ABSTRACT

Title of dissertation: THE ECONOMICS OF WIND POWER AND WHISKY

Ian Brinkmann Page, Doctor of Philosophy, 2015

Dissertation directed by: Professor Roberton Williams
Department of Agricultural and Resource Economics

For many goods, economists tend to ignore the time between the onset of production and the final sale of the good. In some instances, economists do model production with intertemporal considerations such as the extraction of groundwater, population dynamics in fisheries, and manufacturing with costs characterized by learning-by-doing; but even in these cases output at any point in time tends to be unconstrained except when production is limited according to resource availability. In the following three essays, I examine the implications for agent and market behavior when producers cannot perfectly adjust output over time.

In the first two essays, I focus on the relationship between electricity markets and both conventional and renewable power producers. Specifically, in the first essay, I quantify the effect a large level of installed wind power capacity (an intermittent renewable energy) has on power market conditions. Because wind power has virtually no marginal costs for generation, and its output cannot be perfectly controlled, a high penetration level of wind power could potentially lower average prices while also impacting price volatility.

In the second essay, I construct a computational model of a conventional power producer that cannot perfectly adjust its output over time and faces prices that change according to a stochastic process. Then, I measure the impact price volatility has on producers in two ways. First, I analyze changes in their optimal generation strategies in light of price volatility, and then I simulate and track changes in output, profit, and emissions over time.

My third essay pertains to production of whisky. While there are other examples of vintage goods that require a significant amount of maturation, the existing literature typically assumes that there is a unique optimal maturation age for a producer's inventory. However, many Scottish distilleries produce a line of whiskies that vary primarily according to age. I demonstrate that it is possible for a profit-maximizing distillery to mature multiple ages of whisky without market power, but a further exploration of distilleries' product lines indicates the market is far from perfectly competitive.

THE ECONOMICS OF WIND POWER AND WHISKY

by

Ian Brinkmann Page

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:

Professor Roberton Williams, Chair
Professor Anna Alberini
Professor Sébastien Houde
Professor Erik Lichtenberg
Professor Andrew Sweeting

DEDICATION

For my mother and father

and Elina

ACKNOWLEDGEMENTS

I am truly grateful to the many people who have supported me as I worked on my dissertation at the University of Maryland. I am indebted to my advisor, Rob Williams, for his constant guidance, insight, and encouragement. I also wish to thank him for his trust on every occasion I had the privilege of being his teaching assistant. Erik Lichtenberg similarly served as my mentor and helped me with both my pedagogical skills and dissertation. Anna Alberini and Ted McConnell are responsible for innumerable suggestions for improvements to my work. I also thank Sébastien Houde and Andrew Sweeting for serving on my committee.

Many of the graduate students in the department of Agricultural and Resource Economics are also deserving of recognition. Davide Cerruti, Mark Miller, Kota Minegishi, Dan Werner, and Daniel Voica provided me with an invaluable sense of camaraderie and countless opportunities to discuss preliminary work. I am also indebted to all other graduate students who attended my workshops and seminars.

I must also thank the many amazing teachers I have had the great fortune to learn from as a student and mentee. This necessarily includes the faculty in the department of Agricultural and Resource Economics, but also dedicated teachers such as Susanne Quandt and Mary Betancourt. Without their amazing work ethic and passion, I could never have come this far.

Lastly, I feel extremely blessed for the immeasurable support I have received from my family. Elina has been an unwavering source of inspiration and encouragement since we met at Math Camp many years ago. My parents, Annette and Rod, have individually and together shaped the person I am today. As their son, I have been bestowed an incredible number of opportunities to improve myself and realize my potential. And finally, I thank my brother and sister, Derek and Erika, whose kind words have always suggested a certain sense of pride in their middle brother that has carried me through the difficult times. Thank you.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

The Effects of Wind Energy on Power Market Conditions

1. Introduction

Despite the potential environmental and economic benefits of wind power, integrating a large amount of wind power capacity into a power system is generally perceived as difficult. Ideally, generators should be responsive to system needs and can be adjusted as system conditions evolve over the course of a day. Yet wind turbines, unlike conventional fossil fuel-burning generators, rely on a stochastic environmental factor to produce electricity. Subsequently, wind power is intermittent and can be difficult to forecast and control. It is thus unclear whether existing power systems can accommodate a large amount of wind power capacity, as it necessarily leads to a power supply that is less predictable and responsive to system needs. However, installation of wind power capacity is expected to continue to grow in the coming decades as costs decline and older conventional units are retired. It is therefore important to study the current system impacts of wind power so that potential future complications from wind power's variable output are well understood.

For power systems that coordinate supply decisions through power markets, the intermittency of wind power may be detected through changes in market conditions. Unfortunately, research on the market impacts of wind power has been limited so far because of a lack of historical data. Early work by Morthorst (2003) and Parsons et al. (2004) find that the initial, low levels of wind power in Danish and US systems had an insignificant effect on the price for power. More recent studies by Jacobsen and Zvingilaite (2010), Jónsson et al. (2010), Cutler et al. (2011), Woo et al. (2011a,b), Ketterer (2014), and Clò et al. (2015) benefit from several additional years of growth in the wind power sector, and they are able to demonstrate that higher levels of wind power do lead to lower power market prices. However, there is some disagreement as to how to best capture the effects of wind power's intermittency on power system reliability. More recent work emphasizes the consequences of intermittency by determining its impact on price volatility, though there is no consensus as to how to characterize price volatility among these studies. Furthermore, earlier results from Morthhorst and Parsons et al. suggest that wind power's intermittency can affect markets for ancillary services, which are various forms of backup power supply services used to maintain the reliability of the overall power supply. The more recent literature generally omits ancillary services in their analyses entirely.

The purpose of this paper is to completely quantify the impacts of wind power on power market conditions to determine how well existing systems can accommodate the introduction of a large amount of intermittent renewable energy. Accordingly, I divide market conditions into three separate components and estimate wind power's effect on each of them. This includes estimating the impacts of wind power on the price for power, its volatility, and prices for ancillary services. This study specifically focuses on the power system in Texas from 2003 to 2010, during which time installed wind power grew from approximately 1% to 10% of total capacity in the system. The effects of wind power on market conditions are estimated using both Ordinary Least Squares [OLS] and quantile regressions.

The contributions of this work to the existing literature are threefold. Most importantly, it presents a comprehensive overview of wind power's market impacts.

Earlier papers emphasize the effect of wind power on power prices, but they overlook the overall consequences of wind's intermittency by only including either price volatility or prices for ancillary services in their analyses, not both. Additionally, I give considerable attention to the characterization of price volatility and propose my own method for measuring volatility in a way that consolidates previous studies' adopted metrics. Lastly, I use quantile regressions in conjunction with OLS to estimate wind power's effects on market conditions. Jónsson et al. (2010) identify quantile regressions as one of several non- or semi-parametric techniques researchers should use to evaluate wind power's impacts, though to the best of my knowledge, this paper is the first such application.

Results from my empirical estimations confirm that an increase in wind power capacity has a negative effect on power market prices. This effect is especially pronounced on power prices in western Texas, where wind power capacity is concentrated. I find that wind power has no clear effect on price volatility, even though I consider multiple specifications to characterize it. Additionally, results show that wind power decreases prices for ancillary services, indicating that an increase in an intermittent renewable does not necessarily lead to an increase in system costs. These results confirm that that wind power does have economic and environmental advantages, as demonstrated by its ability to lower power market prices and thereby displace power from conventional generators, and that concerns over wind power's intermittency may be overstated, as wind power has no apparent adverse impact on either price volatility or prices for ancillary services.

The remainder of the paper is organized as follows. First, I present an overview of the power system in Texas. This includes a description of the market setup and typical power producers, as well as the application of a theoretical model to demonstrate the effects of wind power market prices. Next, I review the existing literature's methods for quantifying the market impacts of wind power. Then, I discuss my own empirical specifications, as well as detail the characterization of price volatility and summarize the data for my study. Lastly, I present results from the empirical analyses and discuss implications for the future.

2. The Texas Interconnection

2.A. Market Setup

The Texas interconnection is one of three power systems in the contiguous US, though the Texas interconnection itself is quite isolated and very little electricity flows between it and the other two interconnections. It was regulated until 1995, when the state legislature voted in favor of deregulation and allowed for wholesale competition. The Electric Reliability Council of Texas [ERCOT] was then established and made responsible for implementing a suitable market structure that would facilitate competition among power producers, ensure fair access to the power system, and guarantee reliability of transmission.

Scheduling generation in the Texas interconnection occurs through a series of markets, though participation in any one market is not mandatory. Power producers and retailers first have the option of negotiating contracts to schedule supply over an extended period of time. Successful contracts are relayed to ERCOT, which verifies whether the grid can accommodate the scheduled supply. Because market participants are unable to perfectly forecast demand or available generation capacity, contracted generation is often insufficient when compared to actual need. In general, deregulated power systems often operate a real-time market where additional generation services are bought and sold to avoid imbalances between supply and demand. Whereas most power systems refer to this real-time market as a "spot market," ERCOT refers to it as a "balancing market," although the distinction is apparently trivial, as many power companies within interconnection still refer to it as a spot market. Regardless of its designation, this market operates in real-time, settles imbalances between supply and demand, and shapes market participants' perceptions of future power prices when negotiating contracts. In theory, both suppliers and consumers can participate in the balancing market, though in practice demand-side participation is quite low, and demand is assumed to be exogenous with respect to real-time prices.[1]

Historically, the balancing market was divided into four zones. As long as power lines in the system were not congested, power would flow freely between the North, Houston, South, and West zones, and each zone would have the same balancing price for power. However, if there was insufficient transmission capacity between zones, power lines would become congested. Balancing prices would then vary across zones according to load levels, generation, and the availability of transmission capacity. Beginning in December 2010, ERCOT switched from a zonal to a nodal setup to better incorporate transmission constraints and improve market efficiency. With a nodal market, the balancing market is divided geographically to a finer degree, and power prices vary from

---

[1] The assumption of demand's exogeneity with respect to real-time prices is universal in the literature, regardless of which power system is being discussed. In regards to the Texas interconnection, Zarnikau and Hallett (2008) estimate industrial consumers' elasticity of demand with respect to balancing prices and find that it is negligible, only -0.000008 on average. Factors other than price, such as time of day, whether it is a workday, and weather are much stronger determinants of demand.

node to node depending on transmission and demand at each node. Both zonal and nodal markets in ERCOT are depicted in Figure 1.

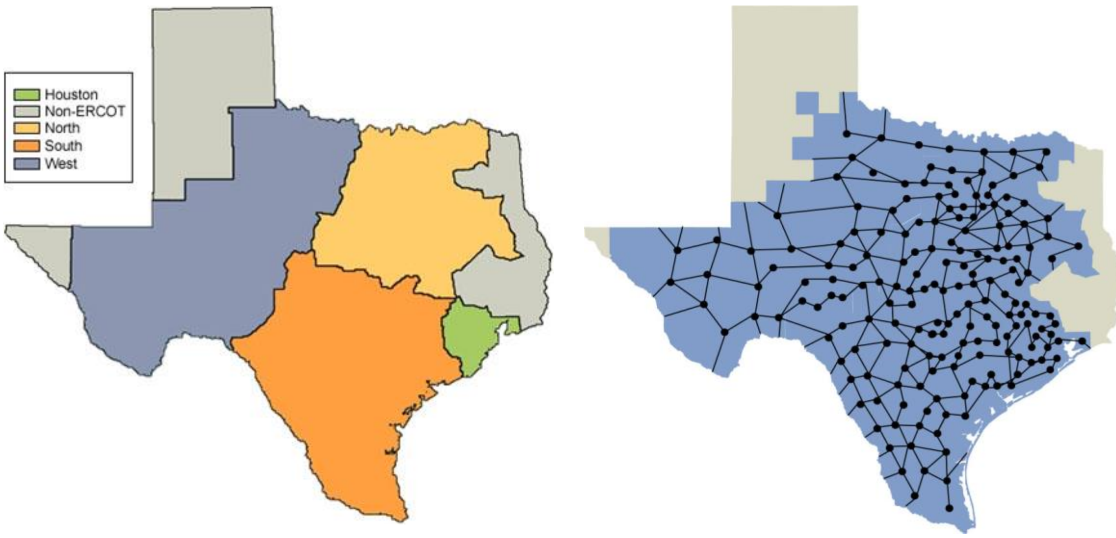

Figure 1. *ERCOT Zonal and Nodal Markets*

Figure 2 demonstrates the possible behavior of ERCOT's balancing markets during its zonal setup. In this example, the prices for each of the four zones were identical between 9am and 10pm, when demand for power was high throughout the state. But prices were noticeably different in the early morning and late at night, and this was especially pronounced in the West zone. The West zone has historically had both the smallest population and highest concentration of wind power capacity. It was often the case that wind farms were more active at night, but during these hours there was insufficient demand for electricity in the West zone and not enough transmission capacity to carry the power to other parts of the state. As a result, the price for power in the West zone was lowered until producers appropriately adjusted their output. In this instance, the price became negative in the West, indicating that the power producers were actually paying for the right to generate electricity. Although this may be counterintuitive, wind power is subsidized per watthour of generation in accordance with the Federal Production Tax Credit. As a result, wind power generation can still be profitable even when the price is negative, so long as the absolute value of the price does not exceed that of the subsidy.

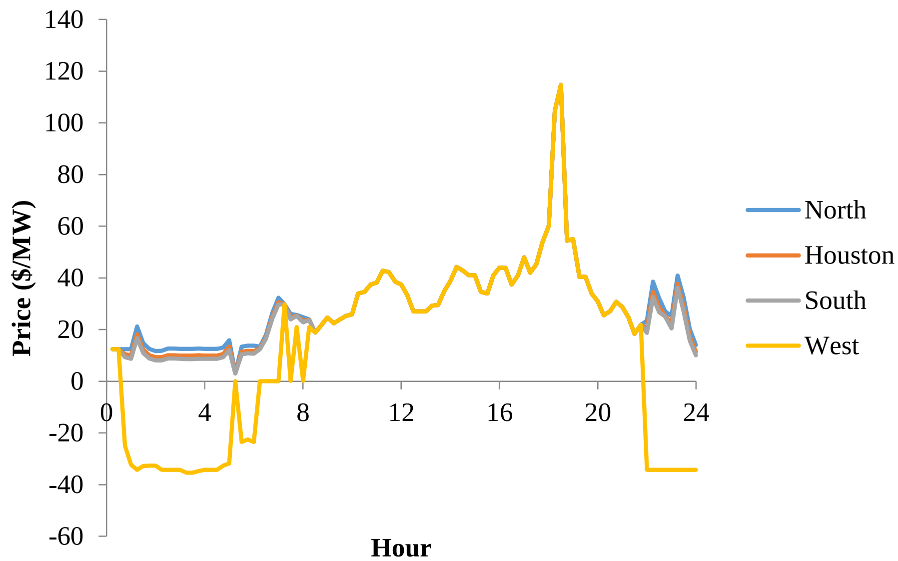

Figure 2. *Balancing Market Prices—January 9, 2009*

Other than the balancing market, ancillary supply services are scheduled one day ahead to maintain the reliability of the power supply. This secondary market exists as insurance in case the balancing market cannot achieve equilibrium, and also because smaller, more rapid changes in producer output are sometimes required and the balancing market cannot facilitate these small changes. Ancillary supply services in ERCOT include "regulation up" and "regulation down," as well as backup generation from "responsive reserves" and "non-spinning reserves." "Regulation up" applies to instances in which a producer adjusts its output so as to increase the frequency of electrical load, and "regulation down" pertains to adjusting output to decrease load frequency. These adjustments to output are small but nearly instantaneous, and are vital to maintaining a close balance between supply and demand over time. “Responsive reserves” and “non-spinning reserves” are both forms of backup generation, and they operate on a longer time scale than regulation. Responsive reserves are backup generators that can either quickly turn on or stop generating electricity. Non-spinning reserves are ancillary generators that are off-line but can begin producing power within thirty minutes.

2.B. Market Participants

The majority of power in Texas comes from fossil fuels. Roughly two-thirds of installed capacity in Texas uses natural gas as its primary fuel, making it the largest source of electricity by far. Coal is the second most common fuel source in the interconnection and accounts for roughly 20% of all installed capacity. Traditionally, coal plants have had much lower marginal costs than natural gas, which has led to coal plants serving base-load demand and being run throughout the day. Additionally, coal plants cannot easily adjust their output and are costly to start up if turned off. Natural gas generators, which tend to be faster-ramping but more expensive to operate, typically serve as peak-load units, though the expansion of efficient combined cycle gas generators and lower natural gas prices have led to the presence of natural gas plants that can effectively serve base-load.

Aside from natural gas and coal, nuclear and wind power are the only other major energy sources in the Texas interconnection. There are two nuclear plants, and together they make up about 5% of all installed capacity. Nuclear generation tends to be the cheapest dispatchable generation available in the system, so it serves as a base-load power source and its output rarely, if ever, responds to power market conditions. Like nuclear, wind power has a very low marginal cost of output, but it is not considered a base-load power source because of its intermittency. Often, wind power output is treated as negative load for analytical purposes because its output is stochastic and does not readily respond to market prices.

The penetration level of wind power varied significantly during ERCOT's zonal market phase. In the beginning of 2003, wind power represented about 1% of all capacity in the system, and by the end of 2010, installed wind power's penetration level was approximately 9.5% of all capacity. Figure 3 demonstrates the evolution of installed wind power capacity over time. Significant factors driving growth in installed wind power in Texas include the Federal Production Tax Credit, the state's Renewable Portfolio Standard, and strong wind resources (Bird et al., 2005). Wind power in the Texas interconnection is comparatively higher than in all other US power systems, though much of it is located in western Texas, where it is far removed from most major cities and industrial hubs, and requires substantial transmission capacity in order to reliably meet demand.

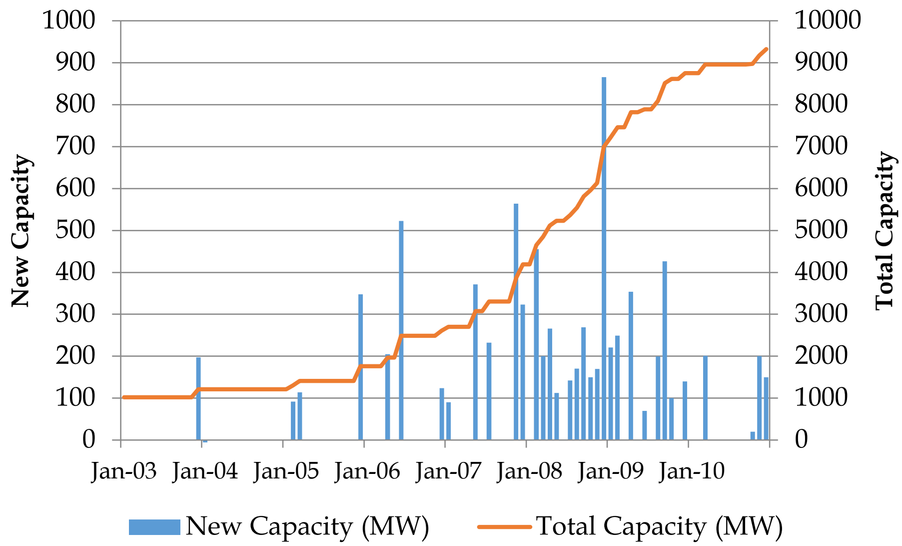

Figure 3. *Wind Power in Texas*

2.C. The Market Impacts of Wind Power

Power markets exist to coordinate generation between producers so that system-wide costs are minimized while also ensuring that the power supply is reliable even in the event of unforeseen circumstances such as a disruption in transmission capacity, an unplanned outage at a plant, or an unpredicted spike in demand. To evaluate the market impacts of wind power, as well as the consequences on total generation and total ancillary services, consider the following model. First, let g be the total amount of generation (measured in Megawatts) from conventional power plants such as nuclear, coal, and natural gas. Furthermore, let $\mathrm{c}(g)$ be the total cost from such generation. If the market efficiently assigns generation responsibilities, lowest-cost generators are activated first, and sequentially more expensive plants are brought online as needed. This naturally translates to $\mathrm{c}(\cdot)$ increasing in g. Assuming that $\mathrm{c}(\cdot)$ is differentiable, and denoting the first-order derivative as c_g, this implies $\mathrm{c}_g > 0$.

Next, let ancillary services be denoted as a (also measured in Megawatts) and total ancillary service costs be calculated according to the function $\mathrm{d}(a, g)$. Ancillary costs are also assumed to be differentiable and increasing in a, $\mathrm{d}_a > 0$. Generation is also included in ancillary service costs because of the relationship between generation and ancillary services. If units are needed for generation, they are run for that purpose and are not available for ancillary services; but if they are not actively generating power for the system (because their marginal cost is higher than the market price), they are available to provide ancillary services. This defines a competitive relationship between generator allocation for g and a. As more generators are used towards g, fewer are available for a, and as a result the units used towards a will be of even higher cost were g lower. Thus, holding a constant but increasing g affects the cheapest units available for ancillary services, $\mathrm{d}_g > 0$, $\mathrm{d}_{ag} > 0$.

Putting the two together, the power market determines total generation and ancillary services to minimize costs

$$\min \mathrm{c}(g) + \mathrm{d}(a, g),$$

subject to two constraints. First, generation from conventional sources plus generation from wind power must equal or exceed current demand

$$g + \omega \geq D,$$

where ω is wind power and D is total demand (also referred to as load). Additionally, the system operator (ERCOT) is concerned about system reliability, hence selects ancillary services to maintain reliability according to the constraint

$$\mathrm{f}(a, L, \mathrm{v}) \geq 0.$$

The reliability function f is increasing in ancillary services ($\mathrm{f}_a > 0$), and decreasing in both net load ($\mathrm{f}_L < 0$) and volatility conditions ($\mathrm{f}_{\mathrm{v}} < 0$). Net load is defined as $L = D - \omega$, and volatility is a non-decreasing function in both load and wind power ($\mathrm{v}_D \geq 0, \mathrm{v}_\omega \geq 0$). Reliability is decreasing in net load because greater demand will generally require additional generators on standby, effectively creating a larger reserve margin of power producers. Furthermore, system volatility (either considered as fluctuations in load or net load, or price volatility) is increasing in both demand and wind power because of their uncontrollable and erratic natures.

Assembling all relevant terms and functions, the general activities of the power market can be summarized using the Lagrangian

$$\mathcal{L} = \mathrm{c}(g) + \mathrm{d}(a, g) - \lambda(g + \omega - D) - \mu\mathrm{f}(a, L, \mathrm{v}).$$

While first order conditions will depend on the specifications of ω and D, as well as the shape of the reliability function f, the most common case encountered in daily operations is described with a strictly interior solution for g^* and a^*, and with both constraints binding:

$$g^*)\ \mathrm{c}_g + \mathrm{d}_g - \lambda^* = 0$$

$$a^*)\ \mathrm{d}_a - \mu^*\mathrm{f}_a = 0$$

$$\lambda^*)\ g^* + \omega - D = 0$$

$$\mu^*)\ \mathrm{f}(a^*, L, \mathrm{v}) = 0.$$

Market prices are set to induce producer output to satisfy the above cost-minimizing conditions. That is, in order to induce producers to generate g^*, the market price must be high enough to cover the operational costs of the highest-cost producer; the price must therefore be equal to c_g evaluated at g^*. The price for ancillary services is similarly identified.

In the model, parameters such as demand and wind power are considered fixed, and variables such as generation and ancillary services can be expressed as a function of these parameters. Yet wind power is only exogenous for the system operator because wind conditions are beyond its control. Wind power can still increase or decrease according to wind conditions or installed capacity. Accordingly, for a marginal change in wind power, g^* and a^* are affected thus:

$$\frac{\partial g^*}{\partial \omega} = -1$$

and

$$\frac{\partial a^*}{\partial \omega} = \frac{f_L - f_v v_\omega}{f_a}.$$

The relationship between wind power and generation is straightforward. Whatever electricity is generated using wind displaces an equal amount of electricity that would have been generated from conventional sources. Furthermore, this marginal decrease in g^* would lead to lower prices for generation, since $c_g > 0$, and prices are set to meet the marginal cost of generation.

