard normal distribution.

Figures 9 and 10 are histograms of hourly evening load observations for Mohawk Valley (Zone E) and New York City (Zone J) from January 1, 2006 to December 31, 2018. In both figures, we have superimposed a normal distribution over the observed distribution. Figure 9 shows that the load in Mohawk Valley has historically followed a distribution that is almost perfectly normal with a kurtosis close to zero and some slight positive skewness. Figure 10, however, shows a more typical distribution of load in the NYISO, especially in the downstate zones. It is highly positively skewed and the high concentration of large load values that lie outside of the normal distribution depict leptokurtosis.

Figure 9: Histogram of evening load in Zone E

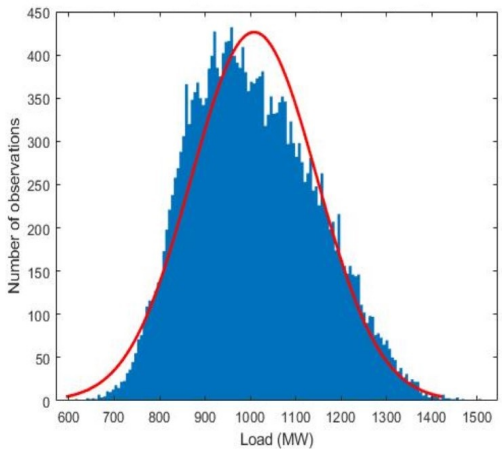

Figure 10: Histogram of evening load in Zone J

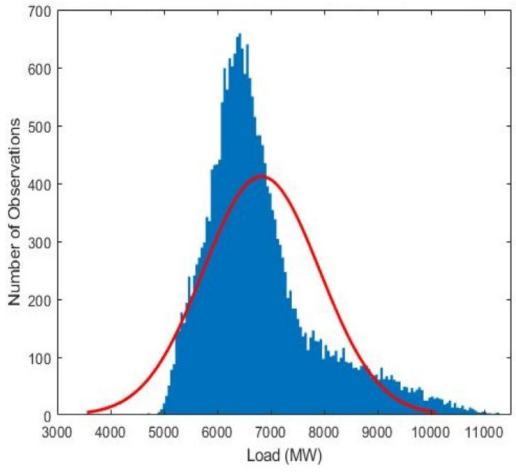

Taken together, the statistics for skewness and kurtosis indicate that there is an elevated risk of "extreme" high-load values in the downstate Zones F to K and Zone B when compared to a normal distribution, that is, most zones exhibit a distribution depicted in Figure 10. Other than in Zone B, the risk of these extremely large load values is generally lower in the upstate zones where kurtosis is small in magnitude or even negative. Figure 9 is more representative of the distribution of load in these upstate zones.

4.4.2 Descriptive statistics: Prices

We now turn our attention to hourly real-time prices. Tables 7–10 present descriptive statistics of real-time hourly zonal power prices in the NYISO from 2006 to 2018. Each table presents the descriptive statistics of prices from a given time period. We see several interesting trends in these tables. First, across all time periods, the mean and median prices in the downstate zones (F–K) are consistently higher than those in the upstate zones (A–E). Average prices in the downstate zones are generally \$10 to \$20/MWh higher than those in the upstate zones and mean prices in Zone K are approximately \$30/MWh higher. Zones J and K consistently exhibit the highest mean and median prices across all time periods due to the fact that these two zones have by far the highest load levels in the NYISO and also frequently rely on gas- and oil-fired generators that have high marginal costs to balance load during periods of high demand. Zone D shows the lowest mean values for prices across all time periods, which reflects the presence of inexpensive hydro and nuclear generators whose capacity far exceeds the local demand for power.

Table 7: Descriptive statistics of **prices**. Time period: **Morning**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	39.46	39.64	41.74	37.96	42.77	48.60	48.81	49.02	49.0	52.59	56.98
Median	32.20	32.89	34.28	31.91	35.00	37.82	38.80	38.82	38.83	40.11	42.17
Mode	30.18	31.83	29.77	26.49	32.43	32.67	23.47	41.37	35.77	40.68	33.83
St.Dev	38.19	36.19	38.76	47.25	40.54	44.99	43.75	44.05	43.96	49.01	54.52
Kurtosis	149.5	168.5	175.3	278.1	179.9	122.6	136.3	133.5	130.8	98.97	106.6
Skewness	8.31	8.25	8.67	-2.95	8.79	7.27	7.63	7.55	7.47	6.75	3.56
Minimum	-294.3	-324.2	-306.2	-1613	-322.9	-153	-156.9	-156	-156.1	-156.7	-1476
Maximum	1135	1183	1209	1236	1283	1266	1236	1232	1224	1231	1210
Observations	14244	14244	14244	14244	14244	14244	14244	14244	14244	14244	14244

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 8: Descriptive statistics of **prices**. Time period: **Midday**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	43.74	41.97	44.08	39.81	45.08	52.23	55.40	56.29	56.57	61.04	70.11
Median	34.49	35.04	36.71	33.97	37.5	40.76	42.45	42.58	42.68	44.85	49.74
Mode	30.77	31.48	35.84	26.98	31.93	38.32	35.03	36.85	37.32	36.6	31.71
St.Dev	47.50	37.70	38.93	40.53	39.70	46.95	56.06	60.41	61.29	62.94	74.64
Kurtosis	129.3	139.6	130.5	133.6	127.4	83.83	108.5	123.2	125.3	87.88	69.64
Skewness	8.17	7.07	6.92	2.66	7.09	5.92	8.26	8.92	9.02	7.40	6.43
Minimum	-716.9	-750.5	-759.1	-1069	-711.8	-743.3	-154.4	-165.4	-142.1	-53.7	-162.3
Maximum	1335	1073	1074	1024	1091	1079	1257	1418	1438	1285	1709
Observations	28488	28488	28488	28488	28488	28488	28488	28488	28488	28488	28488

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 9: Descriptive statistics of **prices**. Time period: **Evening**

Zones	**Zone A** (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
Mean	46.71	46.94	49.01	45.19	50.31	57.43	61.55	62.66	62.93	67.12	80.53
Median	35.56	36.63	38.1	35.6	39.09	42.93	45.25	45.45	45.54	48.16	57.14
Mode	34.49	31.69	37.05	31.29	36.23	35.88	29.34	24.52	33.46	45.03	37.2
St.Dev	47.71	44.36	45.76	46.48	46.94	53.51	61.05	65.27	66.01	66.42	84.39
Kurtosis	91.87	95.76	87.69	85.65	86.52	70.27	81.63	99.52	100.8	68.99	57.61
Skewness	6.15	5.54	5.33	3.77	5.55	5.24	6.84	7.57	7.63	6.26	5.62
Minimum	-666.1	-688.9	-706.7	-815.8	-663.5	-693.4	-227.7	-137.5	-136.9	-138.1	-128.7
Maximum	1183	1202	1192	1178	1247	1288	1327	1499	1527	1389	1800
Observations	23740	23740	23740	23740	23740	23740	23740	23740	23740	23740	23740

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Table 10: Descriptive statistics of **prices**. Time period: **Night**

Zones	Zone A (West)	Zone B (Genesee)	Zone C (Central)	Zone D (North)	Zone E (Moh.Valley)	Zone F (Capital)	Zone G (Hud.Valley)	Zone H (Millwood)	Zone I (Dunwoodie)	Zone J (NYC)	Zone K (L.Island)
Mean	30.07	30.59	31.92	29.89	32.77	39.68	38.88	38.98	38.97	41.02	45.52
Median	27.39	27.78	28.71	27.32	29.33	32.31	32.6	32.59	32.59	33.4	35.37
Mode	29.95	25.8	32.6	25.08	33.54	35.72	35.7	21.75	31.99	32.7	27.98
St.Dev	27.07	27.28	27.95	35.15	29.27	35.36	32.71	32.73	32.66	34.83	40.37
Kurtosis	217.93	133.65	104.87	884.91	102.64	62.63	74.73	70.58	70.35	55.66	41.28
Skewness	5.07	3.64	3.32	-13.58	3.19	4.09	4.11	3.93	3.91	3.60	3.48
Minimum	-469.8	-522.6	-498.9	-2482.1	-519.1	-415.3	-352.7	-350.8	-351.2	-352.1	-357.5
Maximum	1446	1226	1142	1171	1185	1204	1202	1204	1201	1202	1201
Observations	47480	47480	47480	47480	47480	47480	47480	47480	47480	47480	47480

Note: Sample period, hourly observations, 2006:01:01-2018:12:31

Second, across all zones, the mean and median values for prices are lowest during the night, increase during the morning, climb higher during the midday, and reach their highest values during the evening. As with load, these results are overwhelmingly consistent with known consumer behavior and reflect demand fluctuations over the course of a day.

In contrast to the low volatility seen in the load, we see dramatically higher volatility in prices, putting into sharp focus the need for NYISO market participants to hedge against price risk. Indeed some of the statistics presented in Tables 7–10 are extreme by most measures. The standard deviations of prices largely follow the same patterns as the mean and median values. They are considerably higher in the downstate zones than they are in the upstate zones across all time periods, and Zone K exhibits the highest values across all time periods. In each zone, the standard deviations are lowest during the night and, with the exception of Zone D, highest during the evenings. In addition, due to the large presence of windmills and hydro imports from Quebec, we see that Zone D exhibits uncharacteristically high standard deviations during the night and morning compared to other upstate zones. This is because wind speeds are highest and the demand for power is lowest during the night. As a result, unexpected strong winds during the night time can lead to surges in supply that cause large price swings.

Across all zones, price ranges are generally highest during the midday and the evening when demand for power is highest. As discussed in Section 4, prices are normally capped at $1,000 per MWh but since most zones have maximum values around $1,200 per MWh, we see evidence of some extreme weather events where fuel shortages caused generators to incur costs that exceeded the price cap. Ranges do not appear to vary dramatically between zones, although since minimums are generally lower in upstate zones, the ranges in upstate zones are generally a bit higher than in downstate zones. Once again, Zone D stands out from other zones: with minimum values of $-1,613.64 per MWh and -$2,482.04 per MWh for morning and night, respectively, it exhibits enormous ranges. For comparison, the other upstate zones have minimum values of approximately -$400.00 per MWh during the night and -$300.00 per MWh during the morning.

With the exception of the morning and night prices in Zone D, prices in all zones and time periods exhibit high positive skewness—a result that comes as no surprise given the presence of large, positive price spikes. Zone D is negatively skewed due to the large negative price spikes that have occurred during the morning and night. In almost all instances, skewness values are lowest during the nights. With the exception of Zone D, skewness values are

highest during morning in the upstate zones and Zone F. Price skewness values are highest during the midday time period in the remaining downstate Zones G to K.

All prices in all zones exhibit very high levels of leptokurtosis indicating a so-called "fat tailed" distribution consistent with price-spiking behavior. Kurtosis values are generally higher in upstate zones than in downstate zones and tend to be the highest during the morning and midday time periods across all zones. Prices in Zone D during the night exhibit a staggeringly large kurtosis of \$884.9 per MWh , a result driven by an extremely minimum value of \$-2,482.04 per MWh that lies over 70 standard deviations to the left of the mean price in this time period. Perhaps somewhat surprisingly, the corresponding zones to New York City and Long Island (Zones J and K) actually show some of the lowest kurtosis values across all time periods. We suspect that this is because prices rarely become negative in these zones, thereby resulting in statistical outliers that are overwhelmingly positive.

Figures 11 and 12 are histograms of the evening price observations for Zone J (New York City) and the night time price observations for Zone D from January 1, 2006 to December 31, 2018. In both figures, we have superimposed a normal distribution over the observed distribution.

Figure 11: Histogram of evening prices in Zone J

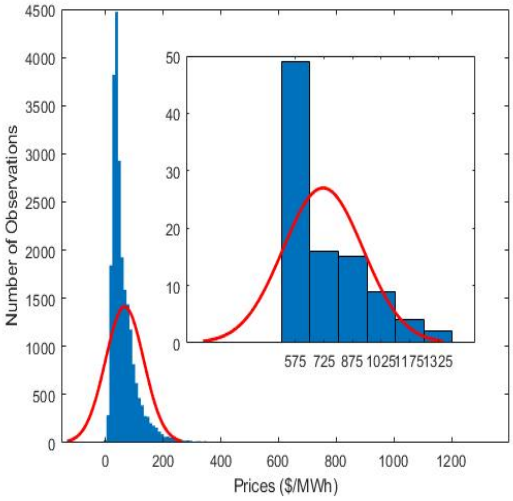

Figure 12: Histogram of night prices in Zone D

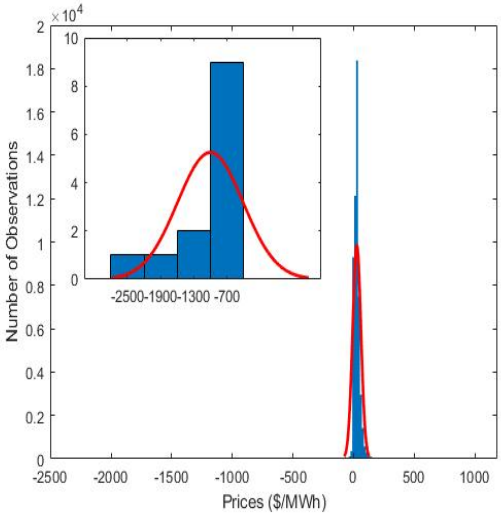

Figure 11 is representative of the price distribution in almost all time periods for all NYISO zones as presented in Table 7. In this figure, the prices are heavily clustered around the median value of \$48.16 per MWh but we can clearly see the positive skewness and large kurtosis of the distribution caused by large positive price spikes. In Figure 12, prices are also clustered heavily around the median value, but several large negative price spikes cause the distribution to be skewed negatively and to have an extremely high kurtosis.

4.4.3 Implications for risk management

The descriptive statistics of hourly real time prices presented in this section underscore the importance of hedging for NYISO market participants that are sensitive to price risk. As we have seen in this section, prices are highly leptokurtic and can exhibit spikes upwards of 3500% higher than the mean price. For LSEs that must purchase large quantities of electricity, such price movements could spell catastrophic losses of millions of dollars over the course of a single hour. LSEs can reduce their exposure to such spikes by purchasing financial derivatives. Though the majority of zones exhibit positively skewed price distributions, we see that in the case of Zone D there is a possibility of extreme negative price spikes during the night, which could be similarly catastrophic for generators. These generators could sell their power forward via futures contracts to hedge against such occurrences.