The effect wind power has on ancillary services is less obvious. Because $f_a > 0$, the sign of $\partial a^*/\partial \omega$ is determined by the numerator. Essentially, there are two competing effects from wind power that could lead to either a decrease or increase in the reliance on ancillary services depending on which is the dominant effect. First, additional wind power decreases net load, leading to less overall need for ancillary services in the system.

However, additional wind power also increases volatility, which increases the need for ancillary services.

Regardless of the net effect on a^*, it is possible that prices for ancillary services could increase or decrease because of a change in wind power. First, if a^* does not increase, ancillary service prices will unequivocally decrease because $\partial g^*/\partial\omega = -1$ and $\mathrm{d}_g > 0$, and $\mathrm{d}_{ag} > 0$. On the other hand, an increase of a^* will put additional pressure on ancillary service costs (and therefore prices). But because a marginal increase in wind power leads to a marginal decrease in generation, this frees up generators that would have been used for g that can now be used towards a. In that case, ancillary service prices might increase or decrease, and it will depend on the change in a^* as well as the structure of $\mathrm{d}(a, g)$.

3. Literature Review

Installed capacity of intermittent renewables has been low in power systems around the world until quite recently. Consequently, there is very little empirical evidence on how well power systems have been able to accommodate wind power, despite concerns about its fluctuating output. Much of the existing research on intermittent renewable energy relies on simulations to determine the potential impacts on system reliability and market conditions (Karki and Billinton, 2001 and 2004; Lund and Münster, 2003a,b; Chen et al., 2006; Lund, 2006; Sensfuß et al., 2008; Green and Vasilakos, 2010; Delarue et al., 2011; Milstein and Tishler, 2011; Di Cosmo and Malaguzzi Valeri, 2014; Shcherbakova et al., 2014), emissions reductions from conventional generators (Benitez et al., 2008; Delarue et al., 2009), or both market conditions and abatement (Holttinen and Tuhkanen, 2004; Luickx et al., 2010). While many of these studies have focused on power systems with liberalized power markets, they rely on simulations to assess the feasibility of intermittent renewable energy either because the capacity being modeled in the system is hypothetical (and often reflects the goals of new policies) or because the relevant market data are not publicly available.

One of the earliest papers that relies on market data to study the impact of wind power on market prices is Morthorst (2003). Morthorst's analysis focuses on the Nord Pool power exchange, which includes the western area of Denmark. Even then, the western area had some of the highest penetration levels of wind power in the world. Yet Morthorst's findings are not particularly conclusive. He demonstrates that there is a general tendency for spot market prices to be lower when wind power is higher, and that spot market prices are higher when generation from wind is low, but no statistically strong relationship is found. And while wind power does increase the need for conventional power producers to adjust their output in response to fluctuations from wind power, there is not much of an effect on the price of this ancillary service. An analysis by Parsons et al. (2004) finds that wind power will increase expenditures for this ancillary service in the US in the future, though these costs were relatively low and expected to remain so for the next several years.

Following Morthorst (2003), Chang et al. (2009), Jacobsen and Zvingilaite (2010), and Jónsson et al. (2010) conduct more recent analyses of wind power market impacts in Denmark. Chang et al. study the integration costs of an offshore wind farm by evaluating correlations between wind power output on market prices and price volatility. They find that wind power from a single offshore site has an insignificant effect on market prices, and that the net effect of its intermittency is negligible. Jónsson et al. are able to demonstrate that total forecasted wind power likewise has an impact on average day-ahead prices, and that it also influences the variance, skewness, and kurtosis of prices. Jacobsen and Zvingilaite present a general overview of price behavior in western Denmark, including differences in average prices and prices' distributional properties between other market areas within the same interconnection.

Due the historically high levels of wind power in Denmark and the availability of market data, the Danish system has received a considerable amount attention. However, as wind power continues to expand in power systems around the world, researchers have been able to analyze the market impacts in a variety of new power systems, including Australia (Cutler et al., 2011), Texas (Woo et al., 2011a,b), Ireland (Di Cosmo and Malaguzzi Valeri, 2012), Germany (Ketterer, 2014), and Italy (Clò et al., 2015). Newer studies likewise favor different methodologies, ranging from presenting a general overview of price behavior (as in Cutler et al.) to an AR-GARCH model that uses wind power as both an explanatory variable and as a determinant of the conditional heteroskedasticity function (Ketterer). A consistent finding is the inverse relationship between wind power and mean prices, though the magnitude of this effect varies.

Woo et al. (2011a,b) present some of the first findings on the impact of wind power on market conditions in the Texas interconnection. Using a two-stage model, Woo et al. (2011a) study the determinants of price divergence in Texas's four zonal balancing markets and find that wind power has a strong and statistically significant effect on the likelihood of price divergence. In another paper, Woo et al. (2011b) focus on the effect wind power has on the balancing markets' mean price and price volatility, where volatility is estimated using the forecast variance formula suggested by Feldstein (1971).[2] Their results indicate that the average marginal effect of wind on prices varies across zones, and that these effects are particularly pronounced in the West zone of Texas, where the majority of wind power is concentrated. They also find that wind power increased the variance of prices, though the effect is somewhat small.

Woo et al. (2011b) present a compelling case study of wind power's market impacts, but additional insight can be gained through further research on the Texas interconnection for several reasons. First, their study omits the market impacts of wind power on ancillary services, even though earlier studies (e.g., Morthorst, 2003 and Parsons et al., 2004) argue that wind power would affect these markets as it becomes increasingly prevalent. It is also important to note that Woo et al. studied the impact of wind power on price volatility by measuring price variance. While it may be common in the finance literature to use variance or standard deviation to study volatility, it does not fully encapsulate the meaning of "volatility" in this instance. An increased penetration of intermittent renewable energy will not only increase variance about the mean. It is also expected to increase the occurrence of extremely high and low prices, as found in Cutler et al. (2011), a concept which is better measured with a distribution's kurtosis and skewness, not just the distribution's variance.

The highly non-linear effect of wind on prices convinced Jónsson et al. (2010) to use a non-parametric model, and they include an analysis of price skewness and kurtosis in addition to price variance and mean prices. However, Jónsson et al.'s analysis of the first four moments relates to daily distributional properties, not intraday price

[2] This method was also adopted by Clò et al. (2015).

movements. A logical extension to the existing literature is therefore to look at the intraday price behavior in Texas, as in Woo et al. (2011a,b), but to include ancillary service prices and use a non- or semi-parametric technique to better analyze potentially non-linear effects on price and price volatility. And variance, skewness, and kurtosis should all be included to better quantify the relationship between wind power and price volatility.

4. Methodology

4.A. Effects on Prices

To assess the effect of wind power on power market prices, the general empirical estimation assumes the form:

$$(1)\ \text{price}_{it} = \alpha_i + \beta_i \text{wind}_t + \gamma_i \text{NGprice}_t + \delta_i \text{load}_t + \theta_i \text{nuclear}_t + \varepsilon_{it}.$$

In the equation above, price_{it} is the reported price for market i during time t, where market i is either one of the four balancing zones (North $[N]$, Houston $[H]$, South $[S]$, West $[W]$) or one of the markets for ancillary services (Regulation-Up $[RU]$, Regulation-Down $[RD]$, Responsive Reserves $[RS]$, Non-Spinning Reserves $[NS]$).[3] Price is a function of operable wind power capacity (wind_t), the previous business day's price for natural gas (NGprice_t), total hourly demand for electricity in the interconnection (load_t), and total output from nuclear generators in the interconnection (nuclear_t). The estimations also include additional month by year fixed effects and hourly fixed effects.

[3] Other analyses of the determinants of electricity prices use either the log of price, the daily average price, or the log of the daily average price. None of these specifications were considered in this paper because of the suspected relationship wind power has with the occurrence of negative prices in ERCOT's balancing markets.

It is expected that wind has a negative effect on balancing market prices ($\beta_i < 0, i \in \{N, H, S, W\}$), though the size of this effect may not be constant across zones. Because wind power is concentrated in the West zone, it is likely that the effect is largest there ($\beta_W \leq \beta_i, i \in \{N, H, S\}$). At the same time, wind power may have a positive effect on prices for ancillary services such as Regulation-Up and Regulation-Down because its intermittency may require greater reliance on these services as wind power becomes more prevalent in the sampled timeframe ($\beta_i \geq 0, i \in \{RU, RD\}$). Prices for Responsive Reserves and Non-Spinning Reserves are likely non-decreasing in wind power capacity ($\beta_i \geq 0, i \in \{RS, NS\}$), since wind power's intermittency may require additional backup capacity to insure system reliability, but the acquisition of these services is typically left to the discretion of the system operator. Alternatively, lower balancing prices would correspond to high-cost generators being displaced by wind power, and these could then be available to provide ancillary services. Thus, it may be the case that wind power actually decrease ancillary service prices ($\beta_i \leq 0, i \in \{RU, RD, RS, NS\}$). But this is conditional on two other results: wind power must decrease balancing prices and it must have little to no effect on price volatility.

Generally, natural gas prices should have a positive effect on all prices ($\gamma > 0$) because natural gas is the most common fuel source in the system and it is used for both base-load and peaking units. Load levels should likewise have a positive effect on balancing prices ($\delta_i > 0, i \in \{N, H, S, W\}$), since periods of high demand will require a greater number of high-cost generators to switch on. Nuclear generation, which represents the lowest-cost dispatchable generation in the system, is expected to have a negative effect on balancing market prices ($\theta_i < 0, i \in \{N, H, S, W\}$) because these units are typically run whenever available, and when nuclear generators are down a greater number of higher-cost generators will have to be run to make up for the deficit in capacity.

4.B. Effects on Price Volatility

While previous research has concentrated on the effect wind power has on average balancing market prices, many of these analyses have also noted that wind power may increase price volatility. Yet there is no standard definition of "price volatility," so previous methodologies have varied. Some studies have presented a general overview of power prices and their tendency to exceed certain thresholds (e.g., Jacobsen and Zvingilaite, 2010, and Cutler et al., 2011). Another common method is to use variance as a substitute for volatility (e.g., Chang et al., 2009, Woo et al., 2011b, and Clò et al. 2015). Yet variance alone does not fully describe the nature of price changes in electricity markets. Power market prices can rapidly spike up- or downward, hence Ketterer's (2014) utilization of a GARCH model.

Figure 4 highlights the importance of defining volatility. The leftmost figure is a histogram of all observed prices for a simulation lasting 96 periods. Jónsson et al.'s (2010) analyses of wind power and power market conditions relies on an analysis of the basic distributional properties of the leftmost figure: mean, variance, skewness, and kurtosis of daily prices. At the same time, using daily distributional properties of prices presents an incomplete picture, since multiple random walks can create the same daily distribution. The center and rightmost graphs show two very different random walks for price that both form the same distribution on the left. Clearly, Price Walk 1 is more "volatile" in the sense that its price changes are more erratic, yet Jónsson et al.'s analysis (or any paper that relies on the daily mean price) cannot differentiate between these price walks. The unpredictable nature of price changes is obviously important when describing volatility, hence most researchers' reliance on the variance of price changes to as the measure of volatility. Even so, there is additional value in considering daily distributional properties of observed prices, since these properties can indicate new price behavior such as the increased frequency of sustained extreme prices, whereas price changes alone cannot.

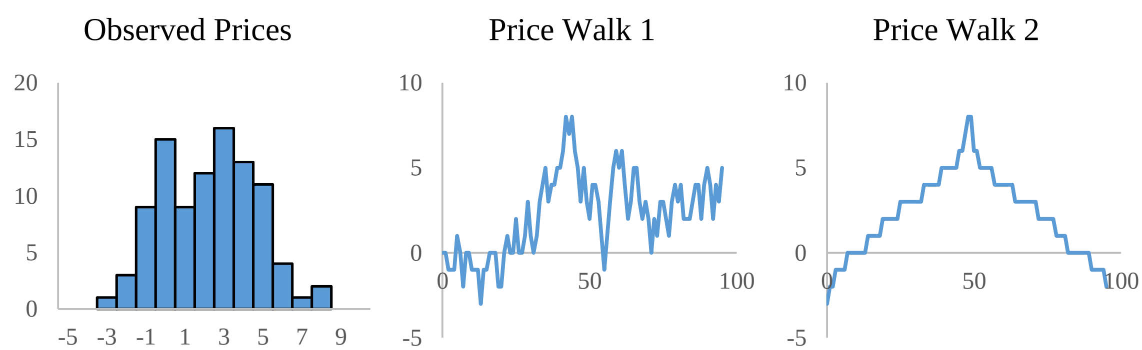

Figure 4. *Price Volatility Components*

In order to better characterize volatility, this paper uses two measures each for variance, skewness, and kurtosis of prices. The first measure (measure A) is calculated using all of the observed prices within a given day. This measure was chosen because the explanatory variables are expected to influence the general shape of price distributions and the occurrence of extreme prices, and is similar to the analysis by Jónsson et al. (2010). Since wind power is also expected to influence the erratic and unpredictable nature of changing prices, a second set of variables (measure B) was created by taking the first difference of all sequential prices in a day and calculating the daily variance, skewness, and kurtosis of these differences. The effect of wind power on price volatility is thus measured according to the specifications in the equation

(2) $\sigma^{z}_{iXt} = \alpha_{ziX} + \beta_{ziX}\text{wind}_t + \gamma_{ziX}\text{NGprice}_t + \delta_{ziX}\text{load}_{zXt} + \theta_{ziX}\text{nuclear}_t + \varepsilon_{it}$.

In the equation above, the zth moment ($z \in \{2,3,4\}$) using measure X ($X \in \{A, B\}$) of zone i during day t is a function of the same variables used to estimate price_{it}, with the exception of load levels. load_{zXt} is calculated as the value of moment z from the distribution of day t according to measure X.

If wind power affects the volatility of balancing prices, it should be the case that β_{ziX} is positive for both measures of variance and kurtosis. If $\beta_{ziA} > 0$ for variance and kurtosis, this will indicate that additional wind power capacity expands the distribution of daily prices, where $\beta_{2iA} > 0$ indicates that prices become more varied about in the mean in general, and $\beta_{4iA} > 0$ indicates an increase in the frequency and persistence of extremely high and low prices. At the same time, even if the distribution of daily prices increases, it may be the case that price movements remain small (e.g., extremely low prices are becoming more common as a result of wind power, but their daily appearance is gradual and predictable). If wind power also affects the unpredictable nature of price changes, it should also be the case that $\beta_{ziB} > 0$ for variance and kurtosis, where β_{2iB} would be comparable to the traditional measure of "volatility" in many analyses (e.g., Chang et al, 2009 and Woo et al., 2011b), and β_{4iB} would describe the tendency of sudden price spikes.

5. Data Description

5.A. Prices

Balancing and ancillary service prices were obtained from ERCOT. Prices for the balancing markets are reported in 15-minute intervals (a standard day has 96 observations) and ancillary service prices are reported at the hourly level. The data span from April 16, 2003 to November 30, 2010. Beginning December 1, 2010, ERCOT switched the balancing market setup from zonal to nodal, effectively creating new data that are not directly comparable to earlier prices. Hence an analysis wishing to utilize more recent data would either have to ignore previous years (and ignore the years with the greatest amount of growth in wind power) or make unrealistic assumptions to simplify and compare both zonal and nodal prices. Similarly, earlier price data were available, but ERCOT revised its market rules in 2003 to influence wind power producers' scheduling behavior (Sioshani and Hurlbut, 2010). It is not clear that the estimated marginal effect of wind power capacity from earlier observations would be the same as later observations, hence earlier data are omitted. Summary statistics for the market-clearing prices are presented in Table 1.

Table 1. *Summary Statistics* – Power Prices ($/MW)

	Count	Mean	Median	Std. Dev.	Min.	Max.
Balancing Markets						
North	267412	47.98	40.30	59.47	-1000.00	2382.50
Houston	267412	49.78	40.68	72.08	-1536.30	3805.70
South	267412	48.84	40.11	79.42	-2292.80	4514.70
West	267412	45.42	39.68	61.10	-1981.80	2320.70
Ancillary Services						
Reg.-Up	66853	13.93	9.66	15.17	0.01	500.03
Reg.-Down	66853	11.43	8.00	14.76	0.01	700.00
Responsive Res.	66853	13.24	8.97	19.18	0	2000.00
Non-Spinning	66853	3.55	0.00	17.34	0	2000.00

Mean balancing prices tend to be between 45 and 50 dollars per megawatt for all four zones, though the median values demonstrate that prices are skewed. Average prices are also somewhat lower in the West, possibly because of the strong presence of wind power in this area and its negative effect on prices. Prices are also slightly higher in the Houston zone, where much of Texas's industry is based. For ancillary services, Regulation-Up, Regulation-Down and Responsive Reserves have similar distributions, though Non-Spinning is considerably different than the other three ancillary services.

Whereas balancing prices are occasionally negative due to imbalances between supply and demand, the minimum values for all ancillary service prices are close to zero and strictly non-negative, indicating that ancillary service prices may need to be treated as left-censored data. Nevertheless, the minimum value is only observed once for Regulation-Up, twice for Regulation-Down, and four times for Responsive Reserves out of more than sixty thousand observations, indicating that accounting for censoring would not significantly improve results, if at all. This is not the case for Non-Spinning prices, whose median price is zero. For the analysis of Non-Spinning prices, OLS estimations are replaced with a tobit model and quantile regressions are only run above the censoring point.[4]

Summary statistics for the two measures of price volatility are presented in Table 2. Both measures of variance and kurtosis indicate that all four zones are subject to considerable variation in day-to-day prices. Yet the reported median statistics, as well as minimum and maximum values indicate that mean levels are very skewed because of extreme values, more so than observed prices. Mean and median statistics for skewness tend to be positive, though the full range includes both negative and positive values. As a consequence, interpreting coefficients for wind power will not be as straightforward in the case of skewness, since both positive and negative values of the coefficients specified in Eq. (2) can indicate an increase or decrease in skewness.

[4] Standard quantile regressions still produce meaningful estimates for quantiles which are above the censoring point (Angrist and Pischke, 2009), though an alternative approach would be the censored quantile regression developed by Chernozhukov and Hong (2002). A similar effect was found using either a quantile regression or a censored quantile regression.

Table 2. *Summary Statistics* – Price Volatility

	Count	Mean	Median	Std. Dev.	Min.	Max.
Observed Prices						
North						
σ_A^2	2785	2991.25	297.76	17328.35	1.94	467000
σ_A^3	2785	1.51	0.78	2.46	-9.38	9.64
σ_A^4	2785	12.54	4.55	18.5	1.21	93.97
Houston						
σ_A^2	2785	4378.71	316.32	28796.56	1.94	959000
σ_A^3	2785	1.63	0.87	2.56	-9.47	9.64
σ_A^4	2785	13.36	4.68	19.88	1.21	93.96
South						
σ_A^2	2785	5350.48	298.94	45893.58	1.94	1530000
σ_A^3	2785	1.52	0.77	2.62	-9.52	9.64
σ_A^4	2785	13.37	4.52	19.94	1.21	93.95
West						
σ_A^2	2785	3102.62	380.28	16853.84	1.94	467000
σ_A^3	2785	1.25	0.65	2.49	-9.59	9.58
σ_A^4	2785	12.00	4.54	17.41	1.05	93.30
Price Changes						
North						
σ_B^2	2785	2902.76	79.81	16217.29	0.19	313000
σ_B^3	2785	0.40	0.38	1.65	-9.63	9.57
σ_B^4	2785	18.82	13.04	14.83	3.14	93.77
Houston						
σ_B^2	2785	3952.58	86.25	19410.83	0.19	369000
σ_B^3	2785	0.36	0.34	1.73	-9.63	9.64
σ_B^4	2785	19.34	13.35	15.45	3.11	93.90
South						
σ_B^2	2785	4370.51	82.52	24848.25	0.19	545000
σ_B^3	2785	0.38	0.36	1.75	-9.62	9.63
σ_B^4	2785	19.25	13.09	15.63	2.91	93.83
West						
σ_B^2	2785	2626.66	111.36	14522.54	0.19	306000
σ_B^3	2785	0.27	0.26	1.70	-9.52	9.56
σ_B^4	2785	19.20	14.20	14.35	3.18	92.93

5.B. Explanatory Variables

Additional data relevant to wind power capacity, natural gas prices, load levels, and output from nuclear generators were collected from several sources. In order to measure wind power capacity in the Texas interconnection, multiple datasets from the Energy Information Administration [EIA] had to be combined to track capacity levels in ERCOT. First, wind farms that were within Texas but not the Texas interconnection had to be identified and discarded, since some generators in Texas are part of either the Western or Eastern interconnection and should be excluded from the analysis.[5] Relevant wind farms were identified using the datasets EIA-906, EIA-920, and EIA-923. These files report general information on all power producers in the US, including whether a Texas-based power plant is located within the Texas interconnection or another power system. Unfortunately, these files do not include information on power plant capacity, hence they were only used for initial identification purposes. EIA-860 reports capacity at the generator-level (a conventional power plant usually consists of multiple generators), as well as each generator's operability status, its initial month and year of activation, the month and year of its retirement (when applicable), and the state the generator is located in. Each year's release was cross-checked with previous years' datasets to ensure consistency of the data, as well as to identify changes in active generators.[6]

Data on the remaining explanatory variables are from ERCOT, the EIA, and the Nuclear Regulatory Commission. Hourly load levels for electricity use in the Texas interconnection come from ERCOT.[7] The natural gas prices used in the analysis are from

[5] Approximate ERCOT borders are depicted in Figure 1.
[6] Two instances were discovered in which wind turbines were lost due to extreme weather events.
[7] While load corresponds to quantity demanded, it is treated as exogenous with respect to price. Although this may be somewhat counterintuitive, since demand for most goods is influenced by price, most

the Henry Hub spot price, as reported by the EIA. Daily generator output for nuclear generators comes from the Nuclear Regulatory Commission, which releases periodic reports of all nuclear generators in the US, including their status and utilization rates. Descriptive statistics for all explanatory variables are reported in Table 3.

Table 3. *Summary Statistics* – Explanatory Variables[8]

	Count	Mean	Median	Std. Dev.	Min.	Max.
Wind Capacity (GW)	92	3.97	2.70	2.96	1.02	9.18
NG Price ($/MMBtu)	1904	6.39	6.15	2.22	1.83	15.39
Hourly Load (GW)	66853	34.99	32.98	8.36	19.66	107.08
Nuclear Generation (GW)	2786	4.72	5.14	0.81	1.35	5.14

6. Results and Discussion

6.A. Balancing Prices

Table 4 presents findings from OLS and select quantile regressions when the dependent variable is the balancing price for each of the four ERCOT zones.[9] The second column includes regression results from OLS for each balancing market, the third column has results from quantile regressions for the twenty-fifth quantile, the fourth column has

consumers do not buy power directly from wholesale power markets. Some industrial consumers do participate in ERCOT's wholesale markets, though findings from Zarnikau and Hallett (2008) indicate that demand's impact on market prices is negligible.

[8] Explanatory variables often have fewer observations than dependent variables because they are recorded at different frequencies (e.g., wind capacity varies month to month, but power prices are reported every fifteen minutes or one hour). In the estimations, observations for explanatory variables are extended so that they match the dependent variables. Summary statistics in Table 3 report data in their original frequencies for illustrative purposes.

[9] Results without month by year fixed effects and hourly fixed effects are provided in Appendix A. Generally, estimated coefficients for wind power were the most sensitive variable to month by year fixed effects. This is likely because wind power capacity is typically increasing from month to month, and without fixed effects to capture unspecified trends, wind would pick up those effects as well. Hourly fixed effects did not influence results much.

results for the median quantile regression, and the fifth column has results for the quantile regression of the seventy-fifth quantile.