We also wish to highlight that the descriptive statistics of prices clearly illustrate the logic behind offering both peak and off-peak futures contracts. Mean prices are significantly lower during the night than other time periods, a fact reflected in the lower prices of power in off-peak futures contracts sold on the CME and ICE. Although they are generally less volatile than peak hours, night time prices still show high levels of skewness and kurtosis, necessitating derivatives capable of allowing market participants to mitigate price risk during these periods of lower demand.

4.5 Methodology

In this section, we will provide an overview of the methods used to analyze the relationship between load and prices in the NYISO. We begin by describing the steps taken to pre-process and filter the raw NYISO market data, which we carry out in order to smoothen the data. We first sort the raw data by aggregating hourly observations into four distinct time periods that span the course of a day. Next, after performing a series of transformations on the sorted data, we calculate three different moving averages for both load and prices: an annual MA, a 90-day seasonal MA, and a weekly MA. Finally, we present three non-linear functions that describe the moving averages of load as a function of a series of periodic expressions and the corresponding MA of prices. The parameters of these functions are estimated using nonlinear least squares.

To avoid confusion, we will refer to the three functions introduced in this section as "models" or "waveform functions." We do this to highlight the fact that these equations are in a sense deterministic since they are built largely from Fourier-like terms and are quasi-periodic. In contrast, we reserve the word "regression" to describe the functions used to analyze the residuals of these models in Section 7. We do this to emphasize the fact that we are using the Box–Jenkins approach to analyze these residuals, an approach that uses purely statistical diagnostic methods.

4.5.1 Data pre-processing and filtering

Before performing our analysis, we carry out several pre-processing and filtration steps to further transform the raw NYISO load and price data. We employ these steps to smoothen the raw data and allow for better parameter estimation from our models. In addition, since LSEs purchase futures contracts for power that have monthly durations, these steps are meant to compress the short-run dynamics of NYISO load and prices into measures that capture more of their long- and medium-run dynamics.

First, we calculate the daily arithmetic mean of load during each time period and across all zones. For example, we calculate the average morning load for Dunwoodie on January 1st, 2006 by summing together the raw load values at hours 7:00, 8:00, and 9:00 and by dividing this sum by three. This procedure is repeated for the midday, evening, and night time periods across all zones and all 4,748 days in the sample period. We add the load values for each time period and divide by the number of hours in the given time period. With four time periods and 11 zones this yields 208,912 observations of average hourly load in each time period.

Next, we transform the raw price data by calculating daily load-weighted average prices across all time periods and zones. We begin this calculation by multiplying the hourly real-time price of electricity by the corresponding hourly load. All total hourly payments within a given time period are then summed together. Finally, this total sum is divided by the total load consumed in the same given time period.

For example, assume that on a given day in New York City consumers demanded electricity as follows: 5,000 MWh at 7:00 AM; 5,500 MWh at 8:00 AM; and 6,000 MWh at 9:00. The respective prices were 30\$/MWh at 7:00 AM, 35\$/MWh at 8:00 AM, and 40\$/MWh at 9:00 AM. The load-weighted average price for the New York City zone at this particular morning observation would then be

$$\frac{(5000MWh \cdot 30\frac{\$}{MWh}) + (5500MWh \cdot 35\frac{\$}{MWh}) + (6000MWh \cdot 40\frac{\$}{MWh}}{(5000MWh + 5500MWh + 6000MWh)} = 35.30\$/MWh$$

Finally, using our transformed load and price data, we calculate a series of moving averages:

$$Moving\ Average\ \Phi_{z,p,t} = \frac{1}{w} \sum_{s=t-(w-1)}^{t} \phi_{z,p,t}$$

where Φ denotes either the moving average of load or prices, ϕ denotes the daily load and prices, w denotes the moving average window, z denotes the zone, p denotes the time period, and t denotes the day of observation. In the three following subsections, we will introduce the moving average data series used in each model and any additional model-specific data transformations.

4.5.2 Annual moving average model

In this section, we introduce a model for the annual MA of load. First, we calculate the annual moving averages of real-time load and prices for each time period in every zone, yielding a total of 88 transformed time series. Since we must use the previous 365 days of the raw data to calculate the first value of the annual MA, the beginning date of the new time series is January 1st, 2007. Like the original values of load and prices, the annual moving averages for load are measured in MWh and the annual moving averages for prices are measures in \$/MWh. As an example of what these new time series look like, Figures 13 and 14 show the annual MA time series for load and prices in New York City (Zone J) during the evening. We first focus our attention on the graph showing the annual MA of

Figure 13: The annual moving average of load in Zone J

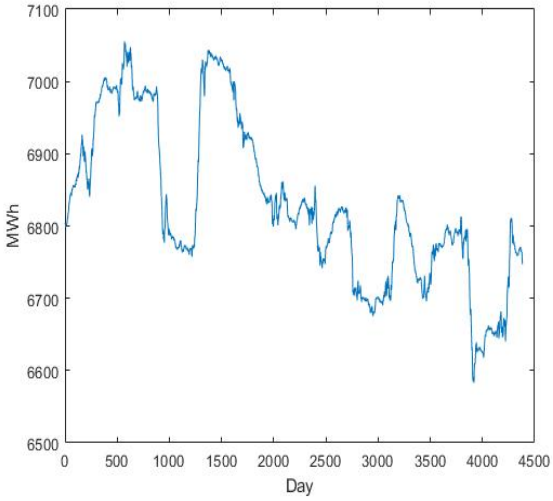

Figure 14: The annual moving average of prices in Zone J

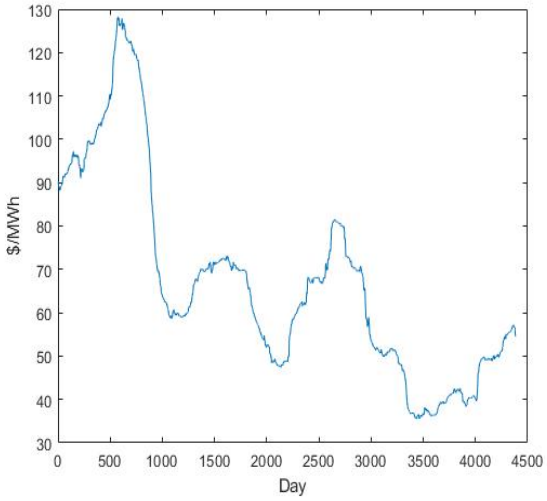

load. We notice immediately that, like the raw load data, the average MA of load appears to follow a quasi-cyclical pattern. Although it is not perfectly uniform across the whole time series, there seems to be a frequency that spans approximately 1,400 days. In contrast, however, the amplitude of this new time series appears quite unpredictable as it shrinks and expands at seemingly irregular intervals. Finally, the data appear to trend downwards over time although the cyclicality persists throughout the time series. The graph depicting the annual MA of prices possesses many of the same characteristics as load. The data follow a downward trending quasi-cyclical pattern with an approximate frequency of 1,400 days and irregular shifts in amplitude.

The prices, however, appear to be somewhat smoother than the load; there are very few sharp spikes or valleys that we see in the time series for load. In addition, there is a precipitous drop in the annual MA of prices that begins around $t = 600$ and persists for the remainder of the series. This point seems to correlate 2009 as the shale gas revolution in the eastern US, when the market was flooded with natural gas, making previously expensive

gas-fired power plants far cheaper to operate. Finally, when comparing both of the graphs, there is a clear correlation between the two series: despite differences in amplitude, they appear synchronous and exhibit the same approximate downward drift.

Taking into consideration our visual inspection of the annual moving averages for load and prices in NYISO, we choose the following non-linear model specification to explain the annual MA of load:

$$\begin{aligned} L_{z,p,t} = \alpha_{z,p} + \beta_{1,z,p} \cdot \exp(\sin(\frac{\pi t}{738})) \cdot \log(P_{z,p,t}) + \beta_{2,z,p} \cdot P_{z,p,t} \\ + \beta_{3,z,p} \cdot t \cdot P_{z,p,t} + \beta_{4,z,p} \cdot \log(t) + \varepsilon_{z,p,t} \end{aligned} \tag{1}$$

where $L_{z,p,t}$ is the annual moving average of load, $P_{z,p,t}$ is the annual moving average of prices, t is a numeric time indicator of the observation, $\varepsilon_{z,p}$ is the residuals, and $\alpha_{z,p}$ and the four $\beta_{i,z,p}$ terms are parameters to be determined. Taken together, these components allow us to capture a large range of potential relationships between the AMA load and prices in the NYISO.

The term $\alpha_{z,p}$ is the "baseline" value of the annual moving average of the load in a given zone and time period. In other words, this term measures the long-run average value of load after controlling for time-dependent fluctuations in weather and prices.
The three coefficients $\beta_{1,z,p}$, $\beta_{2,z,p}$, and $\beta_{3,z,p}$ collectively measure the relationship between the annual moving averages of load and prices.

The expression, $\exp(\sin(\frac{\pi t}{738}))$, is a periodic function and its amplitude is modified by the logarithm of prices to better capture the irregular nature of the load's amplitude as it changes over time. We take the logarithm of these prices to compress the changes of amplitude of the price series as it generally has a higher magnitude than the corresponding amplitude in the load series. When multiplied together, these two components form an interaction term that approximates both the frequency and the amplitude of the annual moving average of load. In practical terms, it allows us to take a more direct measurement of the cyclical relationship between long-run load and prices in a given zone and time period. This relationship is captured in the coefficient $\beta_{1,z,p}$. Because this coefficient is assigned to a term that uses log prices, it measures the average unit MWh increase in the annual moving average of load given a unit percentage increase in the annual moving average of price.

The second coefficient, $\beta_{2,z,p}$, is straightforward to interpret: it measures the linear component of the relationship between the annual moving averages of load and prices. More specifically, $\beta_{2,z,p}$ measures the average response of the annual moving average of load to a one unit increase in the annual moving average of prices after controlling for seasonality and potential long-run changes in the load and price relationships captured in the other terms.

The third coefficient, $\beta_{3,z,p}$ measures the interaction term between time and the annual average prices. $\beta_{3,z,p}$ is therefore a sort of inflationary measurement that allows for changes in the relationship between load and prices to evolve over time. If $\beta_{3,z,p}$ is close to zero,

it indicates that the relationship between prices and load vary little over time, that is, this relationship is approximately the same in 2007 as it is in 2018. If this coefficient is not equal to zero, then we have evidence that the relationship between these two time series is changing over time.

The fourth and final coefficient in the annual model,$\beta_{4,z,p}$, allows for the possibility of load growth or shrinkage over time, which is not directly correlated to changes in power prices. Long-run changes in the demand for power in a given zone may be due to growing or shrinking population, changes in energy efficiency, and so on. If this coefficient is close to zero it indicates that there is little or no change in the baseline load levels over time in a given zone. If this coefficient is positive, it indicates that the baseline level of load is increasing over time in a given zone. If this coefficient is negative, it indicates that the baseline level of load is decreasing over time.

We used a nonlinear least squares to estimate the model parameters. Non-linear optimization in MATLAB requires an initial guess value of the model parameters to facilitate numerical convergence. For all 44 zonal and time-period runs of the annual model, our initial guess value for $\alpha_{z,p}$ was equal to 1,800 and our initial guess values for all four $\beta_{i,z,p}$ coefficients were equal to 1.

4.5.3 Seasonal moving average model

We next present our seasonal model which seeks to estimate the relationship between load and prices that have been deseasonalized at the annual and 90-day time windows. From the perspective of an LSE looking to purchase futures contracts for electricity, we consider this model to be of more importance than the annual model. Since winter and summer are periods of peak electricity demand in the NYISO, it is especially important for an LSE to be able to understand their price risk during these periods. With accurate load forecasts and a model that can accurately forecast the seasonal relationship between load and prices, LSEs can make better-informed hedging decisions given, the current futures market prices.

We begin by dividing the daily observations of load during a given time period and zone by the annual MA of load in the same time period and zone. We then transform this quantity into a percentage by subtracting one. For example, on April 1, 2007, the morning load value for Zone I was 548.8 MW and the annual MA for load on the same date was 723.7 MW. Our percentage of the morning load to the annual MA is therefore $\frac{548.8MW}{723.7MW}$-1=-0.242. Finally, we calculate the 90-day moving average of this percentage to calculate the seasonal moving average (SMA) of load. We repeat this process to calculate the SMA prices for all zones and time periods. We will henceforth refer to these newly transformed data series as the "seasonal" load and prices.

Equation (2) depicts our model for the seasonal MA of load:

$$L_{z,p,t} = \zeta_1(t) + P_{z,p,t} \cdot \zeta_2(t) + \varepsilon_{z,p,t} \tag{2}$$

$$\zeta_1 = \beta_{1,z,p} \cdot \sin(\frac{\pi t}{91}) + \beta_{2,z,p} \cdot \cos(\frac{\pi t}{91}) + \beta_{3,z,p} \cdot \sin(\frac{\pi t}{182}) + \beta_{4,z,p} \cdot \cos(\frac{\pi t}{182}) \beta_{5,z,p} \cdot \sin(\frac{\pi t}{364})$$
$$+ \beta_{6,z,p} \cdot \cos(\frac{\pi t}{364}) + \beta_{7,z,p} \cdot \sin(\frac{\pi t}{728}) + \beta_{8,z,p} \cdot \cos(\frac{\pi t}{728}).$$
$$\zeta_2 = \beta_{9,z,p} \cdot \sin(\frac{\pi t}{728}) + \beta_{10,z,p} \cdot \cos(\frac{\pi t}{728}) + \beta_{11,z,p}.$$

where $L_{z,p,t}$ and $P_{z,p,t}$ are the seasonal moving average of load and prices, respectively, t is a numeric time indicator of the observation, $\varepsilon_{z,p,t}$ is an error term, and the eleven $\beta_{i,z,p}$ terms are parameters to be decided. Because the SMA load and prices are measured in percentage deviations from the annual moving average, we assume a linear intercept equal to zero.

The seasonal model closely resembles a Fourier series, that is, a weighted sum of periodic functions. A Fourier series is a commonly used modeling paradigm for signal processing as it is capable of decomposing complicated waveforms into a sum of weighted sine and cosine pairs.[13]

The first eight terms of our seasonal model are periodic functions that are purely dependent on time. These eight terms are grouped in successive Fourier sine and cosine pairs that cover four frequencies: 182 days, 364 days, 728 days, and 1,446 days. These four pairs of trigonometric terms are intended to capture the medium and long-term trends that may be present in the seasonal data.