Table 4. *Regression Results* – Dependent Variables are Balancing Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
Wind	-4.178**	1.581**+	-0.172+	-2.912**+
NGprice	5.095**	4.286**+	5.079**	5.845**+
Load	3.307**	1.738**+	1.845**+	2.128**+
Nuclear	-3.852**	-2.417**+	-2.189**+	-2.489**+
Intercept	-59.933**	-35.308**+	-33.159**+	-34.796**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
Houston				
Wind	-4.593**	-1.823**+	-3.085**+	-4.829**
NGprice	4.386**	4.226**+	4.970**+	5.574**+
Load	3.760**	1.773**+	1.874**+	2.175**+
Nuclear	-5.418**	-2.649**+	-2.471**+	-2.747**+
Intercept	-62.494**	-31.779**+	-29.959**+	-32.281**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
South				
Wind	-0.215	-1.256**+	-2.256**+	-3.475**+
NGprice	3.838**	4.130**+	4.875**+	5.587**+
Load	3.781**	1.719**+	1.806**+	2.078**+
Nuclear	-5.653**	-2.654**+	-2.574**+	-2.749**+
Intercept	-63.769**	-31.356**+	-28.647**+	-31.305**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
West				
Wind	-7.987**	-0.639+	-0.922**+	-2.913**+
NGprice	4.927**	4.247**+	5.018**	5.795**+
Load	3.103**	1.756**+	1.836**+	2.105**+
Nuclear	-4.988**	-2.827**+	-2.189**+	-2.508**+
Intercept	-46.688**	-32.772**+	-32.091**+	-33.989**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES

*Significant at the 5% level; **Significant at the 1% level
+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

OLS and quantile coefficients can be interpreted in the following manner. According to the OLS model in the second column, the average marginal effect of 1 Gigawatt [GW] of wind power capacity on balancing prices in the North zone is -4.178. That is, if an additional 1 GW of wind power capacity were added to the system, the mean power price in the North zone would decrease by about $4.18. Results from the quantile regressions illustrate that this effect is not constant for all quantiles. For example, the average marginal effect of an additional 1 GW of wind power capacity would only decrease the median price in the North zone by about $0.17, and the effect is not statistically different than zero.

Generally, the estimated effects for all explanatory variables aside from wind power are in line with expectations and previous findings in the literature. Natural gas prices have a strong significant effect on power prices, consistent with the fact that the majority of power in Texas is generated by natural gas-fired plants. They tend to have an average marginal effect on balancing prices between $3.80 and $5, depending on the zone. Natural gas's effect on balancing prices is also larger for higher quantiles, which suggests that natural gas generators play a stronger role in price formation during peak prices. Quantile regressions report that for each zone, the average marginal effect on the seventy-fifth quantile is about $1.50 higher than the marginal effect on the twenty-fifth quantile.

Load levels also have a positive and statistically significant effect on power prices, confirming that times of high demand require more high-cost generators to switch on and produce power. For each of the four zones, the average marginal effect falls between the range $3.10 and $3.80. Interestingly, quantile coefficients for load seem to be relatively flat across the selected quantiles, yet OLS estimates are considerably higher for all four zones. This suggests that the typical marginal effect of load on power prices is between $1.70 and $2.20/GW, but load has a considerably larger effect during peak prices, which substantially increases the mean effect.

The marginal effect of nuclear generation on balancing prices is similar, but not identical, to the absolute value of load's marginal effect. This relationship was suspected because an additional GW of power from a nuclear generation would be very low cost, hence it could be as effective at lowering balancing prices as decreasing load levels by the same amount, since either scenario corresponds to a high-cost generator decreasing its output. Yet the average marginal effect across all four zones is larger and shows more variation than the marginal effect of load. Nuclear generation decreases balancing prices by between $3.70 and $5.70 according to OLS results. However, reported quantile coefficients show two trends worthy of additional consideration. First, average marginal effects for all zones appear relatively flat across quantiles, and the OLS coefficient does not appear to intersect with the quantile regressions' coefficients between the selected range, suggesting nuclear generation may have a larger effect at extreme price levels that accounts for the discrepancy.[10] Furthermore, coefficients for nuclear generation and load tend to be much closer in absolute value for reported quantiles, even though the OLS coefficients for the two variables are considerably different. This suggests that load and nuclear generation usually have a nearly equal but opposite relationship, except that

[10] Additional quantile regressions were estimated up to the ninety-fifth quantile in increments of five to verify the OLS coefficient would only intersect quantile coefficients at extremely high quantiles, and that quantile regression coefficients were relatively flat between the twenty-fifth and seventy-fifth quantiles.

extreme quantiles and the nature of load variability may obscure this on when calculating mean effects.

Most importantly, the information in Table 4 confirms that wind power has a negative effect on balancing prices and highlights a relative disparity in the marginal effects of wind power on balancing prices across the four zones. OLS results find that wind power has the largest effect on prices in the West zone, which is expected because the overwhelming majority of wind power capacity is located in the West zone. There, an additional GW of capacity decreases West balancing prices by $7.79 on average. Estimated coefficients are similar in sign and statistical significance in the North and Houston zones, though smaller in size. The marginal effect of wind power capacity is -$4.18 on North balancing prices and -$4.59 on Houston prices. In the South zone, balancing prices are barely affected by wind power. There, an additional GW of wind power capacity decreases price by less than $0.22, although the effect is not statistically different from zero.

To further explore the effect wind power has on balancing prices, Figure 5 shows the estimated values of β for balancing prices in all four zones using both OLS and quantile regressions results. The dashed and dotted lines depict estimated coefficients from OLS regressions and their 95% confidence intervals, respectively, which are all flat because the standard regression technique assumes that explanatory variables have a constant marginal effect across all quantiles. For all four zones, quantile regressions find that the marginal effects are negative, statistically significant, and slowly decreasing in size between the twenty-fifth and seventy-fifth quantiles. Quantile regression results are also quite similar throughout that range, even though OLS results report substantially different mean marginal effects.

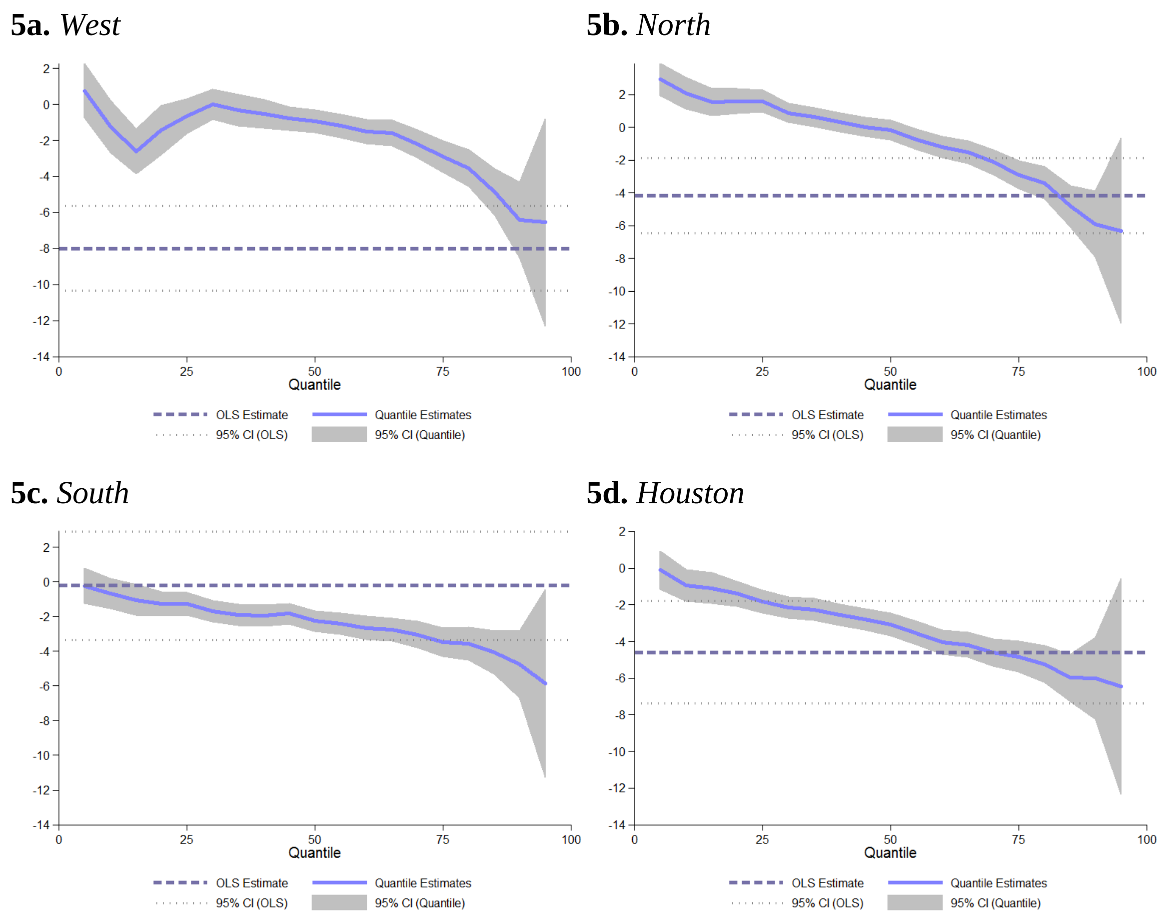

Figure 5. *Marginal Effects of Wind Power across Quantiles*

The difference between OLS coefficients across zones, as well as the difference between OLS and quantile regression coefficients within zones, demonstrates the value of utilizing quantile regressions in conjunction with OLS. For one, quantile regression results indicate that much of the variation in zonal OLS coefficients is attributable to differences in marginal effects at higher or lower quantiles. That is, wind power may normally have an identical effect on all zonal prices, but differences in marginal effects for either high or low prices causes the disparity between OLS estimates. Furthermore, OLS coefficients for several zones appear to be heavily influenced by marginal effects at extreme quantiles. Consequently, relying only on OLS would overstate the variation of average effects of wind power for most prices, and results would not be informative as to the typical impact of wind power on power market prices.

For lower quantiles, quantile regressions highlight a difference in marginal effects across the four zones. In the West zone, wind has a much smaller effect on lower quantiles. This result is potentially inconsistent with the information from Figure 2, and contrasts Woo et al.'s (2011a) finding that wind power is responsible for the price divergence and extremely low prices seen in the West zone. One possible explanation for this discrepancy is that the estimated coefficients are the average marginal effect, and the West balancing prices faced consequences from congestion early on in the analysis's timeframe. If a relatively small amount of wind power already congested transmission lines and effected the minimum price possible, additional units of wind power capacity would have had little if any effect on prices in the lower quantiles. The positive effect on prices in the North zone may be a further consequence of this effect, as congested lines would make it difficult to consistently transmit generation from the lowest-cost producers. Quantile regressions also demonstrate that wind power has no significant marginal effect on Houston and South zones' balancing prices in the lower quantiles.

The signs of these coefficients are similar to those in Woo et al. (2011b) though the exact magnitude of the marginal effects sometimes varies.[11] Woo et al. found that an

---

[11] Woo et al. (2011b) have the same explanatory variables (wind, natural gas price, load, and nuclear output) but their base units are slightly different. For one, their load data pertains to each of the four balancing zones, as opposed to the weather zones ERCOT typically reports load levels in. These were converted for present purposes by calculating a marginal effect of total load for each zone using a weighted

additional 1 GW of nuclear power in the system would decrease power prices by $5 to $7, whereas here it ranges from $3.70 to $5.70. Load and natural gas prices were also found to have a positive effect on prices in Woo et al.'s analysis; an additional 1 GW of load would increase prices by $5 to $6, and a $1/MMBtu increase in the price of natural gas would increase the price of power by $1 to $2. OLS estimates for wind power, on the other hand, are relatively close. According to Woo et al., and an additional 1 GW of wind power capacity in ERCOT would decrease prices by $1.525 in North, $0.975 in Houston, $0.80 in South, and $3.825 in West. It is not known why the estimated coefficients are different from those found here, though Woo et al.'s timeframe is considerably shorter (January 2007 – May 2010) than the analysis here (April 2003 – November 2010), and the present analysis sees considerably more variation in natural gas prices and wind power penetration levels.

6.B. Price Volatility

With price volatility divided into three different moments (variance, skewness, and kurtosis) using two general measures (measures A and B) for four balancing zones, there are twenty-four unique sets of results—not including separate estimation results for OLS and quantile regressions. Accordingly, the results here are restricted to the relationship between wind power and price volatility. Table 5 summarizes the estimations pertaining to the relationship between wind power and the distributional properties of daily prices (measure A). These results are excerpts from each estimation, and only

---

average. Woo et al. also use wind power output (measured in watt hours), instead of installed wind power capacity (measured in watts). The two are compared here assuming an average capacity factor of 25%. Lastly, Woo et al.'s variables are reported at the megawatt-level and were converted to correspond to marginal effects at the gigawatt-level.

display the estimated marginal effect of wind power (β) on each measure of price volatility. Although not included in the table for legibility, natural gas prices have little effect on the variance measure σ_A^2 or kurtosis σ_A^4 for all four zones, though they do appear to increase leftward skewness. Load also has a strong, positive effect on price variance and skewness for all four zones and not much effect on kurtosis. Nuclear generation has no effect on the variances, skewnesses, or kurtoses.

Table 5. *Wind and Price Volatility* (β_i when Dependent Variables are Volatility Measure A)

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
σ_A^2	-1462.731	-67.094+	-160.177+	-1125.490
σ_A^3	0.004	-0.093	-0.738	-0.598
σ_A^4	-1.328	-0.607	-2.775	-4.341
Houston				
σ_A^2	2293.156	-87.522+	-139.217+	-927.191
σ_A^3	-0.475	-0.205	-1.234	-0.420
σ_A^4	3.623	-0.431+	-3.488	-6.201
South				
σ_A^2	10637.106	-81.137+	-158.467+	-998.910+
σ_A^3	-0.580	-0.219	-1.160	-0.561
σ_A^4	-3.641	-0.424+	-2.228	-5.200
West				
σ_A^2	-2444.870	-32.489+	-66.345+	-1438.446
σ_A^3	0.322	-0.045	-0.324	0.213
σ_A^4	-3.636	-0.848+	-1.730	-1.984

*Significant at the 5% level; **Significant at the 1% level
+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Using the variance from measure A for price volatility, estimations find that wind power has a statistically insignificant effect on the variance of prices in all four zones. This size of the marginal effect varies from zone to zone, being quite large in the Houston and South zones and comparably smaller in the North and West zones. Select quantile regression results similarly find no statistically significant effect, though these coefficients are considerably closer to zero and confidence intervals are relatively small for all four zones. This indicates that for most quantiles, wind power likely has no effect on σ_A^2, and at extreme quantiles, wind power may have an effect on the variance of balancing prices, but the hypothesis that they have no effect cannot be rejected.

The skewness of prices' daily distributions were likewise unaffected by wind power according to OLS estimations. Quantile regressions find a positive effect for several of the lower quantiles in the North, Houston, and South zones, confirming that wind power does affect the skewness of daily prices in these other zones. In this case, since the estimated marginal effect of wind power on σ_A^3 is positive for smaller quantiles, and daily skewness for these zones is negative at smaller quantiles, wind power actually decreases the skewness of prices for these quantiles.

The kurtosis of daily prices is also only marginally affected by wind power, and initial OLS estimations find that wind power has no statistically significant effect. However, at lower quantiles, wind power does appear to have a small but statistically significant effect on kurtosis in the North, Houston and South zones. None of the four zones show a significant effect on kurtosis in higher quantiles.

These results actually somewhat conflict with the findings of Jónsson et al. (2010), who study the market impacts of wind power on power prices in western Denmark. Whereas here, wind power has no strong effect on the volatility of daily observed prices, Jónsson et al. find that when a power system's penetration level of wind power grew from 0% to 10%, variance, skewness, and kurtosis tend to decrease. The findings here represent an increase in wind power penetration in Texas for an almost identical range, yet no significant change in variance, skewness of kurtosis could be strongly attributed to changes in wind power levels.

The effects of wind power on the distributional properties of prices changes (measure B) are summarized in Table 6. OLS estimations find no evidence of a change in any of the moments that can be attributed to wind power with reasonable statistical certainty. At smaller quantiles, β is relatively small and is not statistically different than zero. At higher quantiles, estimated coefficients are sometimes statistically significant for skewness, kurtosis, or both. But the effect of wind power on the daily variance of price changes is surprisingly absent in all four zones' estimations. This suggests that, contrary to Woo et al. (2011b), the hypothesis that wind power has no impact on price volatility cannot be rejected.

Table 6. *Wind and Price Volatility* (β_i when Dependent Variables are Volatility Measure B)

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
σ_B^2	-543.921	-17.092+	-40.021+	-474.643
σ_B^3	0.706*	0.553	0.409	0.685*
σ_B^4	-0.937	-1.793	-0.122	4.135
Houston				
σ_B^2	-433.973	-15.420+	-85.103+	-346.736
σ_B^3	0.682	0.651*	0.336	0.468
σ_B^4	-2.959	-3.286	-1.039	-5.094
South				
σ_B^2	4817.846	-13.142+	-39.324+	-528.891
σ_B^3	0.711	0.358	0.224+	0.376
σ_B^4	-3.364	-3.487	-1.315	-4.652
West				
σ_B^2	-1424.045	-33.243+	-163.642+	-746.848
σ_B^3	-0.251	-0.297	0.050	0.092
σ_B^4	0.773	-0.100	-0.787	3.457

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

6.C. Ancillary Services

Unlike the estimations for price volatility, there is an apparent effect of wind power on several of the markets for ancillary services, as demonstrated in Table 7. However, the effect is negative and generally statistically significant, indicating that an increased presence of wind power capacity actually led to cheaper prices for ancillary services. On average, an additional GW of wind power capacity decreased prices for Regulation-Up by $3.28/MW and Responsive Reserves by $4.03/MW. Estimations for Regulation-Down also find a negative marginal effect, by about $0.53/MW, though the effect is not statistically significant. Because of censoring the reported coefficient wind power capacity on Non-Spinning Reserve prices is from a tobit model, and corresponds to an average partial effect of -$0.77/MW.

Table 7. *Regression Results* – Dependent Variables are Ancillary Service Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
Reg.-Up				
Wind	-3.280**	-0.691**+	-1.952**+	-3.356**
NGprice	0.222*	0.004+	-0.102+	-0.057**+
Load	1.025**	0.385**+	0.629**+	0.812**+
Nuclear	-1.570**	-0.631**+	-0.876**+	-1.213**
Intercept	-9.216**	-1.563**+	-2.328**+	-0.633**+
Month*YR FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
Reg.-Down				
Wind	-0.534	-0.873**+	-0.109	-0.067
NGprice	0.454**	0.166**+	0.354**+	0.679**+
Load	-0.065**	0.003+	0.010*+	0.014+
Nuclear	0.418**	0.004+	-0.025+	0.276*
Intercept	3.926**	3.267**	2.481**+	3.336**
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
Resp. Res.				
Wind	-4.029**	-0.788**+	-2.212**+	-4.041**
NGprice	1.833**	0.014+	0.238**+	0.687**+
Load	1.051**	0.390**+	0.628**+	0.817**+
Nuclear	-1.540**	-0.881**+	-1.107**+	-1.339**
Intercept	-20.062	-2.630**+	-4.943**+	-5.125**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	YES	YES	YES	YES
Non-Spinning				
Wind	-2.431*			0.061
NGprice	2.188**			0.014
Load	1.745**			0.420**
Nuclear	-1.262**			-0.194
Intercept	-63.694**			-5.838**
Month*Yr. FE	YES			YES
Hour FE	YES			YES

*Significant at the 5% level; **Significant at the 1% level
+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Additional results for other variables are largely consistent with expectations as well as findings from results from the balancing markets. Higher natural gas prices lead to higher prices for all ancillary services, likely because a majority of ancillary services are fulfilled by gas-fired generators. Higher load levels also increase prices for Regulation-Up, Responsive Reserves, and Non-Spinning Reserves, though they decrease prices for Regulation-Down, probably because fewer generators would need to decrease their output when demand is high. Lastly, nuclear generation tends to have an opposite effect on ancillary service prices as load, indicating that when the nuclear plants are online fewer ancillary services are necessary as more base-load is being satisfied by low-cost generators.

7. Conclusion

Due to the intermittency of wind power output, there is a general concern that a high penetration level of wind power will negatively affect the reliability of the overall power supply, and that this will be realized through an increase in price volatility, a greater need for ancillary services, or both. This paper presents an analysis of wind power's observed market impacts in the Texas interconnection from 2003 to 2010, during which time the power system experienced considerable growth in installed wind power capacity. The estimations account for growth in wind power while also controlling for changes in hourly load levels, daily fluctuations in natural gas prices, and changes in output from nuclear generators. Quantile regressions, whose estimations are more descriptive, are used in conjunction with OLS to quantify the effects on balancing market prices, balancing price volatility, and prices for ancillary services.

The majority of wind power capacity in Texas is located in the western part of the state and, not surprisingly, the largest effect on balancing market prices occurs in the

West zone. The average marginal effect of wind power on all zones' balancing prices tends to be negative and statistically significant, indicating that wind power's effects are not confined to one zone. However, using multiple metrics for price volatility, estimations failed to find a change in volatility attributable to wind power. Furthermore, while estimations did find that wind power had an effect on prices for ancillary services, the effect was negative and statistically significant, contrary to what previous analyses suggested would be the consequences of an increased presence of wind power.

It is noteworthy that the analysis found any effect on power market conditions outside of the balancing markets. Many studies on the market impacts of wind power limit themselves in scope to one power market, typically the balancing market or its equivalent. Some extend their analyses to incorporate ancillary services or price volatility, but none present a comprehensive review of all power market characteristics. This paper analyzes the potential impacts of wind power by considering all three aspects: balancing prices, volatility of balancing prices, and prices for ancillary services; and demonstrates that installing a large amount of wind power capacity can potentially affect all three. Whereas many recent studies omit ancillary services and focus instead on price volatility, in this case prices for ancillary services decreased from a rise in wind power but there were no clear impacts on price volatility. This shows that the effects of wind power's intermittency are not limited to price volatility, so future studies should not restrict themselves to price volatility when trying to assess the consequences of wind power's intermittent nature. Furthermore, the severity of wind power's intermittency appears to be overstated, as demonstrated by this paper's combined results of no impact on price volatility and a decrease in the prices for ancillary services.

Changes in price volatility and prices for ancillary services will have both short and long run effects on power system conditions. In the short run, increases in either price volatility or the need for ancillary services would require conventional generators to run at technically inefficient levels more often. Changes in price volatility and ancillary services will also affect the long term profitability of generators depending on the flexibility of their output. If unaccounted for in cost-benefit analyses of wind power, these short and long term changes may also lead to overstating the potential ability of wind to abate emissions, both because of running generators at inefficient levels (Denny and O'Malley, 2006; Katzenstein and Apt, 2009) and because a heterogeneous change in the profitability of conventional generator types will affect the planning of future capacity. In this case, studies on the environmental and economic impacts of wind power may not need to account for wind power's intermittency in the short run, as evidenced by a decrease in ancillary service prices and no clear effect on price volatility. However, studies on the long term impacts of wind power in Texas should acknowledge that wind power will not only decrease power prices, but that it will also affect the profitability of generators well-suited to providing ancillary services relative to those generators which are not.

Although the power system in Texas is not unique, it should be emphasized that a discrepancy in estimated marginal effects between this study and others by no means negates previous findings. Many power systems' markets are structured differently, penetration levels of conventional generators vary greatly, and even geographic differences in wind conditions could influence results. Regardless, some of the findings from Texas can help guide future research in several ways. First, quantile regressions proved to be a natural addition to the empirical analysis. OLS occasionally found no statistically significant effect, and quantile regressions could explain this as either a non-monotonic effect that varied across quantiles, or a marginal effect that might truly be zero. Second, wind power does have a predominantly negative effect of power prices for the real-time market, but if a wind farm is connected to multiple markets, the effects may not be identical everywhere. Researchers should therefore endeavor to incorporate all markets that a wind farm is realistically connected to in their analyses. Lastly, the impacts of an intermittent renewable energy may manifest themselves differently depending on the system and circumstances. In this instance, wind power increased from approximately 1% to 9.5% of total capacity, and price volatility was not conclusively impacted but ancillary services were. It is not known if comparable results would have been found for similar growth elsewhere (e.g., the power system in western Denmark) or for the same system but with a different amount of growth (e.g., if wind grew from 10% to 20%). In general, researchers will not know how wind power's intermittency will affect power market conditions a priori, so future analyses must necessarily incorporate all the aforementioned characteristics of power market conditions.

Abatement from Wind Power: Do Market Conditions Matter?

1. Introduction

The technology to generate electricity from wind has existed for decades, though it has not experienced widespread application until recently. Consequently, the effects of installing a large amount of wind power capacity into an existing power system are not well understood. While wind power has no marginal cost of output and no emissions from generation, it is an intermittent renewable energy [IRE]. That is, it generates power stochastically and its output is neither perfectly controllable nor forecastable. Thus, wind power might displace power from conventional fossil fuel-burning sources and lead to lower generation costs and emissions for the power system, but at the possible expense of system reliability and (when power systems schedule generation through power markets) greater price volatility.

Historically, utility companies that provided consumers with power also generated the electricity they delivered. And since short run demand for power is inelastic with respect to price, utilities maximized their profits by deciding which generators to use based on the criteria of cost minimization. Today, many power systems have been deregulated and employ a series of wholesale power markets to coordinate supply. Although many utilities continue to operate power plants and sell power directly to consumers, other power producers either sell power solely at the wholesale-level or provide consumers with electricity despite not owning any generators. As a result, power generation is more effectively determined by profit maximization, though researchers still tend to use cost minimization in its place when modeling power system behavior.

An important example of the application of cost minimization to study power system outcomes is the literature that explores the installation of new wind power capacity. Researchers tend to use cost minimization methods not only to track savings in expenditures on fossil fuels attributable to wind power, but also to track changes in emissions. But if results from cost minimization do not accurately describe real market behavior, this weakens findings from existing studies on the abatement potential of wind power. Because conventional generators are heterogeneous in costs and emissions, estimated abatement could have either an upward- or downward-bias if results derived from cost minimization are not realistic. The purpose of this paper is therefore to assess the behavior of a profit-maximizing producer faced with volatile prices that change faster than producers can adjust their output, and determine whether simulations that use cost minimization might produce biased results because they do not accurately reflect producers' decisions.

The existing methodology is assessed in the following manner. First, I provide an overview of the various cost minimization techniques researchers have employed to estimate abatement from intermittent renewables. I then examine how changes in average prices and price volatility can alter the optimal strategy of an ex ante profit-maximizing producer. This concept is first illustrated using a basic computational model and a simple manipulation of a price distribution, then again with a more advanced computational model that accounts for a wide variety of parameters specific to power producers. Agent-based simulations are also run to determine how producer output, emissions, and profit are all affected by a change in power market conditions attributable to wind power. While results from the basic computational model demonstrate that price volatility has an impact on producer output, emissions, and profit, current results from the advanced computational model suggest that realistic changes to price volatility from wind power produce results that are comparable to those derived using ex post cost minimization.

2. Integrating Wind Power and Abating Emissions

Most power systems still have little or no installed IRE capacity and generate the bulk of their power from conventional sources. Consequently, the system impacts of a high penetration of wind power tend to be derived from simulations rather than observed trends. The use of cost minimization is crucial to most of these studies, as it is the preferred technique for modeling the integration of new capacity in a power system and approximating generation decisions. Arguably, substituting individual producers' behavior with a system-wide, cost minimization objective function is appropriate when results are not sensitive to market conditions. This applies to many technical studies on the integration costs and impacts to system reliability from new IRE capacity, as minimum costs will still be meaningful even if estimates are optimistic.

Chen et al. (2006) is one such example that use cost minimization to analyze the integration of a high level of wind power in an isolated power system. They do so by developing an algorithm to determine optimal real-time dispatching decisions from conventional power plants that best accommodates power from intermittents, while factoring in spinning reserves and production costs. Lund and Münster (2003a,b) provide another example by developing a model to estimate optimal system management strategies when a power system has a high penetration of both intermittent wind power and combined heat and power plants, whose output is less flexible than standard power plants. Lund (2006) later applies this model to estimate optimal combinations of three types of IRE (solar, wind, and wave power) when a power system has a target penetration level for cumulative power from renewable energy. Delarue et al. (2011) later demonstrate how wind power would fit into regional planning when choosing the optimal combination of generator types to minimize either expected production costs or exposure to risk of volatile fossil fuel prices. None of these aforementioned studies attempted to measure either the displacement of power from fossil fuels or abatement of GHG emissions, hence their included estimated benefits of IRE (avoided costs and reduction in risk to volatile fossil fuel prices) are still meaningful. More frequently, however, analyses attempt to identify abated emissions, though the complexity with which they model system-wide generation (and account for temporal fluctuations in demand) varies.

Load Duration Curves [LDCs] take demand for power over a period of time and sequence it into descending load levels. Though LDCs cannot account for the variability between two consecutive load levels, they can describe variations in annual, seasonal, weekday, and diurnal demand. Hence LDCs can be a useful in identifying the best combination of power plant types to satisfy load requirements over long periods of time, e.g. months or years. Although LDCs do not account for spontaneous variability or system security, they have been used to study the integration of renewable energy in a power system and measure environmental benefits over time. Lehtilä and Pirilä (1996) measure the abatement potential of renewables in Finland with a model that included a four-step LDC.[12] Similarly, Cormio et al. (2003) use an eight-step LDC in a regional energy planning model to predict generation schedules given the presence of renewables

[12] A true LDC would be continuous, but they are traditionally discretized into “steps” for ease of calculations.

and environmental considerations. Holttinen and Tuhkanen (2004) study the abatement potential of wind power in Nordic countries with a model that was meant to simulate market conditions, though supply and demand curves are actually based on four-step LDCs at a weekly level. As such, their simulated market conditions cannot describe any changes to intraday volatility. All of these studies include cost-minimizing objective functions to predict generation to meet load requirements, but do not actually reflect upon potential volatility from IRE.

Load Duration Curves are perhaps more realistic than a lump-sum variable for demand, but they are still problematic in that they ignore sequential fluctuations in demand. In reality, it may not be possible to perfectly minimize producer costs with the utilization rates predicted by LDCs. Because generators cannot instantaneously and perfectly adjust output in response to fluctuations in demand, it may be the case that the true cost-minimizing load schedule is not feasible, and that second-best generators must sometimes increase, or "ramp up," their output to meet demand. To account for temporal constraints, additional work on the integration of IRE combines hourly load data from power systems with simulated generation from hypothetical IRE capacity. Researchers then use linear programming to determine the cost-minimizing generation schedule from dispatchable sources for load requirements not met by the IRE. Chen et al. (2006) and Delarue et al. (2011) are two previous examples that predict generation scheduling based on linear programming with hourly load levels and cost minimization. However, these papers do not attempt to quantity abatement benefits from IRE.

There is considerable variation in the literature with respect to how researchers using linear programming as a proxy for power market transactions treat uncertainty of load and IRE output in their models. Load forecast errors are typically ignored, though a spinning-reserve constraint is sometimes incorporated to ensure that the predicted cost-minimizing generation schedule does not affect the system's reliability.[13] Examples of linear programming with a spinning reserve include Liik et al. (2003), Li and Kuri (2005), Voorspools and D'haeseleer (2006), Delarue et al. (2009), and Luickx et al. (2010). Note that the spinning reserve constraint is a product of the modeler's specifications and not necessarily indicative of real market conditions, hence it may frequently miscalculate costs and emissions incurred by units on standby.

While the literature on the integration of wind power frequently discusses its intermittency and forecasting difficulties, not all papers consider ex ante wind forecasts. From the perspective of abatement analysis, this is inherently problematic. Spontaneous output from wind power needs to be accommodated by having other active producers decrease, or "ramp down," their generation, which subsequently lowers emissions. Research by Denny and O'Malley (2006) indicates that determining generation scheduling without accounting for day-ahead wind power forecasts could also produce fundamentally different results than those that do, as forecasts influence unit commitment decisions that determine which generators are on when output from wind power must be accommodated. Consequently, the abatement analyses of Voorspools and D'haeseleer (2006), Ummels et al. (2007), Delarue et al. (2009), and Luickx et al. (2010) may be more robust than those of Liik et al. (2003) and Benitez et al. (2008), as the latter model load and wind power output under ex post conditions.

[13] The spinning-reserve margin effectively mandates a relationship between supply and demand in which supply consistently exceeds demand by the modeler's defined amount.

Hourly load data has allowed cost minimization analyses to account for ramping constraints and fluctuations in demand, though whether power markets will necessarily tend to these solutions remains untested. While hourly load data is more detailed than LDCs, predicting load scheduling via cost minimization is not necessarily reflective of new market conditions effected by IREs. To a certain extent, cost minimization implies a decrease in average power market prices. However, using the cost-minimizing load schedule to model the integration of new IRE capacity will ignore any additional effect from an increase in power market volatility. Even if power producers are risk-neutral, a change in market volatility could affect power producers differently based on either ramp rates or marginal costs. Since quantifying abatement from non-dispatchable IRE depends on the technique used for estimating the displacement of fossil fuels, researchers must understand what implications power market volatility has for optimal producer behavior.

3. Modeling Power Plant Output

The existing literature estimates displacement of power from fossil fuels and GHG abatement assuming that the integration of IRE will incur the cost-minimizing generation schedule but without accounting for power market conditions. A model that tests individual power producers' responses to (new) market conditions is therefore useful in that it will indicate whether the overall system will react as predicted by these analyses. In this case, a computational model is more illustrative than a simple theoretical model for several reasons. First, an active power producer must continuously decide whether to ramp up, maintain, or ramp down its current level of output. Hence output decisions cannot be described by a static model without unrealistic assumptions on demand (e.g., ignoring the variation and unpredictability of power prices). Second, the optimal adjustment to output may be technically infeasible because of slow ramp rates or minimum down-times for generators. In instances where a producer cannot perfectly adjust output as desired, any change in market conditions (in either volatility or mean prices over time) is unlikely to elicit a uniform response across power plants that differ by important parameters such as ramp rates and marginal costs.

Since electricity is not storable in large quantities and fuel use represents the majority of generation costs, agent-based models of power generation have historically been dynamic. One of the earliest applications of dynamic programming was actually the analysis of output decisions by hydroelectric dams with a stochastically refilling reservoir. At that time, however, utility companies were still regulated and did not produce power in competitive power markets. Hence dynamic programming applications were mostly limited to unit commitment problems to minimize utility companies' costs from daily generation, as in Lowery (1996). Richter and Sheblé (2000) later updated the system-wide unit commitment model to account for deregulation and competition among producers by replacing the objective function of system-wide cost minimization with one of profit maximization, but in general cost minimization is still widely favored.

The few examples of studies on unit commitment and output decisions by individual power producers in deregulated power systems include work by Arroyo and Conejo (2000), Garciá-González and Barquín (2000), and Valenzuela and Mazumdar (2001). Arroyo and Conejo's model analyze the optimal output decisions of a thermal generator operating on the spot and spinning-reserve markets, though they assume that prices are known with certainty. Garciá-González and Barquín do study the decisions of power producers when faced with price uncertainty, but their method to impute the distribution of hourly prices makes their method admittedly inadequate for studying instances of high volatility, as is particular interest here. Valenzuela and Mazumdar likewise account for price uncertainty for a profit-maximizing power producer, though their analysis fails to account for ramp rate restrictions.

The agent-based model here is presented in two parts. In 3.A., a basic computational model is presented to demonstrate the importance of marginal costs, ramp rates, and price volatility, as well how these different factors affect producer decisions. Whereas the basic model is stylized and not meant to represent an actual generator, the advanced model in 3.B. incorporates additional parameters specific to power plants so that the computational model can be calibrated to real world examples. Both models emphasize the importance of ramp rates in dictating the decisions of an ex ante profit-maximizing producer, and the advanced model also includes minimum down-times and startup costs so that the analysis can incorporate the producer's on/off decision.

3.A. Basic Model

3.A.1. Basic Model: Description

Since a generator's output cannot be perfectly adjusted, its output is treated as a state variable in the model. Output is expressed as the generator-level utilization rate z over the interval $[0,100]$, where $z = 100$ indicates operating at full capacity. Under the assumption of constant returns to scale, a power producer has a constant marginal cost c.[14] Marginal revenue comes from the spot market price for power P, where the spot

[14] Constant marginal cost is prevalent in the literature of greenhouse gas abatement, though most generators actually have increasing concave marginal costs. Constant marginal cost is used here to illustrate the

market price is a state variable that changes over time according to the distribution $\pi(P'|P)$. In this case, it is assumed that time periods are quite short, e.g., 15-minute intervals.

The producer controls only its ramping decision r. At most, the generator can ramp up or ramp down generation by its maximum ramp rate γ, though the exact choice set is further restricted depending on the generator's current utilization rate. For example, since z is between 0 and 100, it must be the case that $r \in [-\gamma, 0]$ when $z = 100$, because the generator cannot increase output beyond full capacity. Likewise, if $z = 0$, then $r \in [0, \gamma]$, as the generator cannot produce negative output. Formally, the power producer's optimal output decision can be described with the Bellman equation

$$(1) \qquad V(P, z) = \max_{r \in R} \ (P - c)(z + r) + \beta \textstyle\sum_i V(P_i, z + r)\pi(P_i|P).$$

The Bellman operator ρ then satisfies Eq. (1) by maximizing a generator's value by adjusting the utilization rate to $z + \rho$. To demonstrate the impact IRE can have on heterogeneous plants, my analysis first focuses on the sensitivity of ρ with respect to γ, c, and the volatility of $\pi(P'|P)$.

3.A.2. Basic Model: Parameterization

Because the spot market price P and the utilization rate z are state variables, both are discretized and bounded from above and below in order for the model to be computationally feasible. In the case of power plant output, z has upper and lower bounds of 100 and 0 by definition. By discretizing z, an additional restriction on the ramping

general effect of market volatility on producer behavior. Section 3.B. includes a more complex model with a realistic marginal cost curve.

decision r is required. Because r determines the size of the state space, r cannot be a continuous variable. The choice set for ramping decisions and utilization rates is therefore restricted to increments of 5. This means that the model allows for generation at $z = 95$ or $z = 100$, for example, but $z = 97.5$ is infeasible. While it is possible to increase the fineness of z and r, such models quickly become more computationally intensive.

The probability matrix π is a 100×100 matrix, where spot market prices range from 1 to 100 and change regularly after short intervals (e.g., 15 minutes). The jth column of the ith row identifies the probability of P_j occurring given the current price P_i. A stylized version of $\pi(P_j|P_i)$ is calculated by weighting prices based on their proximity to the current price using the formula

$$(2) \qquad \pi(P_j|P_i) = \frac{\delta_i - d_j + 1}{\sum_k(\delta_i - d_k + 1)}.$$

In Eq. (2), d_j is the absolute value of the difference between prices P_j and P_i, and δ_i is the difference between P_i and the farthest removed price. This specification allows for probabilities to monotonically decrease the greater the difference between P_j and P_i. A more realistic representation of π would account for additional factors including time of day and characteristics of other active generators in the system. Nevertheless, this initial specification of π allows for a basic analysis of how volatility affects heterogeneous agents separately from a change in average market prices.

Volatility is incorporated into π by re-weighting the distances between prices to calculate probabilities. For the basic model, an increase in market volatility is equivalent to increasing the probabilities of prices farther from P_i (and simultaneously decreasing the probabilities of those prices that are relatively close to P_i). This is done with the adjustment

$$(3) \qquad \pi(P_j|P_i,\alpha) = \frac{\delta_i - d_j + 1 + \alpha}{\sum_k(\delta_i - d_k + 1 + \alpha)}.$$

If $\alpha = 0$, then $\pi(P_j|P_i, \alpha = 0) = \pi(P_j|P_i)$ as in Eq. (2). As α goes to infinity, the distribution $\pi(P_j|P_i, \alpha)$ converges to $1/n$, where n is the number of prices in the distribution. Because the price distribution is bounded from above and below, increasing α will also change next period's expected price, so it is not exactly equivalent to changing a distribution's variance. However, when a random price is simulated over time, the mean price should be 49.5 regardless of the value of α. Five different versions of π were considered in the analysis, with volatility determined by incremental changes to α ranging from $\alpha = 0$ to $\alpha = 0.04$. To demonstrate the effect on volatility from these specifications, $\pi(P'|P = 50, \alpha = 0)$ and $\pi(P'|P = 50, \alpha = 0.04)$ are depicted in Figure 1. While other changes to $\pi(P'|P)$ are possible (e.g., just increasing the probability of $\pi(P' = 100|P)$ or $\pi(P' = 1|P)$), it is difficult to construct other scenarios with an increase in volatility without also significantly altering other distribution properties such as the mean value over time.

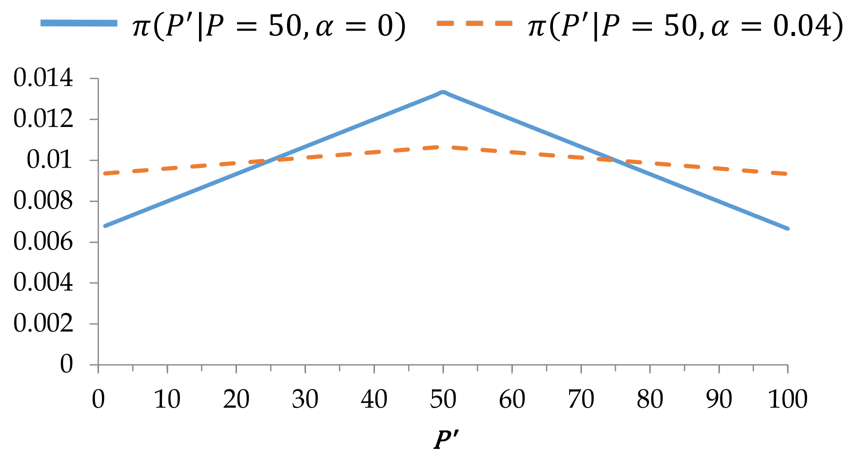

Figure 1. $\pi(P'|P = 50)$

The remaining parameters are identified in Table 1. In addition to five different specifications of π, twenty different ramp rates were chosen, as well as one hundred different marginal costs. Eq. (2) was then optimized by solving for ρ using policy iteration for these two thousand different combinations of parameters.

Table 1. *Basic Model Parameters*

Price volatility (α)	$0, 0.01, 0.02, 0.03, 0.04$
Price (P)	$1, 2, \dots, 100$
Ramp rate (γ)	$5, 10, \dots, 100$
Utilization rate (z)	$0,5, \dots, 100$
Marginal cost (c)	$1,2, \dots, 100$
Discount factor (β)	0.999

3.A.3. Basic Model: Optimal Producer Behavior

Results from select parameter specifications are included here to highlight the impact of price volatility on producer behavior. For simplicity, four different combinations are considered: a plant is either a "low-cost" plant ($c = 40$) or a "high-cost" plant ($c = 60$), and its ramp rate is either slow[15] ($\gamma = 25$) or fast ($\gamma = 75$) with volatility determined by $\alpha = 0$ and $\alpha = 0.04$. With 21 different utilization rates and 100 different prices, there are 2,100 different policy decision that must be calculated to identify ρ, which is simply too many to coherently report in a single graphic. Accordingly, Figure 2 depicts optimal utilization rates, z^*, for these specifications under the volatility schemes $\alpha = 0$ and $\alpha = 0.04$, where the optimal utilization rate is identified based on ρ. Note that ρ is only implicitly included in Figure 2. For example, if a plant with a slow ramp rate were operating at full capacity ($z = 100$) when P suddenly dropped to 1, the producer would want to temporarily shutdown ($z^* = 0$). However, the producer could only ramp down to $z = 75$ ($\rho = -25$) in a single period and would temporarily operate at a loss that period. The next period's ramping decisions would then be determined by the next realization of P, the new utilization rate $z = 75$, and the producer's optimal utilization rate z^*.

[15] In reality, many conventional power plants have even slower ramp rates. Based on the calculations from Benitez et al. (2008), a coal plant would have a ramp rate $\gamma = 8.3$ and a natural gas combined cycle plant a ramp rate of $\gamma = 12.5$. If periods are 15 minutes in length, a "slow" generator with $\gamma = 25$ would be able to ramp up to full capacity in just an hour. Parameters for γ and c were chosen to highlight how slight differences in parameter values can produce marked differences in behavior from volatility, and are not indicative of real world conditions. The more advanced model in Section 3.B. presents results when more realistically calibrated.

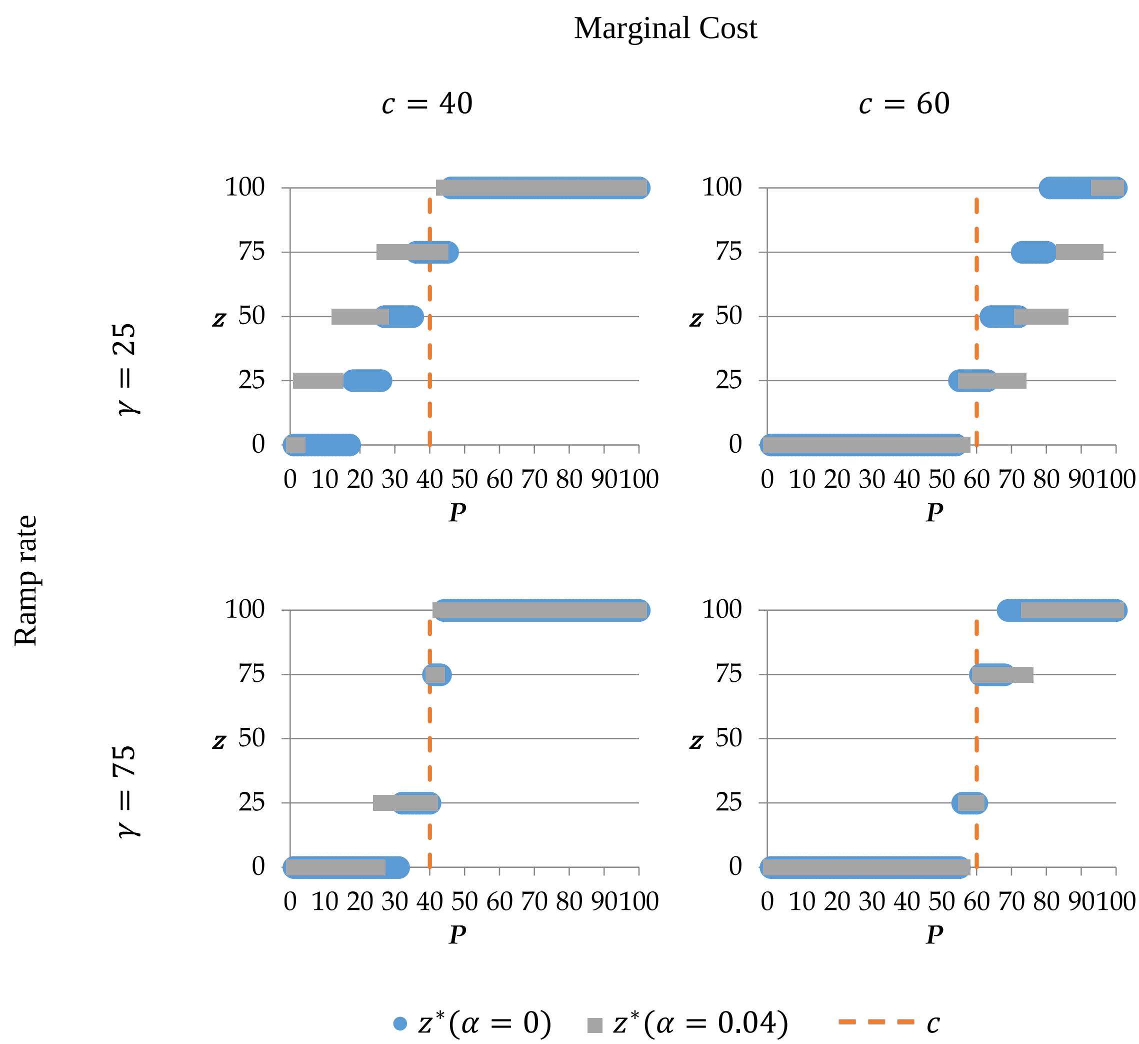

Figure 2. *Optimal Utilization Rates*

Several patterns in producer behavior become apparent in Figure 2. For plants with a slow ramp rate ($\gamma = 25$), producers would prefer one of five different utilization rates depending on the current price. For the low-cost plant ($c = 40$), the producer is often willing to operate at partial capacity for a short term loss in profit. That is, even if the $P < c$, the producer still prefers $z^* > 0$ and operates at a loss. The producer would only want to operate at full capacity when $P \geq 46$, even though the plant operates profitably whenever $P > 40$. On the other hand, the high-cost plant ($c = 60$) with a slow ramping ramp typically prefers to be inactive. So long as $P \geq 55$, it switchs on, even at a loss, though for many $P > c$ it would still prefer not to operate at full capacity.