In addition to the purely time-dependent explanatory variables in the seasonal model, there are three terms that include the seasonal MA of prices. The coefficients β_9 and β_{10} measure the interaction terms between sine and cosine functions with periodicities of approximately four years and the seasonal MA of prices. These two terms effectively extended the model with terms that allow long-term prices to modulate the amplitude of the seasonal load.

The final model term, β_{11}, measures the linear, non-cyclical relationship between the seasonal moving averages of load and prices. This term essentially measures the average change in load correlated with a unit change in prices after controlling for seasonality in the two series.

4.5.4 Weekly moving average model

We now present our third and final model, which maps the relationship between the seven-day moving averages of load and prices. The weekly moving average time window is significantly shorter than the annual and seasonal moving averages and might be useful for an LSE that seeks to take into account shorter term market dynamics when updating estimations for price or load values. LSEs could therefore use this model in tandem with the annual and seasonal models to fine tune estimates for prices or load if they are considering changing their futures

[13] Sine and cosine functions are orthogonal functions and both are necessary to map all possible forms of a periodic function at a given frequency.

contract portfolio at the beginning of a new delivery month.

The process used to create the weekly MA data series is identical to that used to create the seasonal moving averages; for all zones and time periods, we divide the daily average real time load and prices by their corresponding annual MA values and subtract this quantity by one to transform it into relative deviations. Next, we subtract the seasonal MA from these relative deviations. We then calculate the 7-day MA of these data to get our weekly moving average (WMA) load and price data series.

Equation (3) depicts our model for the weekly moving average of load:

$$\begin{aligned} L_{z,p,t} = \Bigg(& \beta_{1,z,p} \cdot \sin\left(\frac{2\pi t}{182.47}\right) + \beta_{2,z,p} \cdot \cos\left(\frac{2\pi t}{182.47}\right) + \beta_{3,z,p} \cdot \cos\left(\frac{\pi t}{182.47}\right) \\ & + \beta_{4,z,p} \cdot \sin\left(\frac{\pi}{a}\right) \cdot P_{z,p,t} \Bigg) \cdot \frac{t}{b+t} + \beta_{5,z,p} \cdot \left(\frac{P_{z,p,t}}{c+t}\right) + \beta_{6,z,p} \cdot P_{z,p,t} + \varepsilon_{z,p,t} \end{aligned} \tag{3}$$

where $L_{z,p,t}$ is the weekly moving average of load, $P_{z,p,t}$ is the weekly moving average of prices, t is a numeric time indicator of the observation, $\varepsilon_{z,p,t}$ is an error term, and the six $\beta_{i,z,p}$ terms are parameters to be determined. In addition, this model includes three scaling parameters, a, b, and c, that we can use for frequency and amplitude modulation. As in the seasonal model, we are again modeling a time series measured in percentage deviations from the annual MA and therefore assume a linear intercept equal to zero.

The weekly model borrows elements found in both the annual and seasonal models. The coefficients β_1, β_2, and β_3 measure time-dependent cyclical fluctuations in the WMA load, precisely like the terms found in the seasonal model.

The coefficients β_4, β_5, and β_6 measure the relationship between WMA load and prices. The coefficients β_4 and β_5 measure time and price interaction terms that are similar to those found in the annual and seasonal models, albeit with some modifications. The interaction term measured by β_4 includes the WMA prices within the trigonometric function, which serves to modulate the frequency of the WMA load as opposed to the amplitude. The parameter a is used to shrink or stretch the frequency of this term to better fit the WMA load. The term measured by β_5 is a hyperbolic interaction term that dampens the amplitude of WMA prices as time increases. The scaling parameters b and c allows us to increase or decrease the dampening effect. As t increases, the dampening effect of c becomes smaller. In this study, we chose to set a to be equal to 1.31, b to be equal to 6.5, and c to be equal to 149.5 when estimating the weekly model coefficients. These values were determined through visual inspection and by iteratively fitting the model to achieve the highest adjusted R-squared values.

Finally, the coefficient $\beta_{6,z,p}$ measures the linear relationship between WMA load and WMA prices, and is analogous to the coefficients featured in the annual and seasonal models.

4.6 Results

4.6.1 Annual model results

Tables 11–14 list the annual model parameter values across all zones and time periods. Unless noted otherwise, all parameters are statistically significant at the 1% significance level. We first discuss the values for α, which represents the adjustment constant of load after controlling for seasonality and price levels in a given zone and time period. The α term is therefore related to the concept of the mean in that it acts as a sort of central moment of the load. Whereas the mean measures the central moment of all observations of a dataset regardless of time of observation, the α term is estimated while explicitly taking into account fluctuations in long-term weather patterns and price levels. Though α is not perfectly synonymous with the mean, the trends in the α term is estimated while explicitly taking into account fluctuations in long term weather patterns and price levels. Though α is not perfectly synonymous with the mean, the trends in the α values largely mirror those seen for the mean values of load in Section 4. Across all time periods, Zone J exhibits the largest α values and Zone H exhibits the lowest. In addition, within each zone, the α values generally tend to be lowest during the morning or the night and highest during the evenings.

Table 11: Annual model parameters across zones. Time period: **Morning**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
α	1836*	1158*	1986*	476*	813*	1379 *	1090*	188*	773*	5644 *	2189*
	(4.567)	(2.03)	(4.385)	(22.741)	(4.418)	(6.499)	(3.04)	(2.519)	(3.033)	(14.75)	(11.51)
β_1	5.47*	0.53*	3.32*	0.77S	-1.45*	5.449*	0.78 *	-1.65 *	4.23*	3.14*	4.50*
	(0.090)	(0.042)	(0.09)	(0.455)	(0.091)	(0.133)	(0.062)	(0.050)	(0.061)	(0.291)	(0.21)
β_2	1.50 *	0.8 *	1.64*	4.98*	1.01*	-0.08*	1.63*	0.49*	-0.16*	3.63*	2.46*
	(0.033)	(0.013)	(0.027)	(0.14)	(0.026)	(0.03)	(0.015)	(0.012)	(0.0152)	(0.067)	(0.049)
β_3	0.0003*	-7.1E-05*	0.0004*	-0.0004*	0.0001*	0.0008*	-0.0004*	-0.0003*	0.0004*	-0.001*	-0.0006*
	(1.2E-05)	(5.8E-06)	(1.2E-05)	(6.6E-05)	(1.1E-05)	(1.5E-05)	(7.4E-06)	(6.1E-06)	(7.2E-06)	(3.4E-05)	(2.2E-05)
β_4	-15.47*	-3.67*	-23.78*	-0.82	12.37*	-12.94*	0.52*	20.50*	-17.89 *	60.36*	25.81*
	(0.586)	(0.263)	(0.574)	(2.937)	(0.577)	(0.859)	(0.402)	(0.333)	(0.401)	(1.955)	(1.489)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.755	0.767	0.839	0.482	0.48	0.579	0.908	0.515	0.630	0.571	0.637

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 12: Annual model parameters across zones. Time period: **Midday**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
α	1863*	1197*	2032*	465*	827*	1424*	1145*	191*	850*	6119*	2220*
	(5.65)	(3.36)	(3.37)	(22.71)	(5.59)	(6.43)	(4.05)	(2.85)	(3.52)	(18.2)	(18.96)
β_1	4.58*	0.61 *	3.23*	-0.31	-1.11*	5.98 *	1.81*	-1.34 *	4.64 *	5.02 *	6.14*
	(0.10)	(0.06)	(0.08)	(0.43)	(0.11)	(0.11)	(0.07)	(0.05)	(0.06)	(0.32)	(0.27)
β_2	2.14*	0.98 *	1.81*	5.25*	1.03*	0.22*	1.74*	0.55*	0.01*	4.69*	3.55*
	(0.04)	(0.02)	(0.02)	(0.13)	(0.03)	(0.03)	(0.02)	(0.01)	(0.01)	(0.07)	(0.06)
β_3	0.0002 *	-6.7E-05*	0.0004*	0.0003*	0.0003*	00007*	-0.0005*	-0.0003*	0.0005*	-0.002*	-0.0008*
	(1.3E-05)	(1.0E-05)	(1.3E-05)	(7.5E-05)	(1.6E-05)	(1.4E-05)	(9.4E-06)	(6.6E-06)	(8.1E-06)	(3.8E-05)	(3.1E-05)
β_4	-7.81*	2.52*	-20.46*	-8.74*	13.11*	-10.91 *	9.16*	22.98*	-21.19*	79.71*	52.09*
	(0.71)	(0.44)	(0.60)	(3.03)	(0.75)	(0.85)	(0.54)	(0.38)	(0.47)	(2.46)	(2.46)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.70	0.553	0.837	0.526	0.439	0.589	0.903	0.488	0.684	0.708	0.649

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 13: Annual model parameters across zones. Time period: **Evening**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
α	1918*	1230 *	2086.*	442*	857*	1526*	1254*	225*	869*	6208*	2519*
	(6.11)	(3.81)	(4.61)	(22.32)	(5.15)	(7.38)	(4.56)	(2.95)	(3.85)	(17.88)	(16.14)
β_1	5.86*	0.93*	3.87 *	0.18	-0.93*	6.63*	2.37 *	-1.35*	4.66 *	4.67*	5.61 *
	(0.11)	(0.07)	(0.08)	(0.43)	(0.09)	(0.13)	(0.08)	(0.05)	(0.06)	(0.31)	(0.235)
β_2	1.32*	0.75*	1.39*	4.76*	0.70*	-0.09*	1.29 *	0.47*	0.02*	3.75 *	2.31*
	(0.04)	(0.02)	(0.02)	(0.12)	(0.03)	(0.03)	(0.02)	(0.01)	(0.01)	(0.06)	(0.05)
β_3	0.0003 *	-2.6E-05*	0.0004*	-0.0003*	0.0001*	0.0009*	-0.0003*	-0.0003*	0.0004*	-0.001*	-0.0006*
	(1.3E-05)	(9.1E-06)	(1.1E-05)	(5.9E-05)	(1.2E-05)	(1.5E-05)	(9.3E-06)	(5.8E-06)	(7.7E-06)	(3.3E-05)	(2.2E-05)
β_4	-10.95*	1.48*	-22.17*	1.53	14.75*	-18.21*	5.04 *	23.24*	-18.34*	67.02*	46.35*
	(0.78)	(0.49)	(0.61)	(2.94)	(0.68)	(0.97)	(0.61)	(0.39)	(0.52)	(2.42)	(2.07)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.62	0.40	0.797	0.50	0.45	0.631	0.85	0.469	0.604	0.651	0.611

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 14: Annual model parameters across zones. Time period: **Night**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
α	1872*	1087*	1989*	554*	760*	1319 *	1087*	174*	722*	5372*	2142*
	(4.81)	(2.34)	(6.63)	(24.64)	(4.79)	(8.55)	(3.79)	(2.52)	(2.95)	(15.76)	(10.59)
β_1	8.28*	1.17*	4.01*	1.02	-1.17*	6.98*	1.19 *	-1.34*	4.53 *	4.32*	4.56*
	(0.10)	(0.51)	(0.14)	(0.51)	(0.10)	(0.18)	(0.08)	(0.05)	(0.06)	(0.34)	(0.19)
β_2	0.60*	0.61 *	1.10 *	6.40*	0.84*	-1.02 *	1.55*	0.32*	-0.34*	3.83*	1.96 *
	(0.04)	(0.01)	(0.05)	(0.19)	(0.035)	(0.05)	(0.02)	(0.01)	(0.02)	(0.08)	(0.05)
β_3	0.001*	0.0002*	0.0009	-6.8E-05	0.0006*	0.0016*	-0.0002*	-0.0002*	0.0006*	-0.001*	-0.0004*
	(1.8E-05)	(9.1E-06)	(2.5E-05)	(9.9E-05)	(1.7E-05)	(2.6E-05)	(1.1E-05)	(7.8E-06)	(9.2E-06)	(4.8E-05)	(2.5E-05)
β_4	-26.55*	-2.64 *	-35.04*	-12.54*	7.35 *	-16.84*	-3.08*	20.13*	-14.82 *	57.65 *	21.57*
	(0.63)	(0.31)	(0.87)	(3.23)	(0.63)	(1.14)	(0.50)	(0.34)	(0.39)	(2.13)	(1.37)
N	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383
R^2	0.735	0.448	0.647	0.48	0.628	0.664	0.803	0.598	0.641	0.39	0.431

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Zones D, G and K are exceptions to this trend: in these zones, α values are highest during the night. More generally, however, we note that within each zone, there is relatively little variation in these α values across time periods. We believe that this is an indication that the remaining four model parameters account for most of the time and price-dependent variations in load and that the estimated α values effectively capture the baseline load levels.

Along with AMA prices, the coefficient β_1 acts to modulate the amplitude of the main waveform function in the annual model. This waveform function measures the long-run cyclicality of the load and has a periodicity of approximately four years. The coefficient β_1 can also be thought of as an amplitude multiplier for the periodic interaction term, adjusting the amplitude up or down depending on how well the amplitude of AMA prices matches that of the AMA load. In addition, the sign of β_1 acts to shift the waveform function left or right to better fit the AMA load.

Across all zones, the β_1 values show no clear time period trends,that is, the relative magnitude of β_1 during the night is not consistently higher than its magnitude during the evenings, and so on. The intra-zonal ordering of β_1 values seem to change randomly from zone to zone. In the vast majority of zones, however, the values for β_1 are the same sign across all time periods: Zones A, B, F, G, I, J, and K show positive values forβ_1 across all time periods and Zones D, E, and H show negative values for β_1 across all time periods. Zone C is the only exception to this trend as it shows a negative β_1 value during the night but positive β_1 values during the morning, midday, and evening.

Zone F shows by far the largest magnitude β_1 values for every time period: it is ap-

proximately one order of magnitude larger than the second largest β_1 values in Zone D. In contrast, Zone E shows by far the smallest magnitude β_1 values in each time period. We find this result particularly interesting given that these two zones are located adjacent to one another but often become isolated from one another due to the Central East constraint as depicted in Figure 8 in Section 2. It is not immediately clear why the β_1 values for these two zones are so dramatically different but it could be simply because prices in the two zones are similar despite large differences in their load profiles. Zone F has a much larger load profile than Zone E and the β_1 values reflect the different amplitudes in the two zones. We also note that the β_1 values in Zones I, J, and K are all very similar across all time periods. This result suggests that the AMA load patterns in the three zones that directly comprise the New York City metro area exhibit approximately the same cyclical dynamics.

The coefficient β_2 measures the linear relationship between load and prices and, as discussed in Section 4, acts as the linear component of the inverse demand elasticity. This is due to the fact that load is extraordinarily inelastic in power markets and causes movements in prices, not the other way around. We would therefore expect all β_2 values to be positive. A large β_2 value indicates that prices are more inflexible given an increase in load and a small β_2 value indicates that prices respond more to an increase in load. In any given zone, the sensitivity of prices in response to movements in load could be the result of many physical grid characteristics. For example, in zones where demand is relatively larger than local supply or areas that can become isolated from supply due to transmission constraints, we would expect to see small β_2 values. In zones where there is an oversupply of power relative to the local demand, we would expect β_2 values to be larger.

Indeed, β_2 values are positive in all but two instances: during the morning and the night time periods in Zone I. This a very strange result as negative value for β_2 implies that increases in demand would actually cause prices to decrease. We note, however, that while this isolated result is strange, it is possible that when combined with the positive morning and night values for β_1 and β_3 in Zone I, the price elasticities for these time periods will turn out to be positive. As for cross-zonal time periods trends, the β_2 values during the midday are largest in eight out of the eleven zones indicating that prices are generally less sensitive to shifts in load during the midday. In contrast, the β_2 values for the evening time period were lowest in six out of the eleven zones and second lowest in four zones. This result is consistent with the fact that load is the highest in most zones during the evening and price spikes happen more frequently as a result.

Across all zones, Zone D exhibits the highest β_2 values for every time period indicating that prices are very insensitive to changes in load. This result makes intuitive sense since Zone D corresponds to a very sparsely populated part of the state of New York that has large hydro and nuclear generators. Perhaps somewhat more surprisingly, β_2 values in Zones J and K, which correspond to the extremely densely populated New York City urban area, are also elevated across all time periods when compared to the remaining zones. Despite

the extreme population density of this area, there are dozens of natural gas and petroleum liquid-fired plants located in these two zones, and the β_2 values are an indication that the supply-to-demand ratio here is higher than expected.

Across all zones, Zone I exhibits by far the smallest β_2 values in magnitude for all four time periods. Zone H also shows consistently small β_2 values that are well below one. Taken together, these β_2 values seem to indicate that the suburbs north of New York City have relatively high demand in comparison to the installed power generation and/or become isolated from cheaper power sources when transmission lines become congested. In the six remaining zones –A, B, C, E, F, and G –β_2 values lie largely between one and two across all time periods.

The coefficient β_3 measures how the relationship between seasonal load and prices changes over time. If β_3 is negative, it is an indication that price sensitivity is increasing: a unit increase in load is causing a progressively larger increase in prices as time progresses. This is because, when combined with β_1 and β_2, a negative β_3 value decreases the combined effect of the coefficients that relate load to prices as time progresses. When we invert the sum of these three coefficients to get an estimate of price elasticity of demand, the denominator $\beta_{1,z,p} \cdot \exp(\sin(\frac{\pi t}{738})) \cdot \frac{1}{AMAPrices_{z,p,t}} + \beta_{2,z,p} + \beta_{3,z,p} \cdot t$ will grow smaller as t increases over the long run, ceteris paribus. A negative β_3 may be the result of load growth, a reduction in generation capacity, or generators offering at higher prices. The opposite is true if β_3 takes on a positive value; a unit increase in load is causing a progressively smaller change in load. Such a result may be due to load shrinkage, an increase in generation capacity, a new transmission line that increases local supply, and so forth.

Across zones, the β_3 values show no apparent time period trends. The ordering of the largest and smallest β_3 value magnitudes seem to change randomly from zone to zone. Within the majority of zones, however, the signs of β_3 values are consistently positive or negative. Zones A, C, E, and I exhibit positive β_3 values across all time periods and Zones F, G, H, J, and K exhibit negative β_3 values across all time periods. Zones B and D show mixed results: the value of β_3 in Zone B is positive during the night and the value of Zone D is positive during the midday but both zones exhibit negative β_3 values in the remaining time periods. Across all zones, Zones F and J show β_3 values with the largest magnitudes indicating that the price sensitivity to load is changing faster in these two zones than the remaining zones after accounting for seasonality and price levels. On the other end of the spectrum, Zones A, B, and C show β_3 values with relatively smaller magnitudes, suggesting that, ceteris paribus, the price sensitivity to changes in load is evolving at a slower pace than the remaining zones.

We now turn to β_4 which measures the "pure" time trend in load after controlling for seasonality and price fluctuations. This parameter is measured in the level-log configuration; a 1% increase in time yields a β_4 change in the AMA load where β_4 is measured in MWh. Positive β_4 values indicate load growth over time and negative values indicate load shrinkage. Though the β_4 values presented in Tables 11–14 may seem high at first glance, it is important

to remember that the β_4 values must be combined with the values of β_1 and β_3 to fully express the time trend of AMA load. In this sense, β_4 is a sort of correction term that captures the linear time trend of AMA load but must be combined with other parameters to fully express the load's time-driven dynamics.

In all but two zones, the signs for β_4 were consistent across all four time periods. Zones A, C, F, and I exhibit negativeβ_4 values across all four periods which indicates load shrinkage over time and Zones E, G, H, J, and K exhibit positive β_4 values across all four periods, which indicates load growth over time. Taken together, the β_4 values indicate load shrinkage in the upstate portion of the state of New York load growth in the downstate portion of the state. These signs are largely consistent with the values of β_3, which also indicate whether price sensitivity was increasing or decreasing over time.

Although Zones B and D both exhibit mixed signs for β_4 across time periods, the evidence seems to be pointing toward load shrinkage in both zones. In Zone B, three of the four β_4 values were negative, only the midday time period β_4 value was positive. In the case of Zone D, the morning and evening β_4 coefficients were not statistically significant but the values for the midday and night time periods were negative, large in magnitude, and highly statistically significant. There are no apparent time period trends in β_4 in the upstate zones; within each upstate zone, the ordering of β_4 values by magnitude for morning, midday, evening, and night vary greatly from zone to zone. In the downstate zones, however, there is a clearer trend where β_4 values during the midday are highest in magnitude, followed by evenings. The β_4 values during the night and morning time periods exhibit the lowest values in all downstate zones except for Zone I. These results are consistent with a larger load growth during the midday and evening hours than in the morning and night time periods.

Finally, we turn our attention to the adjusted R-squared values. We see quite clearly, however, that the annual model seems to work far better for some zones than others. Out of all of eleven zones, Zone G had the highest R-squared values in every time period. R-squared values in Zone G range from 0.80 to 0.91, indicating an excellent model fit across all time periods.

In conclusion, the R-squared values presented in Tables 11–14 suggest that our annual model provides a very good functional fit between the annual moving averages of load and prices in some NYISO zones but a weaker fit in others. There are a multitude of factors that could explain the varying performance of the annual model across different zones and time periods; patterns in consumer behavior, fuel prices, transmission network topology, among others, all vary across zones and can lead to differences in the relationship between load and prices. For LSEs that wish to make more long-term power purchasing arrangements in Zones A, C, F, and G, this annual model would be expected to provide accurate annual price forecasts given an accurate annual load forecast. For the remaining zones, however, the relationship between annual load and prices acts in a way that is not entirely captured by this model. In the next section, we will see that modelling outcomes improve significantly

when we reduce the MA time window to 90 days and control for additional seasonal cycles.

4.6.2 Seasonal model results

Tables 15–18 below present the results of the seasonal model parameter estimations. The first eight coefficients, β_1 to β_8, measure the strength of various medium and long run seasonal cycles in the SMA load. In almost every zone and time period the values for β_1 are positive, the values for β_2 are negative, and both are statistically significant at the 1% significance level, providing strong evidence for the presence of a semiannual seasonal cycle in load. With the exception of β_6, the signs of the remaining seasonal coefficients are more mixed across zones and time periods. The sign of β_6 is almost universally positive across all zones and time periods. These six model coefficients are overwhelmingly statistically significant at the 1% significance level although there are several instances of statistically insignificant coefficients that seem to be scattered somewhat randomly across zones and time periods. Despite these scattered insignificant coefficients, the large number of highly statistically significant coefficients provides evidence of multiple medium-run seasonal cycles in the SMA load.

Table 15: Seasonal model parameters across zones. Time period: **Morning**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.011*	0.016*	0.015*	0.001	0.019*	0.019*	0.028 *	0.040*	0.025*	0.024*	0.031*
	(0.0003)	(0.0003)	(0.0003)	(0.0016)	(0.0006)	(0.0004)	(0.0004)	(0.0008)	(0.0005)	(0.0005)	(0.0005)
β_2	-0.028*	-0.035*	-0.033*	0.010*	-0.052*	-0.045 *	-0.044*	-0.061*	-0.054*	-0.056*	-0.07*
	(0.0003)	(0.0003)	(0.0003)	(0.0016)	(0.0006)	(0.0005)	(0.0005)	(0.0008)	(0.0005)	(0.0005)	(0.0006)
β_3	0.017*	0.002*	0.036*	0.055*	0.051*	-0.016*	-0.027*	0.03*	-0.085*	-0.105*	-0.097*
	(0.0004)	(0.0004)	(0.0004)	(0.0018)	(0.0007)	(0.0006)	(0.0005)	(0.0012)	(0.0007)	(0.0006)	(0.0007)
β_4	0.004*	0.010 *	0.008*	0.010*	0.016 *	0.012*	0.006*	0.020*	-0.014*	-0.016*	-0.013*
	(0.0003)	(0.0003)	(0.0003)	(0.0017)	(0.0006)	(0.0004)	(0.0004)	(0.0008)	(0.0005)	(0.0005)	(0.0005)
β_5	-0.005*	-0.003*	-0.004*	0.027 *	0.007*	-0.006*	0.0008*	0.009*	-0.003 *	-0.005*	-0.0034*
	(0.0003)	(0.0003)	(0.0003)	(0.0016)	(0.0006)	(0.0005)	(0.0004)	(0.0008)	(0.0006)	(0.0005)	(0.0005)
β_6	0.004*	0.004*	0.004*	0.01*	0.005*	0.005*	0.002*	0.024*	0.005*	0.005*	0.009 *
	(0.0003)	(0.0003)	(0.0003)	(0.0017)	(0.0006)	(0.0005)	(0.0005)	(0.0008)	(0.0005)	(0.0005)	(0.0005)
β_7	-0.008*	-0.002*	-0.002*	0.083*	-0.003*	0.002*	-0.0004	0.010*	-0.011 *	0.001*	-0.001*
	(0.0003)	(0.0003)	(0.0003)	(0.0018)	(0.0001)	(0.0005)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_8	0.011*	0.005*	0.006*	0.005*	-0.004*	0.012*	0.009*	-0.003*	0.016*	0.009*	0.013*
	(0.0003)	(0.0003)	(0.0003)	(0.0016)	(0.0006)	(0.0004)	(0.0004)	(0.0008)	(0.0005)	(0.0005)	(0.0005)
β_9	0.044*	0.025*	0.027*	0.026*	0.031*	0.021*	0.024*	0.082*	0.031*	0.031*	0.037*
	(0.0018)	(0.0015)	(0.0013)	(0.0045)	(0.0023)	(0.0017)	(0.0018)	(0.0033)	(0.0019)	(0.0018)	(0.0025)
β_{10}	0.017*	0.035*	0.016*	0.023*	0.023*	0.041*	0.047*	0.055*	0.068*	0.031*	0.046*
	(0.0017)	(0.0015)	(0.0014)	(0.0054)	(0.0025)	(0.0014)	(0.0015)	(0.0028)	(0.0017)	(0.0016)	(0.0017)
β_{11}	0.056*	0.050*	0.065*	0.130*	0.074*	0.063*	0.064*	0.076*	0.035*	0.047*	0.054*
	(0.0012)	(0.0011)	(0.0011)	(0.0045)	(0.0021)	(0.0014)	(0.0014)	(0.0025)	(0.0022)	(0.0015)	(0.0016)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.80	0.822	0.908	0.594	0.852	0.801	0.815	0.774	0.875	0.925	0.907

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 16: Seasonal model parameters across zones. Time period: **Midday**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.016*	0.026*	0.023*	0.003*	0.028 *	0.031*	0.042*	0.056*	0.031*	0.024*	0.041*
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0007)
β_2	-0.037*	-0.051*	-0.045 *	-0.007*	-0.064*	-0.057*	-0.064*	-0.079*	-0.063*	-0.055*	-0.09*
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0005)	(0.0007)
β_3	-0.021*	-0.059*	-0.010*	0.066*	0.009*	-0.074*	-0.114*	-0.074*	-0.139*	-0.123 *	-0.188 *
	(0.0004)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0005)	(0.0007)
β_4	-0.001 *	-0.002*	0.0008*	0.021*	0.009*	-0.0003	-0.014*	-0.0009	-0.024*	-0.018*	-0.028*
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0006)
β_5	-0.0002	-0.002*	-0.002	0.034*	0.007*	-0.0004	0.001*	0.012*	-0.001*	-0.001*	0.001*
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0007)
β_6	0.005*	0.003*	0.005*	0.011*	0.005*	0.004*	0.002*	0.028*	0.006*	0.004*	0.008 *
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0007)
β_7	-0.008*	0.001*	-0.001*	0.076*	0.0002	0.003 *	0.0008	0.014*	-0.01*	0.001*	-0.003 *
	(0.0004)	(0.0005)	(0.0004)	(0.0018)	(0.0007)	(0.0006)	(0.0007)	(0.001)	(0.0007)	(0.0005)	(0.0008)
β_8	0.009*	0.005*	0.006*	0.004*	-0.005*	0.014*	0.014*	0.001	0.019*	0.011**	0.02*
	(0.0003)	(0.0004)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0007)
$\beta 9$	0.102*	0.057*	0.042*	0.174*	-0.02*	0.047 *	0.060*	0.132*	0.069*	0.062*	0.123*
	(0.0061)	(0.0062)	(0.0046)	(0.014)	(0.0076)	(0.0033)	(0.0051)	(0.0071)	(0.0052)	(0.0036)	(0.0067)
β_{10}	0.014*	0.021*	0.004*	0.101*	-0.008*	0.041 *	0.048*	0.072*	0.036*	0.032*	0.040 *
	(0.0018)	(0.0025)	(0.0019)	(0.0067)	(0.0033)	(0.0021)	(0.0025)	(0.0033)	(0.0025)	(0.0018)	(0.0026)
β_{11}	0.094*	0.087*	0.085*	0.221*	0.095*	0.088*	0.098*	0.108*	0.098*	0.070*	0.105 *
	(0.0014)	(0.0017)	(0.0013)	(0.0048)	(0.0023)	(0.0016)	(0.0018)	(0.0025)	(0.0018)	(0.0014)	(0.0019)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.883	0.901	0.88	0.618	0.806	0.896	0.935	0.872	0.947	0.956	0.963