An increase in price volatility affects the two slow-ramping plants in opposite ways. For the low-cost plant, the producer would be more willing to operate at a temporary loss and will often ramp up to higher utilization rates than before. For example, under the low volatility scheme, the low-cost plant would prefer to operate at $z = 25$ when $P = 20$. But when volatility increases, the low-cost plant would prefer to operate at $z = 50$ when $P = 20$. Conversely, the high-cost plant would often require a higher price than before in order to be willing to generate power. This behavior is likely attributable to the symmetry with which volatility increased. An increase in the probability of price spikes would normally encourage the high-cost plant to switch on more often and operate at a temporary loss. But the probability of sudden price drops also increased, meaning that the persistence of high prices is less certain and producers are more likely to be found operating at a loss before being able to adjust output.

Plants with higher ramp rates have greater flexibility in their output. Because they can ramp down quickly, they can minimize losses when $P < c$. As a result, they are more willing to operate at full capacity than their slow-ramping counterparts, and they are able to narrow the range of prices for which they operate at a temporary loss. Increasing volatility does not have much of an effect on either plant type's optimal ramping rules, except that the low-cost plant will be willing to operate at a loss for a slightly greater range of prices, and the high-cost plant will operate at full capacity for a smaller range of prices.

3.A.4. Basic Model: Simulation Results

Although an analysis of producers' ramping decisions with respect to price volatility identifies underlying trends that could alter plant output, changes in ρ alone do not demonstrate the change in plant output over time. While volatility may have an instantaneous effect on ρ, its effect on cumulative producer output is arguably of greater importance. Once optimal ramping rules were determined for each of the two thousand specifications, simulations were run to track change in output from an increase in the spot market's volatility.

To simulate producer output, the state variables P and z first had to be set in a non-arbitrary manner that would also realistically reflect potential scenarios according to ρ. For example, it may be possible for a plant to find itself operating at full capacity ($z = 100$) while the price is at its minimum ($P = 1$), though the appropriateness of such a starting point for a simulation is dubious. Accordingly, the initial price P_0 was drawn at random according to a probability mass function based on averages of $\pi(P'|P, \alpha)$ for all values of P and P'. Following period 0, a random price walk lasting one hundred and sixteen periods was calculated using transitional probabilities defined by $\pi(P'|P, \alpha = 0)$ and $\pi(P'|P, \alpha = 0.04)$. The generator's initial utilization rate was then set at $z^*(P_0)$ and allowed to vary over time according to ρ. In case the first few observations might be sensitive to the simulation setup, the first twenty periods were not used to calculate results. This left ninety-six periods per iteration to calculate average output and profit, the equivalent of one day if each period represents fifteen minutes. One thousand iterations were run for each of the four thousand different calibrations.

Relative changes in plant-level output over time are illustrated in Figure 3. The results demonstrate that greater volatility can lead to a change in output, though this result is not monotonic and varies with both marginal costs and ramp rates. Interestingly, results are almost perfectly inverted across $c = 50$, which corresponds to the approximate mean price over time under both price distributions. For low-cost generators, output typically either remained unchanged or decreased by as much as 5% if the generator had an extremely low marginal cost or a fast ramp rate, possibly because these plants were already operating near full capacity during simulations with a low price volatility. Many of the low marginal cost, slow-ramping generators actually increased their average utilization rate as the volatility increased. Gains tended to be less than 10%, though the shape of the figure peaks at $c = 50$, $\gamma = 5$ with a 20% increase in output. High-cost generators tended to increase their average output provided they were fast-ramping or had extremely high marginal costs. This is likely because many of these generators tended to operate well below full capacity during simulations with a low price volatility, and any increase in output would then produce noticeable results. Conversely, many of the high marginal cost generators with only mid-range or slow ramping rates were unable to take advantage of the increase in price volatility. Reductions in average output were as significant as 40%.

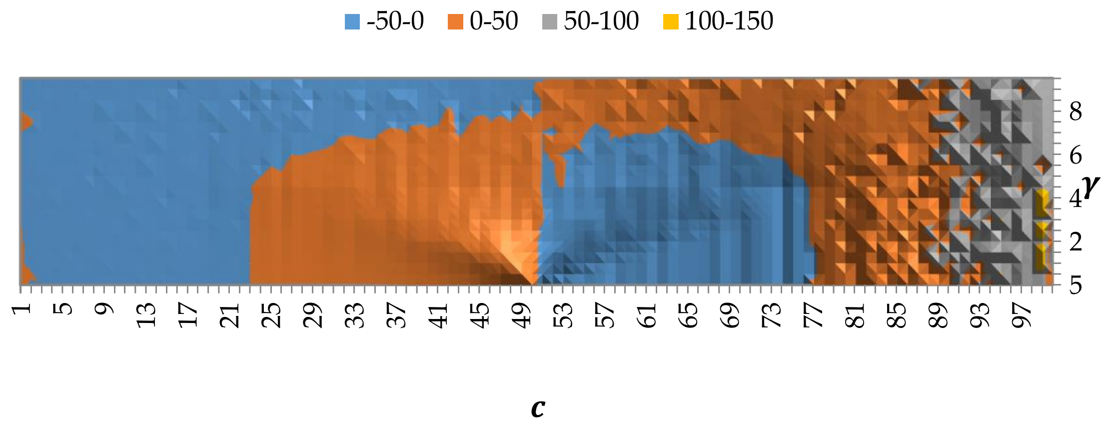

Figure 3. *Changes in Output (%)*

95% confidence intervals for the four previously selected generator specifications are depicted in Figure 4. Although some changes in output are noticeable, none would be described as statistically significant in a strict empirical sense. However, results are practically bounded. For example, slow ramping generators were often operating near full capacity (low marginal cost generator) or were frequently idle (high marginal cost generator). An increase in price volatility shifted each of these generators closer to full capacity or idleness, though it would have been very unlikely for the second confidence interval to shift entirely outside the first. Results may also be sensitive to number of periods per simulation. Nevertheless, the general smoothness of the results in Figure 3 suggest that volatility does have a real effect on generator output, even if this effect appears statistically insignificant for any one particular parameterization.

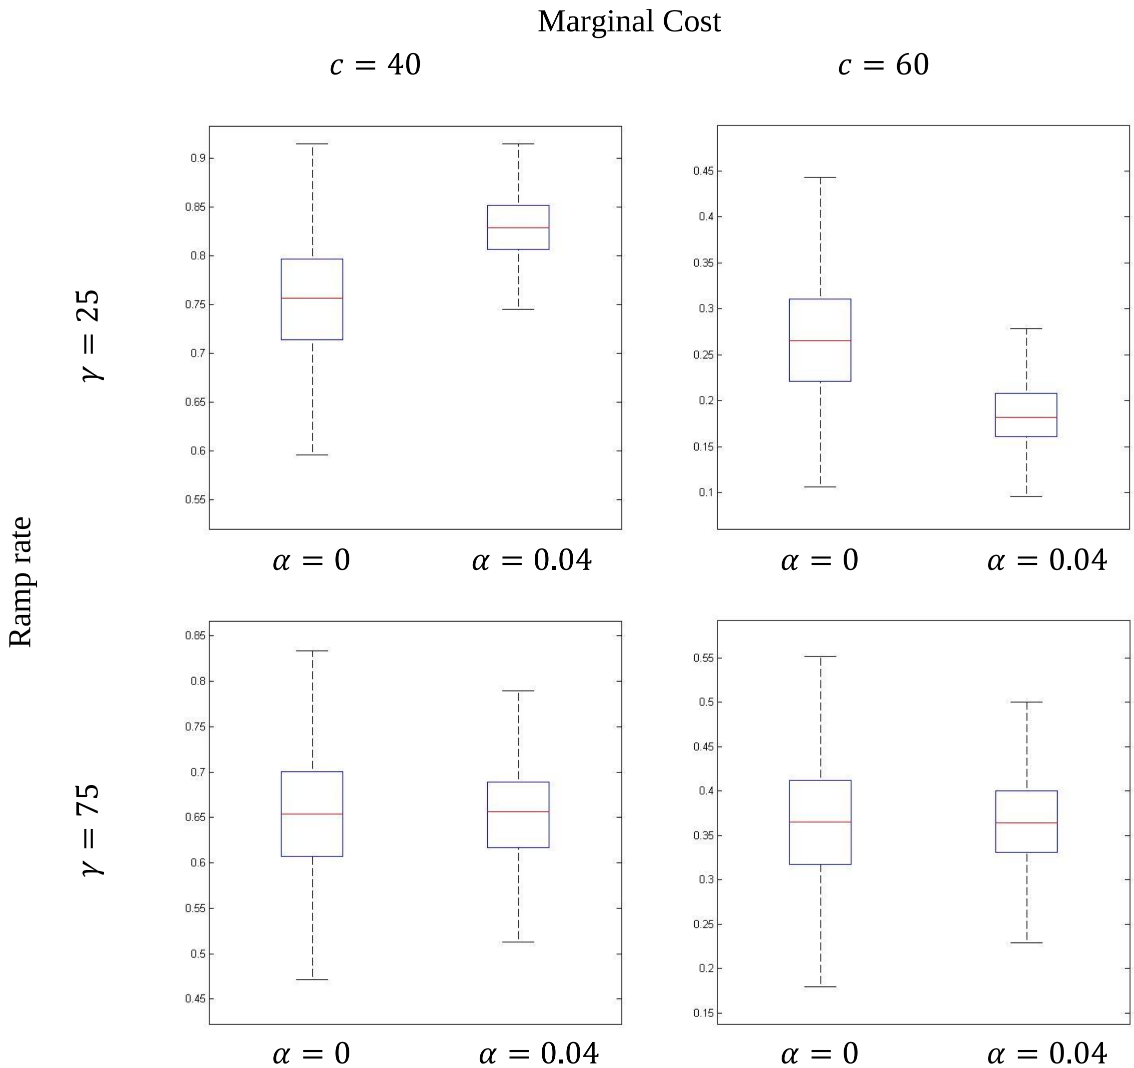

Figure 4. *Select Changes in Output*

The change in price volatility also affected producer profit, though these results are evidently not strongly correlated with changes in output. Relative changes in producer profit are illustrated in Figure 5. Generators with very high marginal costs saw large increases in profit regardless of whether they were slow- or fast-ramping. Slow-ramping plants with marginal costs close to $c = 50$ actually experienced a decrease in profit. All other generators saw modest gains. Figure 6 depicts 95% confidence intervals for profitability of select generator types for both volatility scenarios. While there are some minor differences, changes are rather unsubstantial for these particular parameterizations.

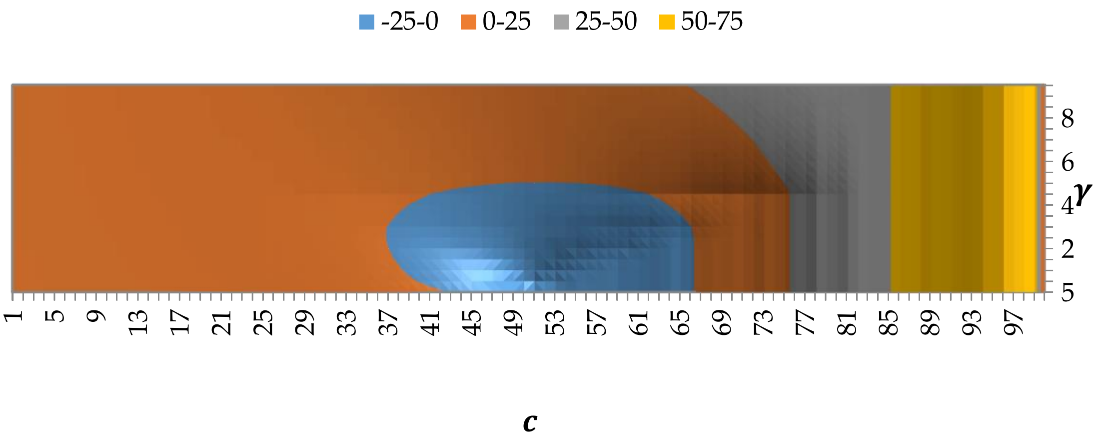

Figure 5. *Changes in Profit (%)*

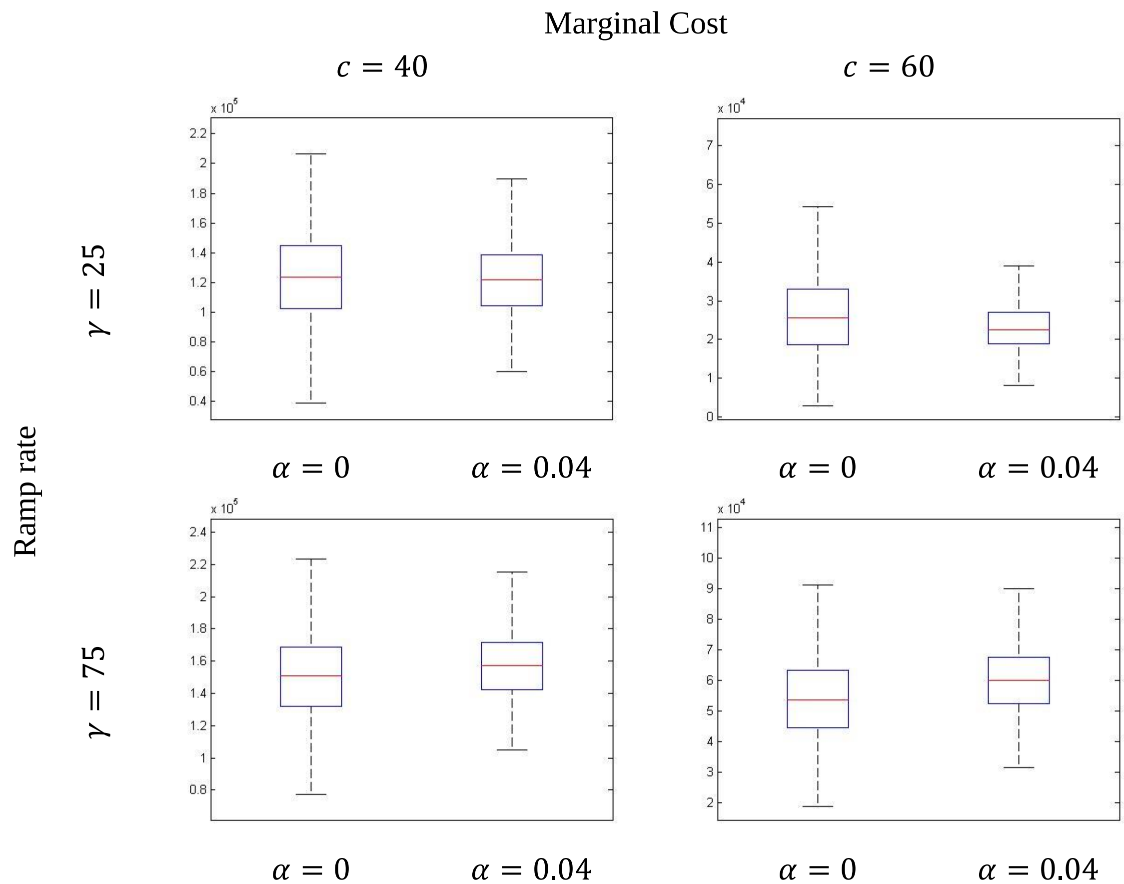

Figure 6. *Select Changes in Profit*

The simulation results can be extended to abatement analysis by considering how output and profit change over time in response to an increase in price volatility. In the short term, capacity is fixed, and abatement would be limited to changes in plant output. Since a change in price volatility has a heterogeneous effect on producers, it is expected that those that increased their output the most in response to a change in volatility would likely not decrease their output as much as predicted using cost minimization techniques, if an increased level of IRE led to greater market volatility. At the same time, producers that decreased their output in response to a change in volatility would likely decrease their output more than the amount predicted by cost minimization. Whether these results would lead to greater or less abatement depends on emissions rates of relevant generators.

In the long term, capacity is not fixed and producers are able to install new generators and retire old ones. In this case, price volatility has implications for abatement analysis through the profitability of various generator types. While some generators experienced a decrease in profitability from the increase in price volatility, others became more profitable. Based on the simulation results, profit-maximizing producers would generally shift away from generators with slow ramp rates. Consequently, the penetration levels of different generator types would change over time in response to the change in price volatility. Since different generator types often have different emission rates, a change in penetration levels will have further implications for long term abatement. Fast-ramping plants tend to be more costly, but they also burn natural gas and have lower emissions rates. Thus, IRE capacity may not abate the same amount during its initial year of operation as in subsequent years due to changes in penetration levels of conventional generators.

3.B. Advanced Model

3.B.1. Advanced Model: Description

The simple model from 3.A. demonstrates the general intuition of how changes in market conditions can affect producer output (and subsequent emissions), though it lacks several parameters and constraints to make it applicable to the case of an actual generator. This section outlines a more complex model that realistically represents the characteristics of fossil fuel-fired generators and incorporates a more detailed power market price distribution into its analysis.

In the basic model, the utilization rate was allowed to range from 0 to 100 without restriction. Results were used to illustrate the general effect of price volatility, but included parameters and calibrations were not reflective of actual power generators. Necessary extensions of this model are the addition of nameplate capacity κ; the establishment of a minimum utilization rate $\underline{z}$, which an active generator cannot operate below without shutting down; imposing a minimum down-time τ if the generator is switched off; and introducing a fixed startup cost F. These new parameters can easily be incorporated into the previous model and solved using policy iteration, though the state space can become quite large depending on size of τ. However, the inclusion of τ is necessary, as it incorporates a power producers' discrete on/off decision into the optimal generation behavior.

The advanced model also relaxes the assumption of constant marginal costs. Many thermal generators actually operate more efficiently when at full capacity. This translates to needing a disproportionately high amount of fuel per MW when operating near $\underline{z}$ relative to the amount required to operate at $z = 100$. As a result, most generators exhibit increasing concave marginal costs. The model is fitted to include such a cost structure using the quadratic cost coefficients c_1 and c_2.[16] An additional startup cost F is also incurred if the generator is off but switched on.

[16] In practice, c_2 is quite small, so the generic assumption of constant marginal cost may not be too inappropriate.

After accounting for the new parameters, the producer has several choices depending on its state. If the generator is on, its instantaneous profit is $P\left(\frac{z}{100}\right)\kappa - (c_1 z + c_2 z^2)$. If active, it can produce up to κ MW by operating at $z = 100$ or as little as $\left(\frac{\underline{z}}{100}\right)\kappa$. But the exact choice set for z in any one period is further restricted by the ramp rate γ, as was the case in the simple model. It can also switch off, which reduces the utilization rate to 0 while simultaneously imposing a minimum of τ periods in which the generator cannot be reactivated. Once an inactive generator has been idle for τ periods, it can either continue to remain off, or be switched on to $\underline{z}$. If activated, its instantaneous profit is $P\left(\frac{\underline{z}}{100}\right)\kappa - \left(c_1\underline{z} + c_2\underline{z}^2\right) - F$. If t is a state variable that identifies how many periods the generator has been inactive and $S(z, r, F, t)$ is a function equal to either 0 or F depending on whether the generator is switched on, then the producer's optimal output decision under the advanced model can formally be described with the Bellman equation

(4) $$V(P, z, t) = \max_{r \in R} \; P\left(\frac{z+r}{100}\right)\kappa - (c_1(z+r) + c_2(z+r)^2) - S(z, r, F, t) + \beta \sum_i V(P_i, z+r, t')\pi(P_i|P).$$

The Bellman operator ρ then satisfies Eq. (4) by maximizing the generators current and expected value by adjusting the current utilization rate as permitted by γ, $\underline{z}$, and τ.

3.B.2. Advanced Model: Parameterization

Because the basic model had only 2,000 different combinations of parameters γ and c, it was relatively straightforward to perform basic optimization and simulations for every possible setting and still be able to interpret results. However, the advanced model employs a much larger array of parameters (γ, c_1, c_2, $\underline{z}$, κ, τ, F) and the resulting set of feasible combinations makes it impractical to optimize and simulate for all possible parameterizations. Accordingly, the parameters of the complex model are set to approximate a 200 MW natural gas generator. Values for c_1, c_2, and $\underline{z}$ were obtained from Wood and Wollenberg (1984) and F was set according to information in Stoft (2002). Parameters are listed in Table 2. Parameters γ, τ, and κ were calibrated for periods of 15 minutes in length. Thus, although the generator can produce 200 MW in an hour, it can only produce one quarter of that in a single period. Similarly, $\gamma = 50$ and $\tau = 4$ correspond to the plant being able to ramp up to full capacity in a half hour and a minimum down-time of one hour, respectively.

Table 2. *Advanced Model Parameters*

Ramp rate (γ)	50
Marginal Cost 1 (c_1)	22.0911
Marginal Cost 2 (c_2)	-0.0341
Min utilization rate ($\underline{z}$)	25
Capacity (κ)	50
Min down time (τ)	4
Startup cost (F)	4000
Discount rate (β)	0.999

Note: κ and τ are calibrated for periods of 15 minutes in length.

The simple price distribution described by π is also refined in the advanced model to better reflect a realistic set of prices and transitional probabilities. Initially, the full set of prices is defined as $\{P\}$ and divided into three separate regimes: low, mid, and high. The low price regime ranges from -25 to 10 in increments of 5. The mid price regime ranges from 15 to 90 in increments of 1. The high price regime includes 100, 200, 500. 1000, and 2000. Prices are considered to be in dollars per Megawatt hour [MWH]. In reality, the price of power can vary in units as fine as cents, though the restriction of a finite and reasonably sized state space for computability necessitates discretizing prices and omitting certain prices (in both low and high regimes) that are uncommon in any case. Nevertheless, the range of included prices is realistic and based on observed prices in the Texas interconnection.

Transitional probabilities are calculated using a normal distribution. Specifically, the probability of a price change as defined by $P_j - P_i$ is calculated with a normal distribution with a mean of 0. Findings of Woo et al. (2011b), who quantified the effect wind power has on average market prices and price variance in the Texas interconnection, are used to simulate the effects of wind power on power market conditions. Specifically, Woo et al. demonstrate that a 10% increase in wind power capacity would decrease power market prices by $3.825/MWH and increase price variance from 1190.3 to 1254.1.[17] Thus, $\{P'\}$ identifies a price regime identical to $\{P\}$ except each observation is smaller by 3.825, and a base case variance of $\sigma_0^2 = 1190.3$ is also used in conjunction with an "increased volatility" scenario with $\sigma_1^2 = 1254.1$. Simulations are then run using all four combinations of P and σ^2, and the generator calibration in Table 2.

3.B.3. Advanced Model: Simulation Results

To adequately measure the impact price volatility has on producers, simulations are run similar to those for the simple model. The initial price P_0 is set using the same

[17] This result is specific to the West zone of ERCOT's market setup, where a majority of Texas's wind power is located.

technique as in 3.A.2., except it is restricted to prices in the mid-range of either $\{P\}$ or $\{P'\}$. Additionally, the random price walk is extended to six hundred and ninety-two periods. This specification allows for twenty periods of adjustment and six hundred and seventy-two periods (seven days) to measure producer profit, output, and fuel use. The initial utilization rate is set to $z^*(P_0)$ if $z^*(P_0) \geq \underline{z}$, and off with the option to switch on otherwise. One thousand iterations are run for each of the four price scenarios.

Confidence intervals showing the impact of different price regimes on output use are depicted in Figure 7. Most of the generation fluctuated between 27,000 and 28,000 MWH under the initial price regime $\{P\}$ and between about 25,000 and 26,000 MWH under $\{P'\}$, approximately a 2,000 MWH reduction attributable to the price decrease from wind power. Since the model has been calibrated to a natural gas generator, this corresponds to roughly 13,900 to 14,400 tCO2 emissions under $\{P\}$ and 12,900 tCO2 to 13,400 tCO2 under $\{P'\}$.[18] Results were evidently not sensitive to either specification of σ^2. This may be because not all possible generator calibrations are sensitive to changes in volatility (as was the case sometimes for the results from the basic model), because the change in volatility was not large enough, or because a change in volatility has no effect on output over time.

[18] Because the generator was parameterized to include concave marginal costs, there are slight differences in instantaneous output and emissions. However, relative changes in cumulative output and emissions in this instance were found to be nearly identical, hence I omit an additional figure with 95% confidence intervals for changes in emissions.