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 17: Seasonal model parameters across zones. Time period: **Evening**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.026*	0.035 *	0.03*	0.004 *	0.034*	0.038*	0.051 *	0.068*	0.04 *	0.028 *	0.051*
	(0.0004)	(0.0005)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0005)	(0.0007)
β_2	-0.04*	-0.053*	-0.046*	-0.011*	-0.062*	-0.057*	-0.064*	-0.078*	-0.065*	-0.056*	-0.086*
	(0.0004)	(0.0005)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0005)	(0.0007)
β_3	-0.009*	-0.036*	0.005*	0.058 *	0.03 *	-0.053*	-0.095*	-0.044*	-0.116*	-0.11*	-0.149*
	(0.0004)	(0.0005)	(0.0004)	(0.0015)	(0.0006)	(0.0006)	(0.0006)	(0.0008)	(0.0006)	(0.0005)	(0.0007)
β_4	0.014*	0.02*	0.019*	0.016 *	0.032*	0.02*	0.004*	0.023 *	-0.005*	-0.008*	-0.001*
	(0.0004)	(0.0005)	(0.0016)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0006)	(0.0005)
β_5	-0.006*	-0.005*	-0.005*	0.028*	0.004 *	-0.007*	-0.001*	0.008*	-0.005*	-0.005*	-0.004*
	(0.0004)	(0.0005)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0006)
β_6	0.005*	0.002*	0.004 *	0.01 *	0.006*	0.002*	0.0007	0.025*	0.005 *	0.003*	0.004 *
	(0.0003)	(0.0005)	(0.0004)	(0.0016)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0007)
β_7	-0.008*	-0.0007	-0.003*	0.074*	-0.003*	0.001	5.1E-05	0.011 *	-0.011*	6.8E-05	-0.003*
	(0.0005)	(0.0005)	(0.0004)	(0.001)	(0.0007)	(0.0006)	(0.0007)	(0.0009)	(0.0007)	(0.0005)	(0.0008)
β_8	0.011*	0.0066*	0.007*	0.008*	-0.003*	0.013 *	0.014*	0.0008	0.0186*	0.01*	0.017*
	(0.0004)	(0.0005)	(0.0003)	(0.0015)	(0.0006)	(0.0005)	(0.0006)	(0.0008)	(0.0006)	(0.0004)	(0.0006)
β_9	0.081*	0.064*	0.054 *	0.136*	0.039*	0.051*	0.064*	0.136 *	0.086*	0.062 *	0.107 *
	(0.0035)	(0.0033)	(0.0024)	(0.0077)	(0.0036)	(0.0031)	(0.0046)	(0.0066)	(0.0049)	(0.0036)	(0.0064)
β_{10}	0.004*	0.023*	0.014 *	0.081 *	0.019*	0.045*	0.052 *	0.068*	0.036*	0.033*	0.027 *
	(0.0024)	(0.0025)	(0.0019)	(0.0065)	(0.0028)	(0.0021)	(0.0025)	(0.0035)	(0.0025)	(0.0019)	(0.0026)
β_{11}	0.085*	0.085*	0.084*	0.21*	0.085 *	0.084*	0.089 *	0.095*	0.094 *	0.067 *	0.093*
	(0.0015)	(0.0017)	(0.0012)	(0.0046)	(0.0019)	(0.0017)	(0.0019)	(0.0026)	(0.0019)	(0.0015)	(0.0021)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.841	0.866	0.895	0.633	0.864	0.872	0.914	0.847	0.928	0.943	0.947

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Across all zones and time periods, the three coefficients that measure the relationship between the SMA load and prices, β_9, β_{10}, and β_{11}, are overwhelmingly positive and statistically significant at the 1% significance level. As with the β_1 term used in the annual model, β_9 and β_{10} measure the long-run cyclicality of the SMA load while using SMA prices for the purposes of amplitude modulation. The fact that these coefficients are positive indicates that the vertical fluctuations of SMA load and SMA prices are synchronous. As with β_2 used in the annual model, β_{11} measures the linear relationship between SMA load and prices after controlling for seasonal fluctuations in the two time series. As we discussed previously, a positive β_{11} value indicates that, over average, an increase in the demand for power yields an increase in prices. This is consistent with the well-documented fact that the demand for power is thoroughly disconnected from real-time power prices and will cause prices to move

Table 18: Seasonal model parameters across zones. Time period: **Night**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.015*	0.022*	0.021*	0.002*	0.027*	0.025*	0.034*	0.051*	0.034*	0.032*	0.038*
	(0.0003)	(0.0004)	(0.0003)	(0.0017)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_2	-0.03*	-0.042 *	-0.038*	-0.011*	-0.061*	-0.054*	-0.054*	-0.073*	-0.067*	-0.066*	-0.079*
	(0.0004)	(0.0004)	(0.0003)	(0.0017)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_3	0.014*	-0.002*	0.036*	0.063*	0.057*	-0.023*	-0.039*	0.02*	-0.099*	-0.119*	-0.107*
	(0.0004)	(0.0004)	(0.0004)	(0.0018)	(0.0007)	(0.0007)	(0.0006)	(0.0011)	(0.0007)	(0.0006)	(0.0007)
β_4	0.004*	0.01*	0.01*	0.017*	0.022*	0.015*	0.005*	0.021*	-0.016*	-0.017*	-0.014*
	(0.0003)	(0.0004)	(0.0003)	(0.0017)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_5	-0.004*	-0.0006	-0.004*	0.037*	0.011*	-0.009*	-0.0007*	0.013*	-0.003 *	-0.004 *	-0.003 *
	(0.0004)	(0.0004)	(0.0003)	(0.0017)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_6	0.0002*	0.003*	0.004*	0.007*	0.005*	0.006*	0.003*	0.027*	0.005 *	0.006 *	0.009*
	(0.0004)	(0.0004)	(0.0003)	(0.0018)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_7	-0.009*	-0.0007	-0.002*	0.088*	0.0009	0.001*	5.0E-05*	0.008*	-0.012 *	0.0005	-0.0015*
	(0.0004)	(0.0004)	(0.0004)	(0.0019)	(0.0008)	(0.0006)	(0.0006)	(0.0011)	(0.0007)	(0.0006)	(0.0007)
β_8	0.014*	0.006*	0.007*	0.002*	-0.003*	0.016*	0.01*	-0.001	0.018*	0.011 *	0.014*
	(0.0003)	(0.0004)	(0.0003)	(0.0017)	(0.0007)	(0.0006)	(0.0005)	(0.0009)	(0.0006)	(0.0005)	(0.0006)
β_9	0.057 *	0.038*	0.033*	0.006*	0.034*	0.025*	0.027*	0.101*	0.047*	0.048 *	0.057 *
	(0.0023)	(0.0024)	(0.0021)	(0.0067)	(0.0038)	(0.0023)	(0.0024)	(0.0041)	(0.0028)	(0.0026)	(0.0035)
β_{10}	0.012*	0.039*	0.018*	0.015	0.027*	0.050 *	0.063*	0.076*	0.043*	0.044*	0.051*
	(0.0019)	(0.0019)	(0.0017)	(0.0074)	(0.0032)	(0.0018)	(0.0018)	(0.0032)	(0.0022)	(0.0020)	(0.0021)
β_{11}	0.072 *	0.068*	0.076*	0.014*	0.102 *	0.069*	0.088*	0.076*	0.076*	0.064 *	0.072 *
	(0.0013)	0.0013)	(0.0012)	(0.0049)	(0.0069)	(0.0017)	(0.0016)	(0.0027)	(0.0019)	(0.0016)	(0.0018)
N	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383
R^2	0.806	0.832	0.902	0.588	0.852	0.793	0.841	0.783	0.896	0.930	0.921

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

rather than the other way around.

The R-squared values presented in Tables 15–18 indicate that the seasonal model generally offers an improved fit in comparison to the annual model. The seasonal model seems to perform particularly well in the downstate zones. Across all models, Zone J exhibits the highest R-squared values during the morning and night time periods and Zone K exhibits the highest values during the midday and evening time periods. In addition, the R-squared values in Zones J and K are very large for all four time periods; both zones exhibited R-squared values over 0.90 in all time periods. Zones G and I also perform quite well with R-squared values that exceed 0.90 for the midday and evening time periods and 0.80 for the morning and night time periods. Of the downstate zones, the seasonal model performed the weakest in Zone H, but even here, R-squared values during the midday and evening time periods are large, hovering around 0.85. Although the seasonal model fit was generally lower in the upstate zones, it still offers improvements over the annual model: Zones A, B, C, and E exhibit R-squared values above 0.80 across all four time periods, and all four model R-squared values for Zone C are closely clustered around 0.90. Zone D was a notable exception as it exhibits the lowest R-squared values of any zone across all four time periods. In addition to the fact that Zone D also exhibited some unusual parameter values when compared to general trends in other zones, the presence of low R-squared values indicates that our seasonal model does not capture the dynamics of the SMA load and prices in Zone D as well as it does in the other ten NYISO zones.

4.6.3 Weekly model results

Tables 19–22 show the weekly model parameter values across all zones and time periods. The coefficients β_1, β_2, and β_3 measure the cyclical, time-dependent components of the weekly model. Across all forty-four weekly models, β_1 values are positive and clustered around 0.45.

In all models except for Zone A during the mornings, β_2 is positive. In addition, the values for β_2 are clustered tightly around 0.10 in the downstate zones in all time periods but exhibit far more variability in the upstate zones. Together, these coefficient values provide evidence of the presence of medium-term seasonal trends in the WMA data.

Table 19: Weekly model parameters across zones. Time period: **Morning**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.029*	0.036 *	0.035*	0.024*	0.059*	0.049*	0.038*	0.047*	0.057*	0.059*	0.073*
	(0.0007)	(0.0008)	(0.0008)	(0.0011)	(0.0016)	(0.0009)	(0.0011)	(0.0016)	(0.0014)	(0.0013)	(0.0014)
β_2	-0.046*	0.064 *	0.059*	0.029 *	0.086*	0.079*	0.09*	0.118*	0.099*	0.097*	0.126*
	(0.0007)	(0.0008)	(0.0008)	(0.0012)	(0.0016)	(0.0015)	(0.0011)	(0.0017)	(0.0015)	(0.0013)	(0.0014)
β_3	0.014*	-0.0004	0.026*	0.057 *	0.036*	-0.015*	-0.023*	0.012*	-0.059*	-0.076*	-0.069*
	(0.0007)	(0.0008)	(0.0008)	(0.0012)	(0.0016)	(0.0016)	(0.0011)	(0.0017)	(0.0015)	(0.0013)	(0.0015)
β_4	0.025*	0.043*	0.042*	0.019*	0.051*	0.035*	0.039*	0.046*	0.039*	0.032*	0.025*
	(0.0013)	(0.0014)	(0.0013)	(0.0016)	(0.0017)	(0.0016)	(0.0018)	(0.0028)	(0.0025)	(0.0026)	(0.0022)
β_5	13.257*	2.058	0.479	5.391*	13.201*	3.317	3.072	32.14*	-6.677*	-0.115	12.507 *
	(1.8318)	(2.0912)	(1.9309)	(2.598)	(2.4848)	(2.1492)	(2.5412)	(3.9183)	(3.4059)	(2.6412)	(3.7895)
β_6	0.009*	0.019 *	0.028 *	0.026*	0.023*	0.033*	0.032*	0.036*	0.021*	0.016*	0.014 *
	(0.0014)	(0.0016)	(0.0014)	(0.0014)	(0.0018)	(0.0016)	(0.0018)	(0.0029)	(0.0025)	(0.0024)	(0.0025)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.637	0.711	0.751	0.557	0.781	0.742	0.724	0.664	0.663	0.738	0.751

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 20: Weekly model parameters across zones. Time period: **Midday**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.039*	0.047 *	0.042 *	0.026*	0.065*	0.053*	0.049 *	0.057*	0.062*	0.057*	0.089*
	(0.0008)	(0.0011)	(0.0008)	(0.0011)	(0.0011)	(0.0011)	(0.0014)	(0.0018)	(0.0016)	(0.0012)	(0.0019)
β_2	0.069*	0.089*	0.081*	0.039 *	0.112*	0.113*	0.121*	0.157*	0.106*	0.087*	0.156*
	(0.0008)	(0.0011)	(0.0009)	(0.0012)	(0.0011)	(0.0012)	(0.0016)	(0.0019)	(0.0018)	(0.0013)	(0.0021)
β_3	-0.018*	-0.042*	-0.007*	0.045*	0.007*	-0.055*	-0.082*	-0.055*	-0.112*	-0.091*	-0.138*
	(0.0008)	(0.0011)	(0.0008)	(0.0011)	(0.0011)	(0.0011)	(0.0014)	(0.0018)	(0.0016)	(0.0012)	(0.0019)
β_4	0.018*	0.037*	0.031 *	0.026 *	0.032 *	0.045 *	0.053*	0.052*	0.051*	0.045*	0.031*
	(0.0014)	(0.0021)	(0.0016)	(0.0019)	(0.0021)	(0.0019)	(0.0024)	(0.0029)	(0.0027)	(0.0021)	(0.0031)
β_5	11.361*	0.392	2.881	-7.674*	20.95*	-2.391	9.093	9.506*	17.722*	15.251*	5.098
	(2.4891)	(3.1672)	(2.7811)	(3.4813)	(3.3581)	(3.5993)	(4.0758)	(4.7821)	(4.3532)	(3.5922)	(4.881)
β_6	0.041*	0.066 *	0.058 *	0.033 *	0.051*	0.044*	0.057	0.067*	0.052*	0.033*	0.073 *
	(0.0021)	(0.0031)	(0.0024)	(0.0026)	(0.0029)	(0.0022)	(0.0014)	(0.0038)	(0.0035)	(0.0025)	(0.0038)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.763	0.804	0.814	0.544	0.831	0.811	0.808	0.778	0.766	0.808	0.812

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 21: Weekly model parameters across zones. Time period: **Evening**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.033*	0.042 *	0.037 *	0.023*	0.057*	0.047*	0.041*	0.043*	0.053*	0.054*	0.074*
	(0.0009)	(0.0011)	(0.0009)	(0.0011)	(0.0011)	(0.0011)	(0.0014)	(0.0018)	(0.0015)	(0.0012)	(0.0017)
β_2	0.082*	0.103 *	0.091*	0.037*	0.113*	0.114 *	0.135*	0.168*	0.122*	0.093*	0.167 *
	(0.0009)	(0.0011)	(0.0011)	(0.0012)	(0.0012)	(0.0012)	(0.0016)	(0.0021)	(0.0017)	(0.0013)	(0.0019)
β_3	-0.009*	-0.029*	0.002*	0.048*	0.022*	-0.043*	-0.072*	-0.039*	-0.087*	-0.082*	-0.113*
	(0.0009)	(0.0011)	(0.0009)	(0.0011)	(0.0011)	(0.0011)	(0.0014)	(0.0017)	(0.0015)	(0.0012)	(0.0017)
β_4	0.036*	0.055*	0.047*	0.026*	0.053 *	0.055*	0.046*	0.047*	0.043*	0.035*	0.03*
	(0.0016)	(0.0021)	(0.0017)	(0.0018)	(0.0021)	(0.0019)	(0.0024)	(0.0029)	(0.0025)	(0.0019)	(0.0028)
β_5	5.881*	-5.322*	-2.797	-8.109*	10.291*	-4.049	2.703	8.806*	4.628	8.921*	-7.386 *
	(2.2161)	(2.7151)	(2.3611)	(2.6252)	(2.7653)	(2.9421)	(3.5513)	(4.2495)	(3.7462)	(3.0651)	(3.3883)
β_6	0.024*	0.041*	0.042*	0.048*	0.032*	0.033*	0.052*	0.065*	0.052*	0.037*	0.053*
	(0.0022)	(0.0027)	(0.0022)	(0.0022)	(0.0026)	(0.0022)	(0.0031)	(0.0037)	(0.0033)	(0.0024)	(0.0035)
N	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384	4384
R^2	0.766	0.798	0.806	0.595	0.818	0.813	0.808	0.781	0.771	0.804	0.822

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

The values of β_3 show far more variation across zones and time periods; they are positive and negative, large and small with no concrete systemic patterns. Across all time periods, however, the β_3 values for Zones I, J, and K are all negative and clustered around -0.10 suggesting that there are some longer term seasonal trends prevalent in the WMA load in these three downstate zones and that these three zones share some similar weekly dynamics.

Table 22: Weekly model parameters across zones. Time period: **Night**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	0.029*	0.039*	0.035*	0.022*	0.063*	0.056*	0.042*	0.051*	0.064*	0.054*	0.077*
	(0.0007)	(0.0008)	(0.0008)	(0.0012)	(0.0012)	(0.0011)	(0.0011)	(0.0018)	(0.0016)	(0.0012)	(0.0015)
β_2	0.052*	0.072 *	0.067*	0.031*	0.101*	0.093 *	0.104*	0.138 *	0.117 *	0.096*	0.137*
	(0.0007)	(0.0009)	(0.0008)	(0.0012)	(0.0012)	(0.0011)	(0.0012)	(0.0019)	(0.0017)	(0.0013)	(0.0016)
β_3	0.012*	-0.0009	0.027 *	0.057*	0.042*	-0.021*	-0.031*	0.004*	-0.071*	-0.082*	-0.077*
	(0.0007)	(0.0008)	(0.0008)	(0.0012)	(0.0012)	(0.0011)	(0.0011)	(0.0018)	(0.0016)	(0.0012)	(0.0015)
β_4	0.042*	0.051*	0.045*	0.026 *	0.062 *	0.041*	0.047*	0.056*	0.055*	0.039*	0.038*
	(0.0014)	(0.0016)	(0.0016)	(0.0019)	(0.0022)	(0.0018)	(0.0021)	(0.0032)	(0.0029)	(0.0021)	(0.0026)
β_5	4.114*	-2.281	-2.306	-4.394	16.534*	9.104*	6.504	53.575 *	-2.828	8.495 *	12.285 *
	(1.8651)	(2.2179)	(2.3971)	(3.1452)	(3.3975)	(2.7494)	(3.2374)	(5.0794)	(4.4868)	(3.3434)	(5.0512)
β_6	0.011*	0.025*	0.038*	0.034 *	0.034*	0.043*	0.041 *	0.041*	0.023*	0.028*	0.028 *
	(0.0016)	(0.0019)	(0.0019)	(0.0019)	(0.0026)	(0.0019)	(0.0023)	(0.0035)	(0.0032)	(0.0026)	(0.0034)
N	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383	4383
R^2	0.731	0.778	0.791	0.558	0.803	0.779	0.774	0.721	0.717	0.788	0.787

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Although it defies straightforward interpretation, β_4 measures what is essentially a frequency modulation term for the WMA load that uses a combination of the WMA prices and a cyclical term. A larger β_4 value indicates that the frequency of the WMA price signal has a larger impact on the WMA load frequency. Across all zones and time periods, the values for β_4 are clustered closely around 0.35. With the exception of Zone G during the midday time period, all β_4 values are positive. Zones A and D generally exhibit the smallest β_4 values and Zones E, H, and I exhibit the largest. Although the time period trends are a bit unstable, the β_4 parameter during the night time tend to be the largest within a given zone indicating that the WMA load and WMA price frequencies are somewhat more aligned during the night time than other periods of the day.

The coefficient, β_5, measures an interaction term between time and WMA prices. However, because the time variable is located in the denominator, the passage of time will have a dampening effect on the relationship between WMA load and prices. This effect will be captured whether β_5 is positive or negative as a negative β_5 value simply acts to reflect the interaction term over the x-axis. The model results for β_5 are quite mixed and unlike most of the other model coefficients we have seen; a sizable portion of the β_5 values is not statistically significant at the 5% significance level. Zones C and G contain no statistically significant coefficients and Zones B and F contain only one statistically significant value, each indicating that WMA load in these zones exhibit little or no time-dependent dampening. In contrast, this time-dampening effect appears to be most pronounced in Zones E and H since both zones show large and statistically significant β_5 values across all time periods.

As expected, the vast majority of values for β_6 are positive. Across all zones, the β_6 values tend to be smallest during the mornings and largest during midday or evenings; seven out of the eleven zones exhibited their largest β_6 values during the midday and three exhibited them during the evenings. Therefore, it appears that, after controlling for cyclical and linear time interactions, WMA prices are more sensitive to changes in WMA load during the mornings and less sensitive to changes in WMA load during the midday and evening time periods. This result makes intuitive sense since load increases quickly during the morning and generators must somewhat rapidly ramp up their production as a consequence, often

times causing large jumps in prices.

Across all time periods, Zones F, G, H, and K exhibit the highest β_6 values and Zones A and J exhibit the lowest. This result suggests that the WMA prices are generally less sensitive to WMA load in Zones F, G, H and K and more sensitive to changes in WMA load in Zones A and J. As a final observation, we note that the values of β_6 are lower than the corresponding coefficient in the seasonal model, β_{11}. We believe that the explanation for this is straightforward: due to the nature of spot prices for power, a single price spike caused by a sudden increase in load will affect the weekly MA value of prices more than it will in a 90-day MA.

Lastly, we turn our attention to the R-squared values. We first note that within almost every zone, R-squared values are lowest during the mornings, higher during the night, and highest during the midday or evenings. Only Zone D deviates from this consistent pattern. What is most apparent, however, is that across all zones and time periods the R-squared value ranges are very tightly clustered between 0.75 and 0.80; only Zone D, which was the worst performing zone, falls far outside of this range with R-squared values clustered tightly around 0.55. This tight clustering of R-squared values for the weekly model stands in contrast to the annual and seasonal models where R-squared values varied far more across zones and time periods. This clustering also seems to suggest that the WMA load and prices, which are data series with significantly less "time momentum" than the corresponding AMA or SMA data series, have a similar relationship from zone to zone and that the weekly model captures these dynamics in largely the same way. Although there are large differences in β_5 across zones and time periods, this parameter's value and statistical significance changes quite dramatically between zones and time periods. Comparing the R-squared values across all three models, we see that the weekly model largely outperforms the annual model, which shows R-squared values that range widely between zones and time periods and are generally much lower. Despite a slightly higher variation in R-squared values, however, the seasonal model generally exhibits higher R-squared values than the weekly model. The seasonal model yielded many R-squared values over 0.80 and several even larger than 0.90.

4.7 Seasonal model residual analysis

In this section, we present the methodology and results of our analysis of the seasonal model residuals. In the parlance of signal processing, these residuals constitute the "noise" component of the seasonal MA model, that is, the left over component of the signal whose behavior is not explained by the quasi-periodic waveform function. As opposed to the annual and weekly models, we chose to analyze seasonal model residuals for two primary reasons. First, of the three models that we developed, the seasonal model yielded the highest average R-squared values. From this result, we surmise that the seasonal residuals will be the least likely to contain omitted systemic information. Secondly, power futures with monthly and seasonal durations appear to be among the most commonly traded hedging instruments.

The seasonal model is therefore likely to be the most applicable model for making real-world decisions about risk mitigation.

4.7.1 The Box–Jenkins methodology

We use the time series analysis methodology developed by Box and Jenkins (1970) to construct regressions of the seasonal model residuals. The core idea of this methodology is that the past values of a given time series variable may contain valuable information about predicting its future values and comprises of three main steps. The first is to ensure that the data in question are stationary, that is, exhibit a constant mean and variance over time. This ensures that the data are deseasonalized and do not exhibit means that are time dependent. The second step is to determine whether past values of the stationary data contain important information about predicting future values and, if they do, choose the optimal number of previous values, also called "lags," of the variable. Finally, after choosing a model specification, the model parameters are estimated and the model residuals are inspected to see if they are white noise, that is, exhibit a zero mean and constant variance.

The Box–Jenkins methodology frequently gives rise to ARMA models, where autoregressive and MA denote two distinct statistical processes. In time series data that follows an autoregressive process, the current value of a given variable can be expressed as a function of its previous value or sum of previous values. In other words, the current value of a given variable is correlated with previous values from the same time series. When this is the case, the current value of this quantity can be computed as a linear function of the previous values, as follows:

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \beta_2 Y_{t-2} + ... + \beta_p Y_{t-p} + \epsilon_t$$

where Y_t is a time series variable, β_1 $...\beta_p$ are coefficients to be estimated, p is the final lag length and ϵ_t is an error term. We will henceforth refer to the lagged variables in an autoregressive process, Y_{t-s}, as AR terms.

In contrast to an autoregressive process, an MA process is one where a given variable can be described as the sum of its previous error terms,ϵ_t. It is formulated as follows:

$$Y_t = \mu + \epsilon_t + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + ... + \theta_q \epsilon_{t-q}$$

where Y_t is a time series variable, μ is the time invariant mean of Y_t, θ_1 ... θ_q are coefficients to be estimated, q is the final lag length, and ϵ_t is an error term. We will henceforth refer to the lagged variables in an MA process, ϵ_{t-q}, as MA terms.

In order to determine the type of model and optimal number of lagged AR and/or MA terms, the Box–Jenkins methodology employs the use of autocorrelation and partial autocorrelation functions (ACF and PACF, respectively). An ACF measures the correlation

between a variable at time t and at all preceding lags, *t-1*, *t-2*, and so on. The PACF measures the correlation between a variable at time t and all preceding lags with the caveat that it controls for all intermediate lags. For instance, a PACF measures the correlation between Y_t and Y_{t-3} while controlling for Y_{t-1} and Y_{t-2}. In this sense, the PACF measures the direct partial relationship between Y_t and any other lag value, Y_{t-p}, in a way that the ACF does not. In the Box–Jenkins methodology, a steadily decaying ACF and PACF with several sharp peaks or rapid decay is indicative of an autoregressive process. A steadily decaying PACF and ACF with several sharp peaks or rapid decay is indicative of a MA process. When choosing lag lengths, it is customary to choose the lags that show statistically significant correlations on the PACF or ACF correlograms depending on whether the process is autoregressive or an MA.

4.7.2 Model selection

We first ran the Dickey–Fuller tests on all forty-four sets of seasonal model residuals to test for stationarity. For all sets of residuals, the Dickey–Fuller test statistic strongly rejected the null hypothesis of a unit root and therefore confirmed that the residuals are stationary across all zones and time periods. In addition, the seasonal model is a waveform function that already captures the seasonality of the SMA load and prices. We have therefore effectively already taken into account seasonality in our data.

To choose the appropriate model ARMA, we generated the ACF and PACF functions for all forty-four sets of seasonal model residuals. In all sets of residuals, the ACFs showed identical patterns; linearly decaying correlation of lagged variables as the lag length increased. Across all sets of residuals, the PACF functions exhibited extremely similar general patterns: large, statistically significant spikes during the first few lags and a quickly decaying dampened sinusoid of lag values thereafter. Taken together, the ACF and PACF correlograms of the seasonal model residuals indicate that they follow an autoregressive process.

Figures 15–18 show two sample ACF functions and two sample PACF functions, that is, those of seasonal model residuals of Zones G and J during the morning time period.