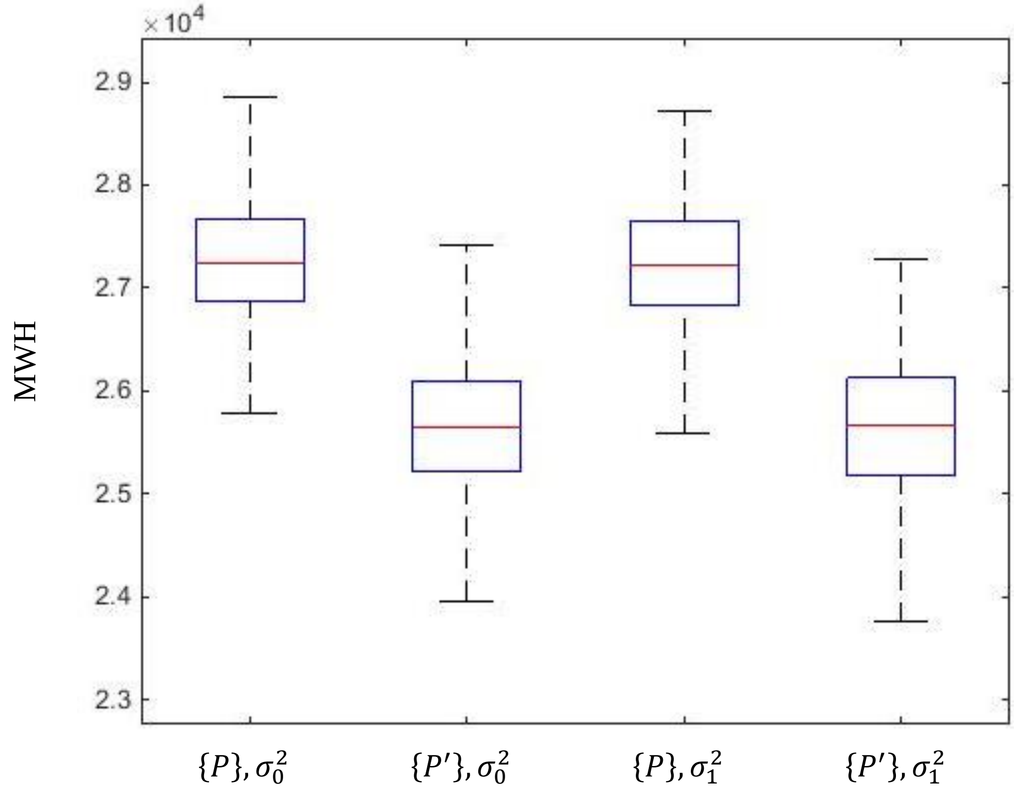

Figure 7. *Changes in Output*

Confidence intervals for profit are included in Figure 8. Profit from the initial price regime tended to be between \$450,000 and \$500,000. For the lower price regime $\{P'\}$, profits are usually between \$350,000 and \$400,000, roughly a 20% decrease in profit from the decrease in price from wind. As with output, results were more sensitive to switching from $\{P\}$ to $\{P'\}$ than to changing the value of σ^2. But the range of profit across scenarios is relatively consistent regardless of which price scenario was used, indicating that these changes in power market conditions have very little effect on a 200 MW natural gas generator.

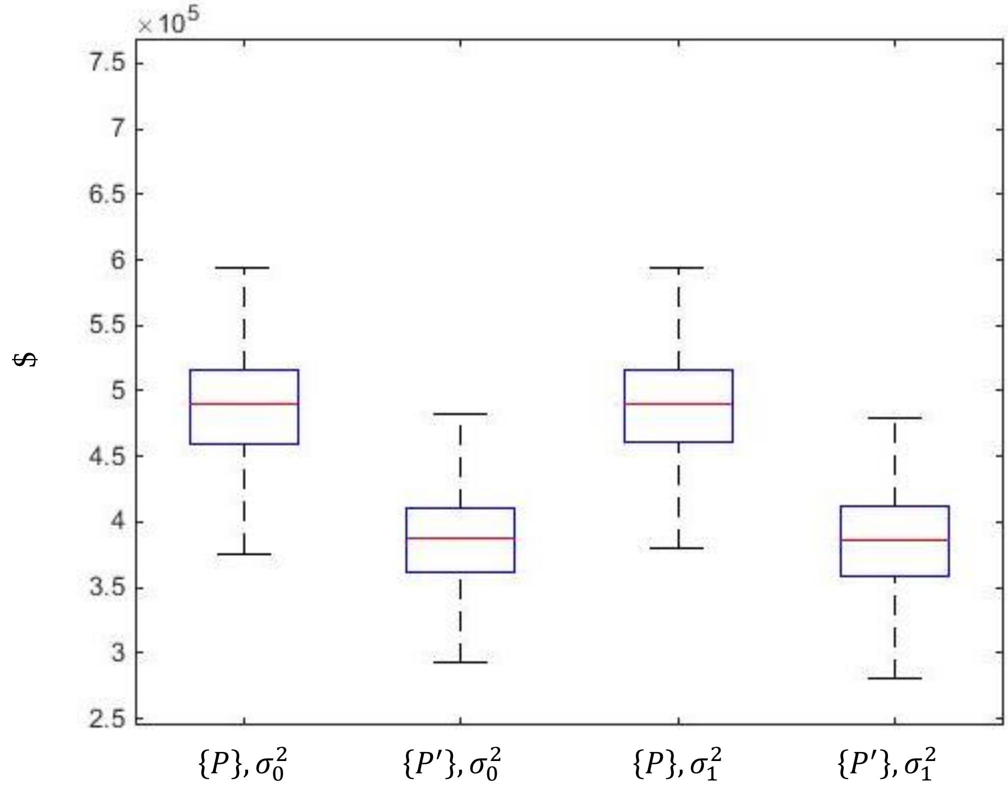

Figure 8. *Changes in Profit*

The simulation results from the advanced model demonstrate the impact market conditions can have on generation, emissions, and profitability of a conventional generator. Admittedly, results indicate that price volatility has little, if any, noticeable impact. However these results are likely sensitive to the specified parameters and transitional probabilities. In this particular case, a 200 MW natural gas generator does not seem to be sensitive to a change in price volatility.

4. Conclusion

Cost-benefit analyses of IRE estimate the displacement of generation from conventional power plants by modeling power system outcomes with ex post cost minimization instead of ex ante profit maximization. Yet the intermittency of wind, solar, and tidal power is expected to increase market volatility as their penetration levels increase, potentially invalidating ex post cost minimization as an appropriate modeling technique because it does not adequately account for changes in uncertainty and risk. This paper offers an analysis of how power market conditions, especially price volatility, affect ex ante profit-maximizing power producers, and therefore assesses the validity of ex post cost minimization for studying the benefits and costs of IRE.

I construct both a basic and an advanced computational model to demonstrate the relationship between price volatility, and producers' ramping rates and marginal costs. With the optimal policy decisions from the basic model, I find that power producers with high marginal costs are more sensitive to market volatility than those with lower marginal costs. However, plants with higher ramping rates also tend to be more immune to any adverse effects from volatility. Generally, an increase in volatility decreases plant idleness, and this effect is especially noticeable in all types of high-cost plants. While simulation results suggest that market volatility will result in relatively more power from high-cost generators, this effect would not be present in results derived using ex post cost-minimization.

Traditional cost-minimizing analyses find that IREs will displace a disproportionate amount of power from natural gas because of its high marginal cost. Yet ex post cost-minimizing techniques are somewhat problematic if intermittents also increase price volatility in a power system, as simulation results from the basic computational model indicate that an increase in price volatility has a non-monotonic effect on power producers depending on their ramping rates and marginal costs. Even so, an application of the advanced computational finds that changes in output, emissions, and generator profitability are not sensitive to a realistic change in price volatility attributable to an increase in IRE. Specifically, the advanced model was calibrated to a 200MW gas-fired generator, and changes in market conditions were derived from Woo et al.'s (2011) study on wind power's effects on average prices and price volatility in the western zone of the Texas interconnection. Through various combinations of alternating high and low price and volatility regimes, the effects of each change in market conditions on output, emissions, and profitability were isolated and explored. In this setting, changes in volatility had no discernible effect on any of the measured outcomes.

Results from this simulation appear to support continued use of ex post cost-minimization because of the lack of consequences from a change in volatility. However, it is difficult to say whether findings from this simulation can be generalized to all cases. Di Cosmo and Valeri (2014), for example, find that an increase in wind power would have a heterogeneous effect on the profitability of various conventional generators, specifically because of the intermittency of wind power and its effect on market conditions. Other calibrations of the advanced computational model may therefore prove to be sensitive to changes in price volatility. Furthermore, the change in price volatility for the advanced model's simulation was rather small. It is possible that a larger change in price volatility, or perhaps if price volatility were characterized beyond variance, simulations would yield different results. The validity of ex post cost-minimization to model the impacts of IREs therefore remains unclear, though results from current efforts do not refute it as an acceptable methodology.

Why Do Distilleries Produce Multiple Ages of Whisky?

1. Introduction

Whisky[19] is one of just a few goods that matures and accrues value over time. Other common examples include timber, which accrues value because trees gain biomass with age, and wine, which accrues value because it can improve in quality even after it is bottled. Whisky, and other aged spirits in general, also has a multi-period production process, but the maturation processes, traits of the final goods, and even the general business decisions of wine and timber producers are not perfectly relatable to whisky. Whisky producers consistently offer a product line in which the primary distinguishing factor within their own brands is product age, which implies the existence of multiple optimal maturation ages for a single producer.

The existence of multiple optimal maturation ages for a single vintage is not usually considered in the literature. In wine economics, researchers explicitly assume a unique optimal maturation age (Goodhue, LaFrance, and Simon, 2009), concede consuming bottles from the same vintage at different years could be optimal in order to track quality improvements in the overall batch (Jaeger, 1989), or employ a model which obfuscates the product's age (Wohlgenant, 1982). In forestry economics, the subject of harvesting trees of heterogeneous ages has been analyzed (Salo and Tahvonen, 2003 and 2004; Uusivuori and Kuuluvainen, 2005), though the emphasis here is on the existence of and convergence towards a steady state with an even distribution of age classes.

---

[19] Both "whisky" and "whiskey" are acceptable spellings, though the usage somewhat depends on the spirit's origin. Because my analysis focuses on Scotch whisky, I opt to use "whisky" throughout the paper, except in cases in which "whiskey" is more appropriate given the context.

Distilleries that sell multiple ages of whisky may be comparable to other general instances of multiproduct firms. In this case, distilleries offer a line of goods that are differentiable vertically (by age) and potentially horizontally (through other product attributes). If consumers differentiate between brands, this would allow distilleries to exert some market power and encourage them to construct product lines so as to maximize their market share and prevent other firms from entering the market. Multiple maturation ages of whisky may therefore be evidence of imperfect competition and limited brand substitutability.

The purpose of this paper is to explore the production decisions of whisky distilleries and determine what market conditions lead them to produce multiple ages of whisky. Accordingly, I begin by constructing a theoretical model with a multi-period maturation process under the assumption that distilleries are price-takers. Results from the model produce several testable conditions to determine the validity of the hypothesis that firms have no market power. Next, I exploit distillery ownership to analyze product line composition and determine how brand substitutability influences product variety at the firm-level.

There are three contributions of this paper to the existing literature. First, it addresses the issue of multiple maturation ages for goods that accrue value over time. Furthermore, it demonstrates how product variety is affected by market concentration. Generally, analyses of market concentration and product variety highlight instances in which two firms merge and they reposition their respective products to avoid cannibalizing profits between brands. Many Scottish distilleries are owned by parent companies that own several more distilleries, allowing for further consideration of how product variety is affected as concentration intensifies. Lastly, the paper exploits the characteristics of single malt Scotch to present a comprehensive and informative measure of product variety. First, product variety can be measured as the total number of products offered by one brand, as is often the case in the literature on product variety. Additionally, products' age statements are used as a second measure to determine how producers' decisions with respect to quality are influenced by market concentration.

Results from the theoretical model are inconclusive as to whether distilleries are price-takers. Although the model potentially enables the calculation of firms' discount rates using prices and product ages without further knowledge of quantity or costs, the subsequent discount rates seem implausible. While a rejection of the calculated discount rates suggests that distilleries are not price-takers, it is possible that factors that the model cannot account for bias results. However, results from the analysis of distilleries' product variety do suggest distilleries are able to exert some market power. Both product line composition and product line size appear to be influenced by closely related brands. This suggests that parent companies are aware of substitutability within their brands, and structure their distilleries' product lines so as to avoid offering comparable products and thus cannibalizing profits between brands.

The remainder of the paper is organized as follows. First, I overview the production process of single malt whisky and explain the significance of maturation. Then, I present a theoretical model and determine under what conditions a price-taking producer would still produce multiple ages of whisky. The model leads to a series of lower- and upper- bounds for observed price ratios, and I am able to use these to estimate firms' discount rates and evaluate the assumption that firms are price-takers. Lastly, I analyze the distilleries' product line composition to evaluate the determinants of a distillery's product variety and assess brand substitutability and market power.

2. Modeling Whisky Production

2.A. Overview

Throughout the paper, I will focus on single malt Scotch whisky, though the analysis could easily be extended to other aged spirits, such as bourbon or rum. One of the key differences between Scotch and other whiskies is that Scotch whisky must be made from malted barley, whereas others (such as bourbon) are mostly made from other grains, such as corn. I will also refer to the spirit as "whisky" regardless of its actual age, though Scotch whisky cannot actually be sold as such until it has been aged for at least three years. In fact, the spirit must satisfy several conditions in order to legally be sold as "Scotch whisky," including a minimum maturation age of three years, a bottled strength of no less than 40 percent alcohol by volume, and no additives are permitted (the sole exception being caramel coloring E150a). I provide a general overview of the production of single malt Scotch whisky here to further elucidate the process.

Single malt Scotch whisky is made with only two ingredients: water and barley. The barley is malted during the initial phase of production by first steeping it in water, then allowing it to dry and germinate. The seeds are next kilned, which halts the germination process and prevents the plant from using its stored sugars. Many distilleries in Scotland toast their malt with peat-fueled furnaces during this phase. Peat, which is decayed organic vegetable matter from bog plants, imparts the smokiness traditionally associated with Scotch whisky. Originally, peat was the primary fuel source available for most distilleries, though the advent of railroads provided them with coke and coal, and thus gave distilleries much greater control over how much peat was used during the kilning process. Consequently, many distilleries' flavor profiles are now only lightly peated.

After kilning, the malt is ground and mixed with hot water, which completes the conversion of the malt's starches into maltose. Yeast is next added to the resulting liquid, and it feasts on the sugars, converting them into alcohol and carbon dioxide over the next several days. After fermentation, the liquid is ready for distillation. Scotch whisky is distilled in a pot still to purify and concentrate the spirit. Pot stills perform distillation in batches, and distilleries must distill the spirit at least twice. Elsewhere, such as Ireland, the spirit must be distilled three times.

After distillation the spirit is a "new-make," which is essentially just un-aged whisky and typically around 70 percent alcohol by volume. The new-make is filled into oak casks and the maturation process finally begins. During maturation, the wood mellows the whisky, imparts some of its flavor, and gives it its color. Generally, older whiskies are smoother and more complex than younger whiskies, as they have been in close contact with the wood for more years. Cask sizes are typically between 180 and 500 liters, with smaller casks being able to enhance the maturation process because of the greater surface area to volume ratio. Because many oak casks are reused, their original contents also add some character to the whisky. While almost all Scotch whisky is aged in casks that previously held bourbon, distilleries sometimes use ex-sherry, port, or rum casks, which impart their own particular flavors.

During storage, the whisky evaporates and the alcohol content typically declines, though to what extent is determined by the climate and the type of warehouse.[20] Whiskies can be bottled at cask strength, which is typically around 55 percent alcohol by volume, though most whiskies are diluted to 40 percent. Although whisky must be matured for at least three years, distilleries are not required to include the whisky's age on the label. If the whisky's age is not included on the label, it is known as a "no age statement" whisky. Standard bottlings have an age statement ranging from 10 to 40 years, though both younger and older whiskies are sometimes available. Unlike wine, whisky will not improve with age once it has been bottled, hence the maturation decision of whisky is at the discretion of the producer, not the consumer.

A distillery will typically produce a "core range" of whiskies, which means that it will consistently bottle and sell the same types of whiskies from year to year, though these whiskies will each possess unique traits to differentiate themselves within the range. The primary distinguishing factor is often age, though distilleries can employ various techniques during the production process to create a whisky that is tangibly different in a way other than additional maturation. Examples include alternating the intensity of peating the malt, maturing whisky in an ex-sherry cask, or bottling at cask strength. The ages in a distillery's core range also tend towards certain numbers. For example, many distilleries produce a 10 or 12 year old [yo.] as their youngest whisky with an age statement, though producing both of these ages (or an 11yo.) is very uncommon. The next youngest whisky produced will typically be between 14 and 16 years old, if another is produced at all. Ages of additional bottlings tend to increase in

[20] The evaporation rate is colloquially known as "the angel's share." In Scotland, it is approximately 2 percent per year.

fixed increments, for example 3 or 5 years. Most distilleries produce three or fewer whiskies, though some distilleries have a core range with more than ten different whiskies.

2.B. Optimal Maturation

Goodhue, LaFrance, and Simon (2009) use a model to analyze the production and aging decisions of a profit-maximizing winery when the firm's output had no impact on the market price of its good. I use a similar model here, with several minor differences. Goodhue, LaFrance, and Simon's model accounts for total quantity of output and includes a convex cost function for the production of its un-aged wines. I instead normalize the quantity of un-aged whisky to 1, since I wish to examine the distillery's decision to mature portions of a batch to different ages irrespective of total quantity. Unlike Goodhue, LaFrance, and Simon, I do include bottling costs, since I wish to distinguish between the marginal revenue of a bottle and its net marginal revenue later in the analysis. I also model time discretely, since distilleries only label their whiskies' ages in whole increments.

Let a denote the age of whisky, where $a = 0, \dots, M$ and M is the maximum maturation age beyond which whisky loses its value. I assume that M is finite but sufficiently large so as to not restrict the optimal solution. The quantity of new-make produced in a year is normalized to 1, and x_a is the portion of new-make to be aged for a years before bottling. Whisky steadily evaporates during its time in storage, so let the evaporation rate of whisky be ε. Then if x_a was initially distilled, only $x_a(1-\varepsilon)^a$ is left after maturation.

Production costs are incurred during three different phases of a whisky's lifetime. Distillation costs are c_D.[21] Given a discount factor of δ and marginal cost c_S, the present value of total storage costs is $\sum_{a=1}^{M}\sum_{v=0}^{a-1}\delta^v c_S x_a$. Note that storage costs are a function of the initial amount of casked whisky and do not decrease even though a percentage of the whisky evaporates. Marginal storage costs are also constant, which is consistent with the modeling assumptions of Krasker (1979), Jaeger (1981), and Goodhue, LaFrance, and Simon (2009) with respect to wine storage. After sufficient maturation, whisky is taken out of its casks and bottled. Marginal bottling costs are c_B, and the present value of total bottling costs is $\sum_{a=0}^{M}\delta^a c_B x_a(1-\varepsilon)^a$.

After the whisky has been matured for a years, it can be sold at the price p_a. While price will naturally reflect consumers' preferences for quality, and product quality increases with age, it is not necessary to assume any further relationship between price and age.[22] In this case, the present value of total revenue is $\sum_{a=0}^{M}\delta^a p_a x_a(1-\varepsilon)^a$. Combining revenue with cumulative costs, the present value of profit from a year's batch of new-make is

$$\sum_{a=0}^{M}\delta^a p_a x_a(1-\varepsilon)^a - c_D - \sum_{a=1}^{M}\sum_{v=0}^{a-1}\delta^v c_S x_a - \sum_{a=0}^{M}\delta^a c_B x_a\,(1-\varepsilon)^a,$$

subject to

[21] Because the analysis focuses on the decision to mature portions of the new-make to different ages, the assumed form of distillation costs turns out to be unimportant. In fact, if the quantity of new-make were not normalized to 1, and distillation costs were instead modeled as a convex function, one could still derive the same results from the model discussed later in this section.

[22] If consumers are heterogeneous with respect to utility of quality, prices may naturally be discontinuous with respect to quality (age). Many countries also regulate the minimum number of years the spirit must aged before it can legally be sold as "whisky." In that case, distilleries are still able to sell the younger spirit, though without the name "whisky" the product price may be considerably different.

$$1 \geq \sum_{a=0}^{M} x_a,$$

$$x_a \geq 0, \quad a = 0, \ldots, M.$$

To find the profit-maximizing combination of maturation ages, I next obtain the Lagrangian $\mathcal{L} = \sum_{a=0}^{M} \delta^a p_a x_a (1-\varepsilon)^a - c_D - \sum_{a=1}^{M} \sum_{v=0}^{a-1} \delta^v c_S x_a - \sum_{a=0}^{M} \delta^a c_B x_a (1 - \varepsilon)^a + \lambda(1 - \sum_{a=0}^{M} x_a)$ and the corresponding first order conditions

$$\partial\mathcal{L}/\partial x_0 = p_0 - c_B - \lambda \leq 0,$$

$$\partial\mathcal{L}/\partial x_i = \delta^i p_i (1-\varepsilon)^i - \sum_{v=0}^{i-1} \delta^v c_S - \delta^i c_B (1-\varepsilon)^i - \lambda \leq 0, \quad i = 1, \ldots, M,$$

$$x_i \, \partial\mathcal{L}/\partial x_i = 0, \quad i = 0, \ldots, M,$$

$$\lambda \geq 0,$$

$$\lambda\left(1 - \sum_{a=0}^{M} x_a\right) = 0.$$

I next use these basic first order conditions to determine when a distillery would find it profit-maximizing to produce multiple ages of whisky. Let x_j^* be the optimal amount of whisky to be matured for j years before being bottled and sold, and assume $x_j^* > 0, j > 0$. Without loss of generality, let $j < k$, where k identifies a second, older age that the distiller also produces, $x_k^* > 0$. I begin with the first order conditions for $x_j^* > 0$ and $x_k^* > 0$ and identify lower- and upper-bounds that the price ratio p_k/p_j must adhere to if production with multiple maturation ages is profit-maximizing.

PROPOSITION 1) If $x_j^* > 0$, $x_k^* > 0$, $j < k$, then

$$\frac{p_k - c_B}{p_j - c_B} > \frac{1}{\delta^{k-j}(1-\varepsilon)^{k-j}}.$$

PROOF: By setting the first order conditions equal to each other and eliminating redundant terms

$$\delta^j p_j (1-\varepsilon)^j - \delta^j c_B (1-\varepsilon)^j = \delta^k p_k (1-\varepsilon)^k - \sum_{v=j}^{k-1} \delta^v c_S - \delta^k c_B (1-\varepsilon)^k.$$

Since $\delta > 0$, it must be the case that $\sum_{v=j}^{k-1} \delta^v c_S > 0$. This means that

$$\delta^j (1-\varepsilon)^j \left(p_j - c_B\right) < \delta^k (1-\varepsilon)^k (p_k - c_B).$$

The above expression can then be rearranged to find the lower limit on the price ratio.

If a distillery produces multiple ages, then the ratio of net marginal revenues is greater than the ratio of discount and evaporation factors. This result is driven by the additional storage costs incurred from aging the whisky longer, and the fact that discounting and evaporation losses tend to make marginal profits from younger whiskies more attractive. In order for the older age to be desirable to a distiller when a younger age is also profitable, the ratio between prices net of bottling costs must exceed this lower bound. Furthermore, $(1-\varepsilon) < 1$ and $\delta < 1$, hence $\delta^j (1-\varepsilon)^j > \delta^k (1-\varepsilon)^k$. This implies that for $\delta^j (1-\varepsilon)^j \left(p_j - c_B\right) < \delta^k (1-\varepsilon)^k (p_k - c_B)$, it must then be the case that $p_j < p_k$. Although one may be tempted to naturally impose the condition $p_0 < \cdots < p_M$, this may not necessarily be true for all ages, even though price should generally increase with quality, and quality is strongly correlated with age. No distillery produces every age of whisky, so the validity of this assumption cannot easily be checked. However, Proposition 1 strongly indicates that older whiskies that are produced should always be more expensive.