Figure 15: ACF of Zone G

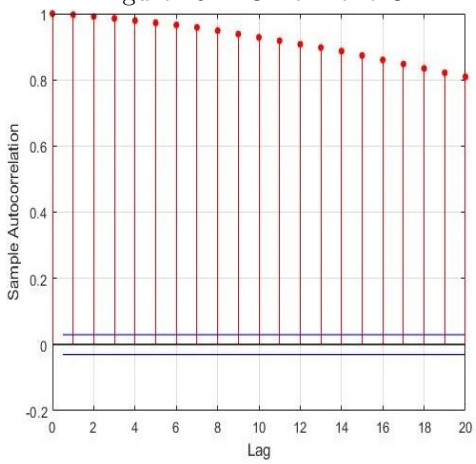

Figure 16: ACF of Zone J

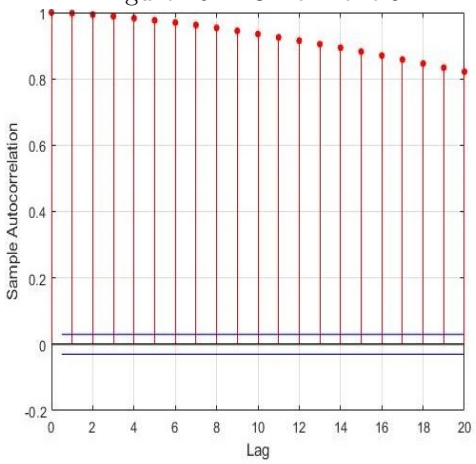

Figure 17: PACF of Zone G

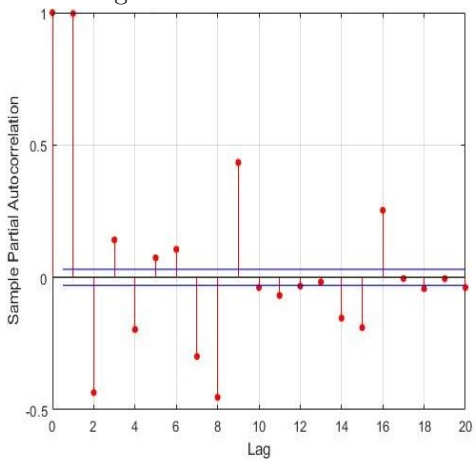

Figure 18: PACF of Zone J

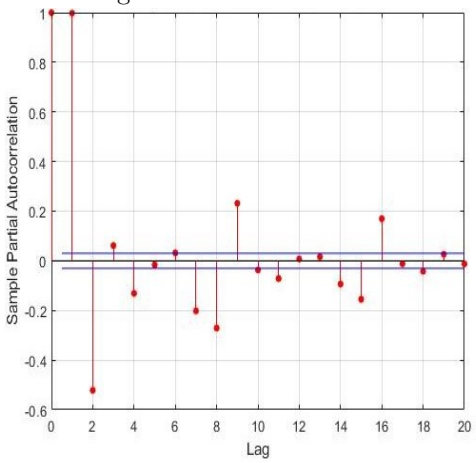

We see from Figures 15 to 18 that the ACF functions are identical but that the PACF functions differ in the number of statistically significant lag terms. There also appears to be a weekly component that is present, that is, the lag at seven, fourteen, and so on are correlated with the current residual value. This is to be expected as the relationship between load and prices on Monday mornings will correlate with past values of the relationship on Monday mornings due to the behavioral patterns of power end-users.

For each zone across all time periods, we fit an AR regression using ordinary least squares to each of the forty-four sets of seasonal model residuals by choosing all statistically significant lags until reaching an insignificant lag. For example, based on the two PACF correlograms above, we fit an AR(9) model to Zone J and an AR(4) model to Zone G during the morning time periods. Although this is a simple model-selection process and excludes several statistically significant lags, we shall see that the adjusted R-squared values following our simple rule indicate a highly sufficient model fit. Because the residuals already have a mean that is approximately equal to zero, we fit these regressions without a linear intercept.

4.7.3 Results

Tables 23–26 below show the results of the AR regressions that run on the seasonal model residuals. The parameter values of these regressions largely match the dampened sinusoid pattern seen in the PACF functions. Twenty-six of these AR regressions have two lags, fifteen have four lags, and three have nine lags. We also see that the regressions for the morning and midday time periods tend to have more lags than those for the evenings and night time periods. It is not exactly clear why the lag lengths tend to be larger during the morning and midday time periods. It could be because SMA load is less volatile during the mornings and midday than it is during the evenings and nights. There would therefore be predictive power in a longer time span of past SMA load values during the morning and midday time periods than in the evenings and nights.

Table 23: AR(2–9) parameters across zones. Time period: **Morning**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	1.35*	1.51 *	1.42 *	1.13*	1.45*	1.52*	1.55	*1.57*	1.55*	1.62 *	1.56*
	(0.015)	(0.014)	(0.015)	(0.014)	(0.013)	(0.014)	(0.014)	(0.012)	(0.014)	(0.013)	(0.012)
β_2	-0.52 *	-0.61 *	-0.65*	0.002*	-0.45*	-0.69 *	-0.69*	-0.58 *	-0.75*	-0.71*	-0.56*
	(0.025)	(0.022)	(0.022)	(0.025)	(0.013)	(0.027)	(0.027)	(0.012)	(0.027)	(0.023)	(0.012)
β_3	0.29*	0.12 *	0.42*	-0.04*	-	0.34*	0.26*	-	0.37*	0.14*	-
	(0.025)	(0.023)	(0.022)	(0.025)		(0.027)	(0.027)		(0.027)	(0.023)	
β_4	-0.12*	-0.02*	-0.19*	-0.09*	-	-0.16*	-0.13*	-	-0.18*	-0.07*	-
	(0.014)	(0.023)	(0.015)	(0.015)		(0.014)	(0.015)		(0.015)	(0.023)	
β_5	-	-0.12*	-	-	-	-	-	-	-	-0.08*	-
		(0.023)								(0.023)	
β_6	-	0.11*	-	-	-	-	-	-	-	0.12*	-
		(0.023)								(0.023)	
β_7	-	0.56 *	-	-	-	-	-	-	-	- 0.64*	-
		(0.023)								(0.023)	
β_8	-	-0.92 *	-	-	-	-	-	-	-	-1.07 *	-
		(0.022)								(0.022)	
β_9	-	0.35*	-	-	-	-	-	-	-	0.43*	-
		(0.014)								(0.013)	
N	4380	4375	4380	4380	4382	4380	4380	4382	4380	4375	4382
R^2	0.995	0.996	0.995	0.999	0.998	0.997	0.997	0.998	0.997	0.998	0.997

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 24: AR(2–9) parameters across zones. Time period: **Midday**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	1.279*	1.49 *	1.46*	1.32*	1.52*	1.59*	1.59*	1.54*	1.54 *	1.56*	1.53*
	(0.014)	(0.015)	(0.014)	(0.014)	(0.012)	(0.012)	(0.012)	(0.015)	(0.012)	(0.014)	(0.015)
β_2	-0.28 *	-0.62 *	-0.62 *	-0.32*	-0.52 *	-0.59 *	-0.59 *	-0.51 *	-0.54 *	-0.66*	-0.52 *
	(0.014)	(0.027)	(0.026)	(0.014)	(0.013)	(0.012)	(0.012)	(0.027)	(0.013)	(0.024)	(0.027)
β_3	-	0.22 *	0.32*	-	-	-	-	0.28*	-	0.17*	0.06*
		(0.026)	(0.026)					(0.027)		(0.025)	(0.027)
β_4	-	-0.09*	-0.16*	-	-	-	-	-0.06*	-	-0.11*	-0.08*
		(0.015)	(0.015)					(0.015)		(0.025)	(0.015)
β_5	-	-	-	-	-	-	-	-	-	-0.09*	-
										(0.025)	
β_6	-	-	-	-	-	-	-	-	-	0.12*	-
										(0.025)	
β_7	-	-	-	-	-	-	-	-	-	-0.47*	-
										(0.025)	
β_8	-		-	-	-	-	-	-	-	-0.81*	-
										(0.024)	
β_9	-	-	-	-	-	-	-	-	-	-0.31 *	-
										(0.014)	
N	4382	4380	4382	4380	4382	4382	4382	4380	4382	4380	4380
R^2	0.994	0.997	0.999	0.996	0.998	0.998	0.998	0.998	0.997	0.997	0.998

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

Table 25: AR(2–4) parameters across zones. Time period: **Evening**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	1.42*	1.52 *	1.48*	1.31*	1.46*	1.61*	1.61*	1.62*	1.61*	1.63*	1.62*
	(0.014)	(0.013)	(0.013)	(0.014)	(0.015)	(0.012)	(0.011)	(0.011)	(0.012)	(0.015)	(0.011)
β_2	-0.42 *	-0.52 *	-0.49 *	-0.30 *	-0.42 *	-0.61*	-0.61*	-0.62 *	-0.61*	-0.83*	-0.62*
	(0.013)	(0.012)	(0.013)	(0.014)	(0.026)	(0.012)	(0.011)	(0.012)	(0.012)	(0.028)	(0.012)
β_3	-	-	-	-	0.07*	-	-	-	-	0.32*	-
					(0.026)					(0.028)	
β_4	-	-	-	-	-0.11*	-	-	-	-	-0.12*	-
					(0.015)					(0.014)	
N	4382	4382	4382	4382	4382	4380	4382	4382	4382	4382	4382
R^2	0.997	0.997	0.999	0.997	0.998	0.998	0.998	0.998	0.998	0.997	0.998

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

The values of the first four coefficients in the AR regressions are remarkably similar across all zones and time periods. We believe that this high degree of similarity in AR regression results suggests that parameters of the seasonal waveform function have captured most of the variation in behavior between zones. In the forty-four AR regressions, all β_1 values are positive and almost universally clustered around 1.50. Similarly, all but one of the values for β_2 are negative and almost universally clustered around -0.50. All but one of the values for

Table 26: AR(2–4) parameters across zones. Time period: **Night**

	Zone A (West)	**Zone B** (Genesee)	**Zone C** (Central)	**Zone D** (North)	**Zone E** (Moh.Valley)	**Zone F** (Capital)	**Zone G** (Hud.Valley)	**Zone H** (Millwood)	**Zone I** (Dunwoodie)	**Zone J** (NYC)	**Zone K** (L.Island)
β_1	1.59*	1.60 *	1.61*	1.32*	1.66*	1.71*	1.71*	1.69*	1.84*	1.71	1.76*
	(0.012)	(0.012)	(0.012)	(0.014)	(0.015)	(0.011)	(0.011)	(0.011)	(0.015)	(0.015)	(0.011)
β_2	-0.60 *	-0.60 *	-0.61 *	-0.32*	-0.71*	-0.71oneS	-0.71*	-0.69 *	-1.01*	-0.93*	-0.77*
	(0.012)	(0.012)	(0.012)	(0.014)	(0.029)	(0.011)	(0.011)	(0.011)	(0.031)	(0.029)	(0.011)
β_3	-	-	-	-	0.16*	-	-	-	0.21*	0.33*	-
					(0.029)				(0.031)	(0.029)	
β_4	-	-	-	-	-0.11*	-	-	-	-0.05*	-0.12*	-
					(0.015)				(0.015)	(0.015)	
N	4381	4381	4381	4381	4389	4381	4381	4381	4379	4379	4381
R^2	0.998	0.998	0.998	0.999	0.999	0.999	0.998	0.999	0.999	0.998	0.999

Note: Standard errors in parentheses. An asterisk indicates significance at the 1% level.

β_3 are positive and clustered loosely around 0.30 and all β_4 values are negative and clustered around -0.10. The AR regression for Zone D during the morning time period is the exception to the otherwise universal trends for β_2 and β_3 as it exhibits positive and negative values for β_2 and β_3, respectively, and neither are statistically significant at the 10% level. After β_4, the steady alternating pattern between positive and negative coefficients appears to break down. In the AR(9) regressions for Zones B and J during the morning and Zone J during the afternoon, the coefficients for lags five through nine seem to switch randomly between positive and negative.

Finally, we note that in all regressions presented here, adjusted R-squared values are extraordinarily high: every regression exhibits values over 0.99. Although these results may seem "too good to be true" at first glance, we keep in mind several important details. First, we initially sorted the raw NYISO load and price into distinct time periods and carried out several data-smoothening transformations that included removing the annual MA trend from the data and a 90-day MA of the resulting deseasonalized data. Second, these transformed seasonal data were then fit to a quasi-periodic waveform function that already yielded high adjusted R-squared values in the majority of zones and time periods, generally clustered around 0.80 but some exceeding 0.90. In other words, the seasonal waveform function has already extracted a good deal of mathematical information from the data series and it comes as little surprise that the resulting residuals are expressed well as a purely time-dependent function of their past values.

In the following section, we validate our modelling approach and test its true usefulness in an applied setting by performing a series of out-of-sample forecasts. The results of these forecasts provide evidence that the high adjusted R-squared values from the AR regressions presented here are indicative of a truly good model fit.

4.8 Seasonal Model Out-of-Sample Forecasting

We next test the predictive power of the combined seasonal model and corresponding AR regressions by performing a series of out-of-sample forecasts. We use hourly NYISO real-time load and prices from 2019 and 2020 as our forecasting test data. Before running the forecasts, these raw data are sorted and filtered in exactly the same manner as the raw data from 2006 to 2018 as discussed in Section 4. We then calculated the predicted values of the test data

using the model parameters that we estimated using 2006 to 2018 NYISO data as presented in Tables 14–18. Finally, we use the model parameters estimated by the AR regressions in Section 4 to create predicted model residuals and subtract these from the observed seasonal model residuals. Note that for each out-of-sample forecast, we use the seasonal model and AR-regression parameters specific to each zone and time period.

To measure the accuracy of the seasonal model's in-sample fit and out-of-sample forecasts and compare forecast accuracy across all zones and time periods, we calculate the mean absolute error (MAE) for each forecast. The MAE is defined as

$$MAE = \frac{\sum_{i=1}^{n} |\tilde{y_i} - y_i|}{n}$$

where $\tilde{y_i}$ is the predicted value of the variable given from the model and y_i is the observed value of the variable we wish to predict, and n is the number of observations in the given data sample. From this equation, we see that as the distance between the predicted and observed values of the variable increases, the MAE also increases. Consequently, a higher MAE value indicates a poor out-of-sample forecast whereas a lower MAE value indicates a good out-of-sample forecast. In addition, because the MAE uses the absolute value of these distances, it penalizes over- and under-forecasted values identically. In the context of this analysis, a larger MAE value of the out-of-sample error terms indicates a weaker ability of a seasonal model to forecast the relationship between seasonal load and prices.

The out-of-sample MAE values measure how well the seasonal model predicts the relationship between the seasonal load and prices in a given zone and time period. Figures 19 and 20 depict the in-sample and out-of-sample forecast error terms, $\epsilon_i = \tilde{y_i} - y_i$, for evening time period in Zone A and Zone J, respectively. The in-sample error terms are colored in blue and the out-of-sample forecast error terms are colored in red.

Figure 19: Zone A evenings

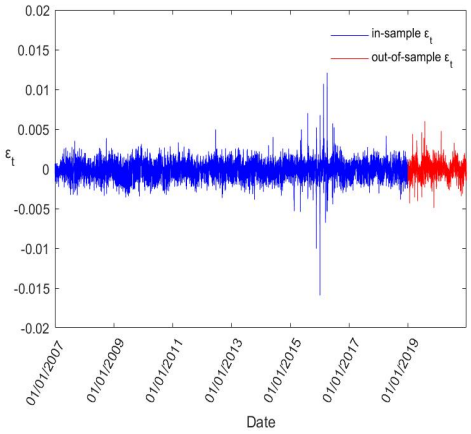

Figure 20: Zone J evenings

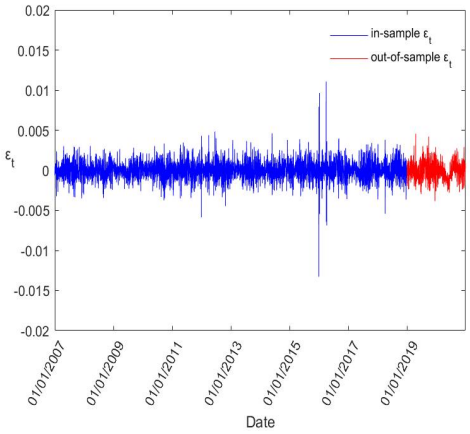

In both figures, we see that the average error value for the entire timespan of the data is approximately zero. With the exception of a few very large error terms, the volatility

remains fairly constant throughout the time span as well, including during the out-of-sample period. In addition, we can clearly see the impact of the Covid-19 global pandemic during the spring of 2020, where the MAE values shrink to almost zero. We suspect that this is because load patterns became more predictable as people were forced to stay home.[14]

4.8.1 Forecasting results

Table 27 shows the MAE values for all in-sample and out-of-sample model residuals across all NYISO market zones and time periods. These MAE values are presented as decimal representations of percentages.

Table 27: Mean absolute error of in-sample model residuals and out-of-sample forecast residuals in NYISO market zones

Zone	**in-sample**				**out-of-**	**sample**		
	Morning	**Midday**	**Evening**	**Night**	**Morning**	**Midday**	**Evening**	**Night**
Zone A	0.021	0.065	0.060	0.042	0.024	0.044	0.037	0.030
Zone B	0.025	0.043	0.046	0.036	0.014	0.021	0.025	0.024
Zone C	0.036	0.048	0.048	0.040	0.013	0.019	0.022	0.022
Zone D	0.024	0.074	0.063	0.053	0.018	0.078	0.069	0.078
Zone E	0.035	0.054	0.044	0.044	0.014	0.019	0.018	0.022
Zone F	0.028	0.035	0.035	0.029	0.013	0.013	0.016	0.020
Zone G	0.026	0.041	0.037	0.034	0.014	0.017	0.018	0.025
Zone H	0.032	0.047	0.049	0.041	0.020	0.034	0.038	0.032
Zone I	0.016	0.046	0.048	0.031	0.011	0.021	0.027	0.021
Zone J	0.022	0.029	0.026	0.027	0.016	0.024	0.024	0.025
Zone K	0.025	0.067	0.067	0.014	0.015	0.046	0.058	0.029

Table 27 has several patterns that are worth noting. First, all MAE values for both the in-sample model residuals and out-of-sample forecast residuals are less than 0.1%. Second, across all zones and time periods, the out-of-sample MAE value tends to be lower than the corresponding in-sample MAE value. This is despite the fact that we use static model parameters estimated from the in-sample data to perform the out-of-sample forecasts. We speculate that this is a result of the fact that we only forecast over two years of data in contrast to the twelve years of data used for the in-sample. With a forecast window that is short relative to the in-sample training period, there are theoretically fewer opportunities for large, unexpected market outcomes that yield large MAE values such as those seen in Figures 19 and 20 during 2016. We note, however, that our model yields these low MAE values despite significantly lower than expected load levels in the NYISO during the spring and summer of 2020 due to the COVID-19 pandemic (NYISO, 2020). Taking these facts into consideration, the low out-of-sample MAE values relative to the larger in-sample MAE values provide evidence that our modelling approach is capable of accurately forecasting

[14]NYISO, "COVID-19 and the Electric Grid: Load Shifts as New Yorkers Respond to Crisis." April 14, 2020. Available at: https://www.nyiso.com/-/covid-19-and-the-electric-grid-load-shifts-as-new-yorkers-respond-to-crisis.

the relationship between the seasonal load and prices in the NYISO, even when unexpected events occur. Next, across all zones and time periods, the MAE values for both the in-sample and out-sample model residuals tend to be smaller during the morning. For the in-sample residuals, the MAE values tend to be higher during the midday and evening time periods and then shrink during the night.

Finally, our modelling approach yields lower out-of-sample MAE values in some zones than others. Zones B, C, E, F, G, and J exhibit out-of-forecast MAE values below 0.026% across all time periods and Zone I exhibits out-of-forecast MAE values below 0.028% across all time periods. In contrast, Zones A, D, H, and K exhibit MAE values that are considerably higher than this 0.026% threshold. These elevated values are especially apparent during the midday and evening time periods, both of which are part of the on-peak portion of the day. Across all zones, Zone D exhibits the highest MAE values during the midday, evening, and night time periods, each of which is approximately three times higher than the corresponding MAE values in the NYISO market zones where the model fits best. Zone A exhibits the highest forecast MAE value for the morning time period.

4.8.2 Out-of-sample forecast implications for risk management

The results of these out-of-sample forecasts have implications for LSEs or other market participants who wish to mitigate spot market risk by purchasing futures contracts for power. The signal processing analysis in this study is a tool that can yield an independent estimate of power prices given a load forecast or vice versa.

In seven of the eleven NYISO zones, the MAE values for the out-of-sample forecasts have all been approximately the same and below a threshold of 2.6E-4. Although this outcome is statistical in nature, it is nevertheless evidence that our signal processing methodology is able to filter and model the load and price data in these zones such that the relationship between these two time series exhibits approximately the same risk profile in each zone. Higher MAE values for the out-of-sample forecasts in Zones A, D, H, and K, however, indicate that the relationship between average seasonal load and prices in these zones is more difficult to predict than in the remaining zones. In other words, the relationship between these two data series is less predictable using our signal processing approach and an accurate seasonal load or price forecast in these zones is less likely to produce an accurate forecast for the corresponding seasonal price or load level.

An LSE looking to mitigate spot market risk should choose more aggressive hedging strategies in Zones A, D, H, and K than the other zones. Such a strategy could mean hedging a larger share of the expected load in these zones even if futures market prices are higher than the seasonal model estimates. For commodity market speculators, the results of our analysis indicate that market participants who face spot market price risk might be willing to pay a higher risk premium for futures contracts in these zones.

4.9 Conclusion

In this paper, we presented a signal processing approach to modeling the medium-term relationship between load and prices in the NYISO wholesale electricity market. Although signal processing is a modeling paradigm that has long been commonplace in engineering and other applied sciences, it has more recently received increased attention in the field of quantitative finance. It is primarily comprised of three modeling steps: a filtration step, where raw data is transformed to allow for easier model fitting; determining a quasi-periodic waveform function that describes the filtered signal's frequency and amplitude; and modelling any remaining noise component of the signal that is not well described by the waveform function.

The highly seasonal cyclicality and fairly stable stochastic drift seen in electricity market load and prices make them good candidate variables for signal processing. In addition, there are real-world risk management applications of modeling the relationship between these variables as many power market participants seek to hedge their price risks by purchasing futures contracts. We chose to carry out our analysis on the NYISO because, in addition to being publicly available and easily accessible online, NYISO market data is well maintained, organized neatly into eleven market areas, and goes back to 2006.

We first transformed raw NYISO market data into moving averages in order to smoothen the data, which constituted the filtering step. Next, we fit a series of quasi-periodic waveform functions to these filtered data series in order to describe the cyclical component of the signal they exhibit. Finally, after extracting the modelled waveforms from the filtered data, we modelled the model residuals, which comprised the signal noise, by using the Box–Jenkins time-series modeling methodology.

To validate the predictive capability of our modelling approach, we performed a series of out-of-sample forecasts on test data from 2019 and 2020 using our seasonal model. The results provided evidence that the approach developed in this paper was able to forecast the relationship between load and prices in the NYISO with a high degree of accuracy. These predictions could act as an independent assessment of whether futures contracts for power are trading at a premium, discount, or fair market value and assist a NYISO market participant in making more informed hedging decisions.

Several avenues of future research could build upon the findings presented in this paper. A simple extension could be to roll the parameter estimation process of the waveform functions so that model coefficients are updated with each new observation. Further research could also seek to combine models with different MA time windows in the hope that a multiple-window approach will improve forecasting. Ambitious researchers with access to historical NYSIO power futures data and experience with programming algorithms could seek to back test hedging strategies generated by our model and compare it to other risk management models. Finally, signal processing could be used to model markets for other commodities that exhibit strong seasonal fluctuations such as agricultural products like corn, wheat, and

soybeans. Data on the demand for such products will likely be of lower frequency than power market load data but the results would be no less interesting.

References

[1] Aggarwal, S. K., Saini, L. M., and Kumar, A., "Electricity price forecasting in deregulated markets: A review and evaluation," *International Journal of Electrical Power & Energy Systems*, vol. 31, no. 1, pp. 13-22, 2009.

[2] Amjady, N., "Short-term bus load forecasting of power systems by a new hybrid method," *IEEE Trans. Power Syst.*, vol. 20, no. 1, pp. 333-341, 2007.

[3] Benth, E. F., and Schmeck, D. M., "Pricing futures and options in electricity markets," The Interrelationship Between Fin. and Energy Markets, pp. 233-260, 2014.

[4] Bernstein, M.A., and Griffin, J., (2006), "Regional differences in the price-elasticity of demand for energy, " *National Renewable Energy Lab.(NREL), Golden, CO (United States)l.*

[5] Bo-Juen, Ch., and Chang, M. W., "Load forecasting using support vector machines: A study on EUNITE competition 2001," *IEEE Trans. Power Syst.*, vol. 19, no. 4, pp. 1821-1830, 2004.

[6] Box, G., and Jenkins, G., (1976), "Time-series analysis: forecasting and control. London: Holdens Day, " *Inc.*

[7] Chang, P.-C., Fan, C.-Y., and Lin, J.J., "Monthly electricity demand forecasting based on a weighted evolving fuzzy neural network approach," *Int. J. of Electric. Power & Energy Syst.*, vol. 10, no. 9, pp. 963-974, 2010.

[8] Charytoniuk, W., Chen, and M. S., "Very short-term load forecasting using artificial neural networks," *IEEE Trans. Power Syst.*, vol. 15, no. 1, pp. 263-268, 2000.

[9] Chen, J., Deng, S. J., and Huo, X., "Electricity price curve modelling and forecasting by manifold learning," *IEEE Trans. Power Syst.*, vol. 23, no. 3, pp. 877-888, 2008.

[10] Con Edison., (2019), "Annual Report, " *Con Edison.*

[11] Craan, G. E., (2019), "NYISO Energy Marketplace, " *New York Market Orientation Course.*

[12] Christiansen, W. R., "Short-term load forecasting using general exponential smoothing," *IEEE Trans. Power App. Syst.*, no. 2, pp. 900-911, 1971.

[13] Deryugina, T., and MacKay, A., and Reif, J., (2017), "The long-run elasticity of electricity demand: Evidence from municipal electric aggregation, " *American Economic Journal: Applied Economics Forthcoming.*

[14] Engle, R.F., Mustafa, C., and Rice, J., "Modeling Peak Electricity Demand," *Journal of Forecasting*, vol. 11, pp. 241–251, 1992.

[15] Feehan, J. P., (2018), "The long-run price elasticity of residential demand for electricity: Results from a natural experiment, " *Utilities Policy*, Vol.51, pp. 12-17.

[16] Ghiassi, M., and Zimbra, K. D., "Medium term system load forecasting with a dynamic artificial neural network model," *IEEE Trans. Power Syst.*, vol. 17, no. 3, pp. 626-632, 2002.

[17] Goude, Y., Nedellec, R., and Kong, N., "Local short and middle term electricity load forecasting with semi-parametric additive models," *IEEE Trans. on Smart Grid*, vol. 5, no. 1, pp. 440-446, 2013.

[18] Han L., Peng, Y., Li, Y., Yong, B., Zhou, Q., and Shu, L., "Enhanced deep networks for short-term and medium-term load forecasting," *IEEE Access*, vol. 7, pp. 4045-4055, 2018.

[19] Klemola, A.,i and Sihvonen, J., (2015), "Covered Option Strategies in Nordic Electricity Markets, " *Journal of Energy Markets, September.*

[20] Klüppelberg, C., Meyer-Brandis, T., and Schmidt, A., "Electricity spot price modelling with a view towards extreme spike risk," *Quantitative Finance*, vol. 10, no. 9, pp. 963-974, 2010.

[21] Lloyd, J. R., "GEFCom2012 hierarchical load forecasting: Gradient boosting machines and Gaussian processes," *International Journal of Forecasting*, vol. 10, no. 2, pp. 369-374, 2014.

[22] Lucia, J. J., and Schwartz, S. E., "Electricity prices and power derivatives: Evidence from the nordic power exchange," *Rev. of Derivatives Res.*, vol. 5, no. 1, pp. 5-50, 2002.

[23] Nepal. Sh., (2015), "Signal Processing in Finance, " *Electrical and Computer Engineering Design Handbook.*

[24] New York ISO Annual grid & Markets Report., (2019), "Reliability and a Greener Grid, " *New York Independent System Operator.*

[25] Patton, B. D., and LeeVanSchaick, P., and Chen J., and Naga, P. R., (2020), "State of the market report for the New York ISO markets 2019, " *Potomac Economics.*

[26] Stoft, S., (2002), " *Power System Economics*, IEEE Press, Piscataway NJ, USA.

[27] State of the Market Report for PJM., (2014)," *Monitoring Analytics, LLC* , Independent Market Monitor for PJM, USA.

[28] State of the Market Report for PJM., (2020)," *Monitoring Analytics, LLC* , Independent Market Monitor for PJM, USA.

[29] J. W. Taylor, and P. E. McSharry, "Short-term load forecasting methods: An evaluation based on European data," *IEEE Trans. Power Syst.*, vol. 22, no. 4, pp. 2213-2219, 2007.

[30] J. W. Taylor, "Short-term load forecasting with exponentially weighted methods," *IEEE Trans. Power Syst.*, vol. 27, no. 1, pp. 458-464, 2011.

[31] J. Taylor, and B. Buizza, "Neural network load forecasting with weather ensemble predictions," *Electric Power Syst. Res.*, vol. 76, no. 5, pp. 302-316, 2006.

[32] Washington Utilities and Transportation Commission..(2017), "Docket UG-132019: Policy and Interpretative Statement on Local Distribution Companies' Natural Gas Hedging Practices".