PROPOSITION 2) If $x_j^* > 0$, $x_k^* > 0$, $0 < j < k$, then

$$\frac{p_k}{p_j} < \frac{\sum_{v=0}^{k-1} \delta^v}{\left(\sum_{v=0}^{j-1} \delta^v\right)(\delta^{k-j})(1-\varepsilon)^{k-j}}$$

PROOF: Take the first order conditions for x_j^* and x_k^* and solve for c_S

$$\frac{1}{\sum_{v=0}^{j-1} \delta^v}\left(\delta^j p_j (1-\varepsilon)^j - \delta^j c_B (1-\varepsilon)^j - \lambda\right)$$

$$= \frac{1}{\sum_{v=0}^{k-1} \delta^v}\left(\delta^k p_k (1-\varepsilon)^k - \delta^k c_B (1-\varepsilon)^k - \lambda\right).$$

Because $j < k$, $\delta^j c_B (1-\varepsilon)^j / \sum_{v=0}^{j-1} \delta^v > \delta^k c_B (1-\varepsilon)^k / \sum_{v=0}^{k-1} \delta^v$ and $\lambda / \sum_{v=0}^{j-1} \delta^v \geq \lambda / \sum_{v=0}^{k-1} \delta^v$, taking out these terms leaves only

$$\frac{\delta^j p_j (1-\varepsilon)^j}{\sum_{v=0}^{j-1} \delta^v} > \frac{\delta^k p_k (1-\varepsilon)^k}{\sum_{v=0}^{k-1} \delta^v}.$$

The above expression can then be rearranged to find the upper limit on the price ratio.

Even though Proposition 1 establishes that older whiskies must fetch a higher price, Proposition 2 identifies their upper-bound. If the price ratio were to violate this upper-bound, then $x_j^* > 0$ would not be optimal, as the producer could earn greater profits by shifting production to the older whisky. Note that if $\delta = 1$, the distiller does not discount future values and the ratio simplifies to $k/j(1-\varepsilon)^{k-j}$, which is the ratio of ages weighted by the evaporation losses that occur between years j and k. With discounting, the price ratio is the ratio between the present value of marginal storage costs weighted by both evaporation losses and time preferences between years j and k.

Propositions 1 and 2 have several important applications. If the evaporation rate, bottling costs, and the discount factor are known, the limits in Propositions 1 and 2 can be calculated exactly, and the observed price ratios should fall between the lower- and upper-bounds. Alternatively, if the evaporation rate is known but distilleries' discount factors are not, Proposition 2 can be used to estimate the maximum discount factor (or minimum discount rate) possible, such that observed price ratios still adhere to the upper-bound. Proposition 1 does require knowledge about distilleries' marginal bottling costs, but if they are assumed to be negligible ($c_B = 0$), then Proposition 1 can be used to calculate distilleries' minimum discount factor (or maximum discount rate), such that observed price ratios adhere to the lower-bound.[23] If distilleries are not price-takers, then market power unaccounted for in the model could bias the distilleries' calculated discount factors (or discount rate) to unreasonable levels. Propositions 1 and 2 therefore allow me to explore the possibility that distilleries are not price-takers using only observed prices and the evaporation rate.

3. Analyzing Whisky Prices

It is known that whisky aged in Scotland incurs annual evaporation losses of 2 percent, $\varepsilon = 0.02$, during the aging process. However, distilleries' discount rates are not readily known, so constructing the bounds as described by Propositions 1 and 2 and

[23] For $c_B > 0$, the estimated minimum discount factor (maximum discount rate) would be even lower (higher).

directly testing observed price ratios is not a viable option. Instead, I opt to use observed retail prices to estimate distilleries' discount rates. To do so, I collect data on single malt whisky prices to construct price ratios, then solve for the discount rate r, where $\delta = 1/(1+r)$, for which the price ratio would exactly equal the lower- and upper-bounds. Because the price ratio must be strictly greater than the lower-bound and less than the upper-bound, the calculated discount rates identify the range of possible discount rates for which the decision to produce multiple ages could be profit-maximizing without market power.

Although there are approximately one hundred distilleries in Scotland, I employ some selective criteria that preclude many of them from the analysis. First, I am only interested in distilleries that have been active for at least the past twenty years. I ignore younger distilleries because their stocks may not yet be mature enough to sell older whiskies, or young distilleries may produce multiple ages to explore the profitability of various ages, or they may even sell a portion of their immature stock to raise revenue during their initial years of operation. Similarly, distilleries that were previously mothballed and only recently resumed production may not have the ages and quantities on hand that they would prefer. I also only consider whiskies in a distillery's core range, and for which age is the primary distinguishing factor across products. That is, limited and special editions, travel retail products, and special whiskies (e.g., those bottled at cask strength) are omitted from the analysis. I must also exclude distilleries which produce fewer than two ages of whisky.

Without access to producer-end prices, I must rely on retail data for the analysis. This somewhat influences my interpretation of observed price ratios, as consumer-end prices will include markups not paid to producers, and hence should not influence the distillery's decision to produce multiple ages of whisky. But in the case of markups that occur at a fixed rate, this has no impact on the analysis. Let P_i be the retail price of a bottle of whisky aged i years and p_i be the price paid to the producer. Let $P_i = (1+v)p_i$, where v represents a price markup such as the Value Added Tax. The observed price ratio is $P_k/P_j = (1+v)p_k/(1+v)p_j = p_k/p_j$, which is unbiased. On the other hand, retail prices may include a constant markup, which does bias the price ratio. Let $P_i = p_i + s$, where s is a fixed price markup such as shipping costs. Then the observed price ratio is $P_k/P_j = (p_k + s)/\left(p_j + s\right) < p_k/p_j$. Consequently, price ratios constructed from retail data may be downward-biased, which means that any calculated minimum discount rate r from P_k/P_j will also be downward-biased.

Data were collected from Master of Malt, one of the largest online retailers of single malt whisky in the world. I chose this retailer for several reasons. First, because it is one of the largest retailers, it is expected that differences between their prices and producer-end prices will be small. Their available stock of whiskies is also extensive, which is a relative strength compared to many other online inventories because price ratios require at least two observations per distillery, and many other retailers either do not carry a particular brand, or do not have enough whiskies from a particular distillery available. Because Master of Malt is based in the UK, shipping charges between producers and the retailer are expected to be low, and their prices will not include additional charges from importation, a distinct advantage over using price data from a US-based retailer. Prices do include the Value Added Tax, though as previously discussed, this has no significant impact on price ratios. The final dataset includes whiskies from twenty-four Scottish distilleries. On average, each distillery has about three observations. Table 1 includes further summary statistics on the data.

Table 1. *Summary Statistics – Retail Prices*

	Obs.	Mean	Median	Std. Dev.	Min.	Max
Price (£)	72	91.24	46.26	144.59	23.9	895
Age (years)	72	17.46	15.5	7.12	10	40

After constructing the price ratios, I solve for the distilleries' minimum discount rate, $\underline{r}$, using Proposition 2, and maximum discount rate, $\overline{r}$, using Proposition 1 and assuming $c_B = 0$. Because many distilleries have observations for more than just two whiskies, most distilleries have more than one estimated $\underline{r}$ and $\overline{r}$. For example, a distillery with four observations would have six different price ratios and as many as six different estimates for $\underline{r}$ and $\overline{r}$. But because Propositions 1 and 2 establish the distillery's minimum and maximum discount rate, results can be simplified by finding the minimum and maximum values of all calculated $\underline{r}$ and $\overline{r}$'s for each distillery, as these identify the minimum and maximum discount rates that satisfy Propositions 1 and 2 for all observed price ratios. For instances when $\overline{r}$ was less than $\underline{r}$, the distillery's minimum marginal bottling cost, $\underline{c_B}$, was calculated such that $\underline{r} = \overline{r}$, as higher values of c_B lead to higher estimates for $\overline{r}$. Table 2 summarizes the results for the twenty-four distilleries based on the data.[24,25]

Table 2. *Estimated Discount Rates*

Distillery (*Brand*)	Ages	$\underline{r}$ (Ages k/j)	$\overline{r}$ (k/j)	$\underline{c_B}$ (k/j)
Aberlour	10, 12, 16, 18	0.087 (18/16)	0.056 (12/10)	7.38 (12/10)
Arran	10, 14	-0.130 (14/10)	0.034 (14/10)	0
Auchentoshan	12, 18	0.052 (18/12)	0.101 (18/12)	0
Balvenie	12, 14, 17, 21	0.125 (17/14)	0.033 (21/17)	47.74 (21/17)
Bowmore	12, 15, 18	0.150 (15/12)	0.105 (18/15)	13.60 (18/15)
Bunnahabhain	12, 18, 25	0.164 (25/18)	0.071 (18/12)	19.49 (18/12)
Caol Ila	12, 25	0.074 (25/12)	0.105 (25/12)	0
Dalmore	12, 15, 18	0.185 (18/15)	0.079 (15/12)	19.37 (15/12)
Dufftown (*Singleton*)	12, 18	-0.011 (18/12)	0.064 (18/12)	0
Fettercairn	24, 30, 40	0.119 (40/30)	0.049 (30/24)	75.61 (30/24)
Glen Moray	10, 12	-0.025 (12/10)	0.082 (12/10)	0
Glenfarclas	10, 15, 21, 25, 30, 40	0.062 (25/21)	0.044 (30/25)	21.90 (30/25)
Glenfiddich	12, 15, 18, 21	0.183 (21/18)	0.051 (15/12)	20.37 (18/15)
Glengoyne	10, 12, 18	0.036	0.091	0

[24] The relationship between the distillery and its brand(s) is not always straightforward. Springbank distillery, for example, has three separate brands that are differentiated, among other things, based on the number of times the spirit is distilled. The Singleton brand, on the other hand, includes whiskies from three different distilleries, though only whiskies from the Dufftown distillery are included here.

[25] Some distilleries bottle their whiskies at slightly different strengths across ages. I also constructed a dataset which adjusted for alcohol content, though these results were not substantially different.

		(18/12)	(18/12)	
Glenlivet	12, 15, 21, 25	0.159 (21/15)	0.080 (15/12)	12.23 (15/12)
Highland Park	12, 18, 40	0.102 (40/18)	0.099 (18/12)	0.84 (18/12)
Jura	10, 16	-0.022 (16/10)	0.070 (16/10)	0
Laphroaig	10, 18	0.008 (18/10)	0.081 (18/10)	0
Macallan (*Sherry line*)	18, 25	0.212 (25/18)	0.216 (25/18)	0
Pulteney (*Old Pulteney*)	12, 17, 21, 30	0.114 (17/12)	0.091 (21/17)	10.25 (21/17)
Springbank (*Springbank*)	10, 15, 18	0.102 (18/15)	0.043 (15/10)	18.43 (15/10)
Tobermory (*Tobermory*)	10, 15	0.053 (15/10)	0.116 (15/10)	0
Tomatin	12, 15, 18, 30	0.037 (15/12)	0.049 (18/15)	0
Tomintoul	10, 14, 21	-0.061 (14/10)	0.031 (21/14)	0

Table 2 identifies four distilleries with a minimum discount rate that is negative. This does not actually mean that these distilleries have irrational time preferences, merely that I cannot rule out the possibility that their discount rate is exceptionally low or even zero. According to the data, eleven of the sampled distilleries must also have a discount rate of at least 0.1, or else their decision to produce some of their younger whiskies could not be described as profit-maximizing. It is surprising that the estimated minimum discount rate is so high for many distilleries. There are several possible explanations that may account for such high minimum discount rates. It may be the case that the older, more expensive whiskies are somehow biasing estimates upwards. But all products selected for the analysis are within each distillery's core range, and are therefore standard bottlings. Furthermore, the dataset includes five distilleries with at least one 30+ yo.whisky, yet two of these still have an estimated minimum discount rate less than 0.1. In fact, only one distillery's minimum discount rate was determined by a price ratio that included at least one 30+ yo. whisky. The inclusion of exceptionally old whiskies is therefore not the source of upward-bias in $\underline{r}$.

Another possibility is that a distillery's youngest whisky is under-priced to attract consumers to the brand. This would lead to abnormally high estimates for $\underline{r}$ for prices ratios that included the youngest age in the line, yet this is decidedly not the case. High minimum discount rates for most distilleries are concentrated among prices ratios between whiskies in the middle of the range. Of the seven of the distilleries that have at least four expressions represented in the analysis, five of them have a high minimum discount rate because of a price ratio between two whiskies in the middle of the distillery's age range. The exact relationship between age and $\underline{r}$ is therefore unclear, though it seems that neither very cheap young whiskies nor expensive old whiskies are the dominant factor.

The maximum discount rates, $\overline{r}$, as calculated according to Proposition 1 and assuming $c_B = 0$, show about as much variability as $\underline{r}$. But an inconsistency often appears for many distilleries' minimum and maximum discount rates, as data from the price ratios frequently report $\overline{r} < \underline{r}$. Naturally, $\overline{r}$ was calculated assuming $c_B = 0$, and higher marginal bottling costs could adjust the calculated maximum discount rate such that $\underline{r} < \overline{r}$. Yet for most distilleries with an initial values of $\overline{r} < \underline{r}$, the necessary minimum bottling cost appears unrealistically high. Of the twelve distilleries that require $c_B > 0$ in order for $\underline{r} < \overline{r}$ to hold, ten distilleries would need to have a minimum marginal bottling cost of at least £10. In the case of Balvenie, whose minimum bottling cost is £47.74, c_B would then exceed the price of their youngest whisky in the dataset, indicating that it would never be profitable to bottle their youngest expression, a clear contradiction in the assumption that all firms are profit-maximizing.

While the original model finds that a firm can produce multiple ages as a profit-maximizing solution, the results from observed price ratios of whisky do not support this. It is worth noting that these conditions were derived assuming the distillery has perfect foresight with respect to prices, that costs are stable throughout time, and that there is no difference in marginal costs between ages. Krasker (1979), Jaeger (1981), and Ashenfelter (2008) demonstrate how uncertainty with respect to future vintage quality can influence wine prices over time. But because whisky is produced in a more controlled environment, it is much less sensitive to weather patterns, hence quality uncertainty is not a significant issue with respect to aging whisky. On the other hand, both demand and input prices can fluctuate, and Jaeger (1981) and Wohlgenant (1982) found evidence suggesting that such uncertainty can likewise influence the production and maturation decisions of a winery. For whisky, the price of barley has the greatest potential to fluctuate unpredictably from year to year, though this element of uncertainty would not affect the distillery's decision to mature a batch of new-make to multiple ages according to the model. It may then be the case that producers offer a diverse product line to hedge against demand uncertainty (Carlton and Dana, 2008; Chen, Yeh, and Hu, 2011). Lastly, if marginal costs, especially marginal storage costs, are not constant across years, then calculated discount rates would be biased.

An alternative explanation for biased calculated discount rates and multiple maturation ages in this particular market is the possibility of imperfect competition. There are approximately one hundred distilleries in Scotland, the largest of which accounts for less than four percent of total industry capacity. I previously imposed the assumption that distilleries were price-takers and that p_a was thus fixed. Instead, if the price the distillery faces for a whisky aged a years is $p_a(x_a)$, where $p_a(x_a)$ is a function decreasing in x_a and at least once differentiable, then multiple maturation ages can still be optimal ($x_j^* > 0$, $x_k^* > 0$, $0 < j < k$), provided that the price ratios adhere to the lower- and upper-bounds

$$\frac{p_k(x_k^*) - c_B}{p_j(x_j^*) - c_B} > \frac{1}{(\delta^{k-j})(1-\varepsilon)^{k-j}} + \left(\frac{1}{(\delta^k)(1-\varepsilon)^k\left(p_j(x_j^*) - c_B\right)}\right)\left(\delta^j p_j'(x_j^*) x_j^* (1-\varepsilon)^j - \delta^k p_k'(x_k^*) x_k^* (1-\varepsilon)^k\right)$$

and

$$\frac{p_k(x_k^*)}{p_j(x_j^*)} < \frac{\sum_{v=0}^{k-1}\delta^v}{\left(\sum_{v=0}^{j-1}\delta^v\right)(\delta^{k-j})(1-\varepsilon)^{k-j}} + \left(\frac{\sum_{v=0}^{k-1}\delta^v}{p_j(x_j^*)\delta^k(1-\varepsilon)^k}\right)\left(\frac{\delta^j p_j'(x_j^*) x_j^* (1-\varepsilon)^j}{\left(\sum_{v=0}^{j-1}\delta^v\right)} - \frac{\delta^k p_k'(x_k^*) x_k^* (1-\varepsilon)^k}{\left(\sum_{v=0}^{k-1}\delta^v\right)}\right).$$

Propositions 1 and 2 were useful because they only required the evaporation rate and observed prices, but estimating δ and r using the new upper-bound requires knowledge of x_a^* and the marginal change in price, or the elasticities of price for both p_j and p_k. To my knowledge, no such estimations have been done to date, necessitating the omission of the second term in the original estimations. However, the second term in both expressions has the potential to bias estimates of δ and r if it is omitted. For example, if $\delta^j p_j'(x_j^*)x_j^*(1-\varepsilon)^j/\sum_{v=0}^{j-1}\delta^v > \delta^k p_k'(x_k^*)x_k^*(1-\varepsilon)^k/\sum_{v=0}^{k-1}\delta^v$, the second term will be positive and estimates of $\underline{r}$ that are not calculated with the second term will be upward-biased. Consequently, a distillery's market power has the potential to influence the estimated discount rates.

4. Market Power and Product Differentiation

The preceding section demonstrates that the market for single malt Scotch whisky is likely not perfectly competitive and that distilleries are not price-takers. While this may not be surprising for some of the biggest brands (e.g., Glenfiddich and Glenlivet), results from Table 2 indicate that even many lesser-known brands from small distilleries are able to exert some market power. Market power and product line size may be further affected by the degree of substitutability between brands: because of regional differences in inputs (i.e., the distillery's source of water), as well as production practices specific to distilleries and regions, one distillery's 12yo. single malt Scotch is necessarily different than another's 12yo. single malt Scotch. A distillery may therefore offer multiple expressions of whisky to capture a greater share of the market, with particular products targeting different segments that the distillery is well-suited to serve. However, it remains unclear which product attributes other than age are important to consumers and therefore brand differentiation.

In the field of industrial organization, particular attention has been given to firms' product variety, with analyses that include identifying where in a particular state-space firms strategically locate their products in relation to one another. The general profile of a distillery's whisky is often classified with as many as a dozen flavor categories,[26] and distilleries have further opportunities to distinguish products within their brand using various cask finishes, bottling their whisky at different cask strengths, etc. Hence, precisely defining a suitable state-space to describe Scotch whisky characteristics is impractical in theory, and infeasible in practice, even if only a subset of these flavors determines market segmentation. However, distilleries in Scotland are categorized based on which region they are located in. While regional classification does not necessarily bind the distillery to particular production practices, certain trends do emerge. For example, distilleries from Islay produce some of the smokiest whiskies, whereas those produced in the Speyside region are sweeter and considerably less smoky. Almost half of all Scottish distilleries are located within the Speyside region, with another quarter of distilleries in the Highland region. The remaining regions each have eight or fewer distilleries. Regions are illustrated in Figure 1.

[26] For example, researchers in the Department of Mathematics & Statistics at the University of Strathclyde profile distilleries' whiskies based on the following flavor categories: body, sweetness, smoky, medicinal, tobacco, honey, spicy, winey, nutty, malty, fruity, and floral. Further information and the data are available at https://www.mathstat.strath.ac.uk/outreach/nessie/nessie_whisky.html (Last accessed 1/20/2015).

Figure 1. *Scotch Producing Regions*

There are several possible scenarios for brand substitutability that would each uniquely affect how distilleries determine their product lines in order to maximize profits. First, consumers may not care about the various attributes that distinguish the single malts amongst themselves and treat all brands as interchangeable. That is, even if there are differences between brands, consumers regard a whisky from one distillery as if it is identical to a whisky with the same age from another distillery. In this scenario, anCnoc 12yo. (a Speyside), Bowmore 12yo. (an Islay), and Highland Park 12yo. (an Island) are all considered perfect substitutes, as consumers assume they are homogenous for their purposes. Or consumers may give more weight to some flavor characteristics than others. Then most consumers would regard Aberlour 12yo. and Balvenie 12yo. (both Speysides) as better substitutes for anCnoc 12yo. than either Bowmore 12yo. or Highland Park 12yo. Lastly, consumers may have considerably stark preferences with brands' flavor characteristics, and view each brand of single malt Scotch as a unique product. In effect, brands are not substitutes for another, and each producer operates as a monopoly.

In the aforementioned scenarios, it is assumed the firms are aware of consumers' preferences and selection criteria, and have chosen the composition of their product lines accordingly. A distillery's product variety is therefore a reflection of its substitutability, or location, with respect to other brands. In this context, a firm's "product variety" may be impacted either through the total number of products offered, the product positioning in regards to important attributes, or both. For example, Alexander (1997) and Watson (2009) find that when new firms enter a particular segment of the market where other firms are already competitive, the total number of products offered by each firm decreases. But this is not necessarily an indication that cumulative product variety amongst all brands decreased. Mazzeo (2002), for instance, demonstrates that product positioning and quality levels are both important criteria, and that if firms enter a particular market segment they may offer a product at a different quality level so as to distinguish themselves from potential competitors.

Product positioning in response to brand substitutability would be further reflected in instances when a single firm owns multiple distilleries. Decades of growth in the industry have led to instances of firms establishing new distilleries adjacent to distilleries they already own, as well as firms buying distilleries from one another to create a portfolio of distilleries. This creates a scenario in which distilleries are owned by parent companies, and many distilleries have "sister" distilleries whose whiskies may be close substitutes. By the end of 2013, only sixteen of the ninety-nine distilleries were owned by a company with no other distilleries; the rest were owned by a parent company that owned at least two distilleries. For example, Diageo owned twenty-eight distilleries, and these distilleries may have selected their product lines so as to avoid competing against other Diageo-owned Scotch brands.

Some attention has been given to the effect of mergers on the market for blended Scotch whiskies. Specifically, Ashenfelter and Hosken (2010) study the merger between Guinness and Grand Metropolitan to form Diageo, and they find that it led to a significant price increase for brands such J&B, Johnny Walker, and Scoresby. However, the effect of mergers on single malt Scotches remains unstudied. Generally, it is understood that mergers tend to increase product variety in the overall market. Merging firms have sudden incentive to reposition their products farther away from each other to better differentiate their brands so as to avoid cannibalizing profits while also capturing a greater market share for the parent firm (Berry and Waldfogel, 2001; Gandhi et al., 2008, Draganska, Mazzeo, and Seim, 2009; Sweeting, 2010). Unfortunately, the general methodologies of these analyses rely heavily on data such as costs, product prices, and a straightforward way to calculate distances between firms (in instances when firms literal location is of interest) and/or distances between goods (in terms of quality levels and other quantifiable characteristics).

In the case of single malt Scotch whisky, the question of producing multiple ages of a vintage good appears to be a question of market power and substitutability between brands. In order to gauge substitutability, analyzing distilleries' product line variety in light of sister distilleries is therefore key. Because of a lack of cost data and a superfluous number of flavor characteristics, the necessary strategy is therefore to use distilleries' regional classifications as a proxy for a measurement to characterize differences in brands' flavor profiles, and determine how sister distilleries' product lines affect the products offered by distilleries as compared to those without sister distilleries. Specifically, I estimate the probability that distilleries offer particular expressions (Subsection 4.A.) as well as the overall number of whiskies they offer (Subsection 4.B.), while also accounting for the presence of sister distilleries and potential substitutability of products.

4.A. Brand Substitutability and Product Positioning

In order to measure product positioning, whiskies are classified according to the single malt Scotch categories from the San Francisco World Spirits Competition.[27] Categories include: whiskies 12 years and younger, whiskies between 13 and 19 years, whiskies 20 years and older, and whiskies with no age statement [NAS]. A more robust measure of product positioning would perhaps account for all yearly increments in age, as well as basic attributes such as alcohol content and cask-finishes, but discretizing a continuum quality levels to better facilitate analysis is common in practice (e.g., Mazzeo, 2002).

A linear probability model is used to determine what factors influence a distillery's decision to offer at least one expression of whisky in a given age category. The general setup is as follows. The probability that distillery i offers age class A is

[27] This international competition's categories were selected for several reasons. First, it is one of the premiere spirits competitions in the world, so its chosen categories are recognized by the industry and are not arbitrary. Other competitions and industry experts merely suggest categories for Scotch that differentiate products based on prices, classify single malts into two groups depending on whether a bottle has an age statement or not, or offer no categorization scheme at all. The categories from the San Francisco World Spirits Competition present a logical breakdown between single malts based on ages.

determined by the distillery's capacity, what regional style of Scotch it produces, and whether a sister distillery offers a whisky in the same category:

(1) $A_i = \alpha + \beta Capacity_i + \gamma Speyside_i + \delta Sister_i + \varepsilon_i$.

In the equation above, A_i is equal to one if distillery i has at least one product of category type A and is equal to zero otherwise. $Capacity$ is the distillery's capacity, measured in millions of liters of pure alcohol per year, $Speyside$ is a dummy variable equal to one if the distillery is located in the Speyside region, and $Sister$ is the number of sister distilleries that offers a whisky in the same age category. It is expected that a distillery's capacity will have a nonnegative effect on the probability of offering a particular age ($\beta \geq 0$), since greater capacity corresponds with the ability to produce more output annually. If the distillery is located within Speyside, this may decrease the probability that it produces a particular age ($\gamma \leq 0$), since Speyside is the most populous region for distilleries and this segment of the market may therefore be more competitive. Lastly, if a parent company owns multiple distilleries, it may be less inclined to have its brands competing against each other and cannibalizing its profits. In this case, a distillery is less likely to offer age category A if a sister distillery already produces that age ($\delta \leq 0$).

One extension is to see if all sister distilleries affect each other's product positioning equally. Because of regional style differences, distilleries may be serving different parts of the overall market, and individual brands might not be close substitutes if they are located in different regions. Accordingly, a second specification splits the $Sister$ variable into two variables: $SisterRegion$ and $SisterNon$, where $SisterRegion$ is the number of sister distilleries in the same region that offers at whisky in age category A, and $SisterNon$ is the number of sister distilleries in a different region offering category A. The linear probability model then becomes

(2) $A_i = \alpha + \beta Capacity_i + \gamma Speyside_i + \delta_1 SisterRegion_i + \delta_2 SisterNon_i + \varepsilon_i$.

As with the first specification, it is expected that if a sister distillery offers category A, distillery i is less likely to do so in order to avoid competing against its parent company's brands ($\delta_x \leq 0$; $x = 1, 2$). However, if consumers do not perceive all whiskies as homogenous, then sister distilleries in the same region are likely to have a stronger effect than those distilleries in a different region ($\delta_1 \leq \delta_2$).

A linear probability model is useful for this estimation because of the obvious endogenous relationship between distillery i offering A when accounting for the products of sister distillery j, since distillery i would be considered a sister distillery for observation j. Accordingly, the initial specification is executed as a two-stage least squares model, with total sister capacity and both the total numbers of sister distilleries located in Speyside and elsewhere as instruments for the sister distillery variables. For the second specification, which distinguishes between sister distilleries by region, instruments include total sister capacity in the same region, total sister capacity in other regions, and the total number of sister distilleries located in Speyside and elsewhere.

Data on distilleries' product lines, capacities, and ownership come from the Malt Whisky Yearbook 2013, an annual publication containing industry-relevant information and statistics. Product lines were also cross-checked with brands' own websites in order to ensure accuracy. Summary statistics are presented in Table 3.

Table 3. *Summary Statistics – Product Positioning*

	Count	Mean	Min.	Max.
Dependent Variables				
$A \leq 12$	99	0.566	0	1
$13 \leq A \leq 19$	93*	0.441	0	1
$A \geq 20$	93*	0.430	0	1
NAS	99	0.596	0	1
Exogenous Variables				
$Capacity$	99	3.224	0.02	12.5
$Speyside$	99	0.485	0	1
Endogenous Variables				
$Sister$				
$A \leq 12$	99	4.677	0	13
$13 \leq A \leq 19$	93*	3.806	0	10
$A \geq 20$	93*	2.656	0	7
NAS	99	4.152	0	12
$SisterRegion$				
$A \leq 12$	99	1.758	0	6
$13 \leq A \leq 19$	93*	1.280	0	4
$A \geq 20$	93*	0.968	0	4
NAS	99	1.232	0	5
$SisterNon$				
$A \leq 12$	99	2.919	0	12
$13 \leq A \leq 19$	93*	2.527	0	10
$A \geq 20$	93*	1.688	0	7
NAS	99	2.919	0	11

*Six distilleries are omitted because they were founded less than 13 years ago.

The data show that product ages are well-distributed among distilleries. A randomly selected distillery has roughly a fifty percent chance of producing any of the four age categories, with slightly higher probabilities for both the youngest and No Age Statement categories and slightly lower probabilities for all ages above 12 years. At the same time, roughly eighty percent of all distilleries have at least one sister distillery that offers a product in the same age category. Even accounting for regional differences between sister distilleries, more than half of all distilleries are paired with at least two sister distilleries (one inside its region and one outside) that produces the relevant age category.

Empirical results are presented in Table 4.[28] For the youngest age category, $A \leq 12$, neither the distillery's capacity nor its regional style appear to have much impact on a distillery's decision to offer this type of whisky. Coefficients are not statistically different from zero, and they tend to be small. Sister distilleries' products do appear to have an impact, though. If the distillery has a sister distillery which offers a whisky 12 years or younger, it is less likely to offer it as well. Furthermore, when sister distilleries are differentiated according to their regional location, results suggest that a sister distillery in the same region has the larger impact, as the coefficient for $SisterRegion$ is negative and statistically different from zero, whereas the coefficient for $SisterNon$ is smaller in absolute size and not statistically significant. This indicates that parent firms are aware that offering similar products from sister distilleries will cannibalize their own profits, but regionally different whiskies do not appear to be close substitutes based on product variety decisions, hence they are not of concern to the parent firm.

For the mid-age category, distillery size has a positive effect on the probability that a distillery will produce a whisky between 13 and 19 years of age. The average marginal effect for an additional one million units of capacity is between 6.3% and 6.8%, depending on the specification for modeling sister distillery product positioning. Regional styles also appear to matter in this instance, as Speyside distilleries are considerably less likely to offer this category, by as much as 35.8%. Furthermore, sister

[28] Results from an alternative specification are presented in the Appendix. In those estimations, the variables for sister distilleries' product offerings are replaced with simple dummy variables. Dummy variables are equal to one so long as the distillery has at least one sister distillery that offers a whisky in age category A.

distilleries do not appear to influence product positioning with respect to this particular age group, as only one of the coefficients for sister distilleries is statistically significant, and only at the 10% level. This is somewhat surprising, as the opposite results were found for the younger category, but signals that the relationship between sister distilleries is not necessarily constant across all age categories.

For the oldest age category, $A \geq 20$, distillery size does not seem to have an effect. The initial specification does find that Speyside distilleries are generally less likely to produce older whiskies, as are those with sister distilleries selling a product of comparable age. But after controlling for regional classifications of sister distilleries, these effects lose their statistical significance. As with other specifications, no significant relationship between sister distilleries in other regions is apparent.

Whiskies without an age statement represent a fourth class of whisky that are harder to differentiate because their attributes are not as clearly communicated as those with an age statement, but almost certainly represent a more heterogeneous mix of whiskies because of the various production techniques distillers use to sell their product in lieu of an age statement. Even the NAS category appears to be sensitive to sister distilleries' product lines, however, exactly as seen in the case of whiskies 12 years and younger. The effect appears to be concentrated on sister distilleries in the same region, and there is no apparent effect from sister distilleries in other regions. In general, Speyside distilleries are much less likely to produce NAS whiskies.

Across age categories, capacity only seems to have a significant effect on mid-age whiskies. For other age categories, I cannot reject the null hypothesis $\beta = 0$, that small distilleries are just as likely as large distilleries to produce a whisky of the given age category. However, the qualities of Speyside whiskies and sheer number of distilleries in that region do appear to be factors in the NAS age category and both categories for whiskies at least 13 years old. In general, the signs of coefficients for sister distilleries product variety confirmed expectations: $\delta < 0$. When a parent company owns multiple distilleries, and one of those distilleries produces a particular age category, the company's other distilleries are less likely to produce a whisky in the same category. This supports the hypothesis that merging companies have incentive to reposition their products so as to avoid competition amongst themselves. Furthermore, because $\delta_1 < \delta_2$, it appears that parent companies are aware of the substitutability of their different brands, and that some brands are better substitutes for others as measured by their regional styles. If these distilleries were independently owned and had no sister distilleries, it seems many of them would be more likely to offer a greater variety of whiskies.

Table 4. *2SLS Results – Product Positioning*

	$A \leq 12$		$13 \leq A \leq 19$		$A \geq 20$		NAS	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Intercept	0.753***	0.690***	0.490***	0.472***	0.590***	0.537***	0.922***	0.871***
	(0.092)	(0.087)	(0.098)	(0.095)	(0.101)	(0.108)	(0.085)	(0.085)
$Capacity$	-0.013	-0.016	0.063***	0.069***	0.034	0.031	-0.013	0.008
	(0.022)	(0.020)	(0.024)	(0.024)	(0.024)	(0.025)	(0.020)	(0.021)
$Speyside$	-0.094	0.138	-0.358***	-0.301***	-0.265***	-0.175	-0.329***	-0.276***
	(0.101)	(0.130)	(0.103)	(0.108)	(0.101)	(0.162)	(0.093)	(0.095)
$Sister$	-0.019**		-0.015		-0.049***		-0.030***	
	(0.010)		(0.013)		(0.018)		(0.010)	
$SisterRegion$		-0.106***		-0.099*		-0.091		-0.217***
		(0.026)		(0.056)		(0.069)		(0.068)
$SisterNon$		0.019		0.016		-0.016		0.034
		(0.016)		(0.022)		(0.034)		(0.022)
Cragg-Donald F-statistic	994.124	86.049	1330.214	32.171	480.688	10.975	532.034	12.813

***Significant at the 1% level; **5%; *10%
Standard errors are reported in parentheses

4.B. Product Line Size

In addition to measuring distilleries' product variety by determining what types of whisky each produces, product variety can also be a measure of the total number of products offered (Alexander, 1997; Berry and Waldfogel, 2001; Watson 2009). In the context of single malt Scotch, the interaction between market organization and firms' product line size is expected to produce several effects. First, segments of the market with a higher concentration of distilleries (e.g., Speyside) are likely more competitive than areas with only a few distilleries, so product lines should be smaller for these distilleries than those in areas where distilleries have fewer rivals (Watson, 2009). Additionally, for instances in which multiple companies are owned by a parent firm, companies have incentive to position their products to increase the distance between them, thereby decreasing competition between their brands but in a manner that sometimes affords them the opportunity to consolidate product line size while also discouraging entry of other firms (Alexaner, 1997; Berry and Waldfogel, 2001). Thus, it is expected that the effect of a sister distillery's product line size will be non-positive.

The distribution of product line sizes by distillery is summarized in Figure 2. The mean number of products per firm is about 5.6, with a median of 3 and the full range between 0 and 25 for all distilleries. Distilleries that do not offer any products are not inactive. Rather, they represent a number of distilleries whose single malts go solely towards blended whiskies.[29] In theory, nothing prevents distilleries from putting all of their whisky towards single malt Scotch bottlings. But in practice, developing single malt

[29] Most distilleries use some of their total stock towards blends, or sell their single malt stock to third parties who blend it with whiskies from other distilleries.

brands may be price-prohibitive, and their products may have to compete in segments of the market that are well-saturated with other competitors who could serve as close substitutes. Almost all of the distilleries that offer no single malt Scotch are located in the Highland and Speyside regions, the two most populated regions.

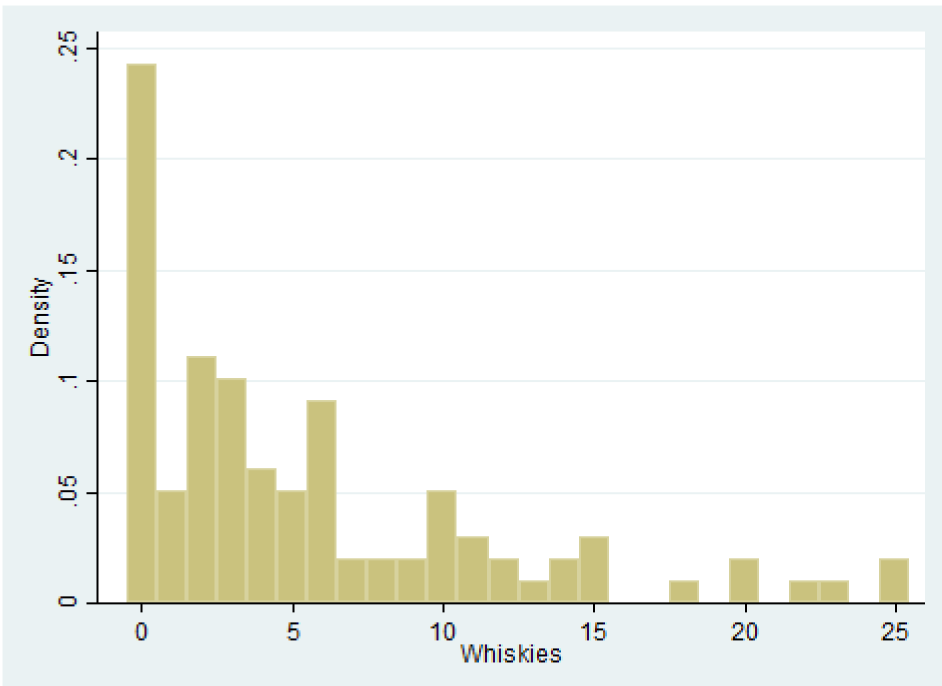

Figure 2. *Number of Expressions per Distillery*

A negative binomial model is used to estimate distilleries' product line size due to the discrete nature of the data and the apparent overdispersion that would invalidate a Poisson model. As in section 4.A., the model takes two forms to account for regional variability in the effects of sister distilleries. The basic negative binomial model has the conditional mean

(3) $\mu_i = \exp(\alpha + \beta Capacity_i + \gamma Speyside_i + \delta SisterCapacity_i)$

and conditional variance

(4) $\mu_i + \theta\mu_i^2$.

Similar to the estimations in section 4.A., it is expected that capacity will have a positive effect on distillery's product line size ($\beta > 0$), because larger distilleries are able to produce more output, but distilleries may produce a greater variety of expressions to avoid over-saturating particular market segments. Distilleries in the Speyside region are likewise expected to have smaller product line sizes ($\gamma < 0$) because there are so many active distilleries in this region producing comparable products. Ideally, the estimation would also include the number of products offered by sister distilleries, similar to the specification in 4.A., however dealing with endogenous variables in a count model is decidedly more difficult than instrumenting in the case of a linear probability model. Accordingly, total sister distillery capacity is used as a proxy for sister distilleries' product line size. Because capacity should increase product line size, δ should have the same sign as the unbiased coefficient for sister distillery product line size ($\delta < 0$), however, its size and the corresponding average marginal effects obviously cannot be interpreted literally. Even so, if capacity is endogenous, it should demonstrate the general effect (if present) of product line sizes from sister distilleries.

A second specification distinguishes between sister distilleries located in the same region and those located elsewhere:

$$(5)\ \mu_i = \exp(\alpha + \beta Capacity_i + \gamma Speyside_i + \delta_1 SisterCapRegion_i + \delta_2 SisterCapNon_i),$$

and using the same conditional variance as in Eq. (4). The same hypotheses apply with respect to capacity and regional differences ($\beta > 0, \gamma < 0$). Generally, it is expected that the product line sizes of sister distilleries in the same region will have a larger effect than the effect of sister distilleries in other regions. It is also likely that regional differences determining the substitutability of brands would also lead to the relationship $\delta_1 < \delta_2$ for product lines. However, it is unclear whether this will necessarily translate when using total sister capacity by region as controls, hence I make no such assumption a priori.

Results for the negative binomial models in Eq. (3) and Eq. (5) are summarized in Table 5. The estimate for θ is positive and statistically different from zero, confirming that the conditional mean is overdispersed, thus a negative binomial distribution is more appropriate than a Poisson in this context. Both estimations report that capacity has a positive effect on product line size, though the effect is apparently small and only statistically significant at the 10% level. When sister capacity is not regionally differentiated, the coefficient for the Speyside dummy variable is negative and statistically significant, though controlling for regional differences in sister distilleries appears to diminish this effect, possibly because of a correlation between distilleries located in the Speyside region and incidence of sister distillery capacity. The initial estimation finds that sister distilleries' total capacity does have a negative effect on product line size. Furthermore, by differentiating sister capacity based on regions, results for Eq. (5) find a negative and statistically significant effect for both regional variables, and that the effect for sister distilleries in the same region appears to be larger.

Table 5. *Negative Binomial Results – Product Line Size*

	(3)	(5)
Constant	2.080***	1.993***
	(0.182)	(0.182)
$Capacity$	0.077*	0.081*
	(0.044)	(0.043)
$Speyside$	-0.417*	-0.151
	(0.231)	(0.273)
$SisterCapacity$	-0.016***	
	(0.003)	
$SisterCapRegion$		-0.028***
		(0.007)
$SisterCapNon$		-0.011**
		(0.004)
θ	0.828***	0.790***
	(0.171)	(0.164)

***Significant at the 1% level; **5%; *10%
Standard errors are reported in parentheses

Section 3 identified conditions which led me to reject the hypothesis that distilleries are price-takers. Results from sections 4.A. and 4.B. further characterize the relationship between brands, market power, and substitutability. The market for Speyside whiskies appears to be the most competitive segment of the market, with the fewest opportunities to introduce products without fierce competition due to nearby substitutes. Distilleries in the Speyside region are generally less likely to produce a whisky of any particular age category, and also tend to produce fewer whisky expressions overall. Furthermore, regional differences suggest limitations of the substitutability between brands, and this is reflected in the product line decisions of distilleries owned by parent companies. Generally, parent companies seem to be inclined to structure their brands' product lines so as to avoid offering close substitutes of their own goods. This leads to instances in which sister distilleries are less likely to offer whiskies of a particular age if a sister distillery in the same region already produces such a whisky. But this is usually not the case if the sister distillery is located in another region. In that case, product variety offered by a sister distillery in a different region appears to have no impact. Additionally, the overall size of product lines also appears to be affected by sister distilleries' product line sizes. This negative relationship implies that increased ownership among a smaller number of parent companies has led to a market equilibrium in which many distilleries offer fewer products than they would were they independently owned. Because of brand substitutability, they have carefully positioned and limited the total number of products to best suit the needs of the parent company.

5. Conclusion

This paper presents an analysis of the production decisions of whisky distilleries. Production models for goods such as wine and timber are inadequate when trying to study the production of whisky. The dynamic production models in the forestry literature focus on the increasing quantity of timber available over time, whereas the quantity of casked whisky decreases over time and simultaneously improves in quality. Wine similarly improves with age, though it tends to be characterized as having a unique optimal maturation age, and consumers can easily continue the maturation process after purchase. Whisky distilleries, on the other hand, typically do not age all of their whisky to a uniform age, but will instead bottle amounts after different years of maturation.

I find that the decision to produce multiple maturation ages of whisky can be consistent with price-taking firms, and identify natural upper- and lower-bounds for the price ratios of whiskies from a single distillery. Because the evaporation rate of casked whisky is known, I am able to use the upper-bound and observed retail prices to estimate the minimum discount rate of twenty-four distilleries. I am also able to use the lower-bound to calculate maximum discount rates and, for several instances, the minimum marginal bottling costs of distilleries. I find that the minimum discount rates for many distilleries are actually quite high, and that maximum discount rates are illogically low. This suggests that the single malt Scotch whisky industry is not perfectly competitive, in spite of the presence of so many active distilleries.

I also perform an analysis of distilleries' product lines to determine how general product characteristics, distillery size, and the presence of a parent firm affects product variety. These analyses find several crucial elements that define the market for single malt Scotch. First, distillery size does not impact the producer's decision to produce whiskies of a particular age. However, distilleries in the Speyside region are less likely to produce most ages of whisky because there is a high concentration of distilleries in the Speyside region who all produce a similar product. While every distillery is said to produce a unique whisky, it appears that single malts are comparable enough by consumers' standards, such that firms do not readily offer particular products if there is too much competition. This is further reflected in the effects of sister distilleries on a producer's decision to offer a particular product: if two distilleries are owned by a parent company, each will offer products so that they are not close substitutes to those of the sister distillery. These results are generally only significant for sister distilleries in the same region, however, which further supports the notion that not all single malt Scotches are perfect substitutes for one another. Finally, market concentration also has an effect on the total number of products a distillery is likely to offer. If the distillery is owned by a parent company with multiple distilleries, the distillery is likely to have a smaller product line size were it independently owned. These findings suggest that distilleries produce multiple ages of whisky to capture a greater market share, but that the existence of close substitutes from competing brands causes the distillery to decrease its overall product variety.

Based on my analysis, it is extremely unlikely that the Scotch whisky industry is perfectly competitive despite the almost one hundred distilleries active in Scotland. Instead, consumers differentiate between many of the brands based on product characteristics, which creates an opportunity for market power. Firms' market power would then explain why distilleries produce multiple ages of whisky, and evidence based on regional styles and parental ownership support this. Future work should consider further developing a model with monopolistic competition, product quality, costs, and consumers' utility maximization problem to better estimate maturation decisions of multiproduct firms using a structural approach; regrettably there is insufficient data as of yet to do so. I also expect that the aged spirits considered could likewise be broadened to include other spirits, including Irish whiskey, bourbon, and even rum.

APPENDIX A

Table A1. *Regression Results* – Dependent Variables are Balancing Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
Wind	-0.535**	-0.611**+	-0.466**+	-0.411**+
NGprice	6.437**	4.953**+	6.041**+	7.226**+
Load	2.087**	1.287**+	1.304**+	1.380**+
Nuclear	-3.365**	-2.025**+	-2.163**+	-2.732**+
Intercept	-48.309**	-29.471**+	-29.245**+	-28.124**+
Month*YR FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
Houston				
Wind	0.043	-0.555**+	-0.372**+	-0.307**+
NGprice	7.713**	5.249**+	6.435**+	7.684**
Load	2.356**	1.309**+	1.328**+	1.416**+
Nuclear	-4.706**	-2.157**+	-2.417**+	-3.047**+
Intercept	-60.030**	-31.198**+	-31.149**+	-30.424**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
South				
Wind	0.347**	-0.447**+	-0.316**+	-0.224**+
NGprice	7.607**	5.187**+	6.321**+	7.626**
Load	2.356**	1.265**+	1.282**+	1.365**+
Nuclear	-4.013**	-2.221**+	-2.353**+	-2.831**+
Intercept	-64.771**	-29.940**+	-29.989**+	-30.423**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
West				
Wind	-1.567**	-1.334**+	-0.770**+	-0.564**+
NGprice	5.564**	4.810**+	5.920**+	7.197**+
Load	2.109**	1.446**+	1.348**+	1.396**+
Nuclear	-3.517**	-1.496**+	-2.177**+	-2.749**+
Intercept	-41.254**	-35.580**+	-29.404**+	-28.128**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO

*Significant at the 5% level; **Significant at the 1% level
+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A2. *Regression Results* – Dependent Variables are Balancing Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
Wind	-4.767**	1.481**+	-0.880**+	-2.913**+
NGprice	5.338**	4.281**+	5.214**+	6.007**+
Load	2.810**	1.689**+	1.630**+	1.793**+
Nuclear	-3.456**	-2.351**+	-1.993**+	-2.022**+
Intercept	-56.705**	-38.242**+	-32.470**+	-32.568**+
Month*YR FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
Houston				
Wind	-5.232**	-1.620**+	-3.725**+	-4.924**
NGprice	4.663**	4.224**+	5.044**+	5.755**+
Load	3.197**	1.724**+	1.674**+	1.853**+
Nuclear	-4.968**	-2.563**+	-2.362**+	-2.397**+
Intercept	-58.992**	-35.109**+	-28.962**+	-30.380**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
South				
Wind	-0.942	-1.282**	-2.840**+	-3.490**+
NGprice	4.138**	4.129**	4.925**	5.789**+
Load	3.169**	1.669**+	1.611**+	1.777**+
Nuclear	-5.165**	-2.598**+	-2.417**+	-2.347**+
Intercept	-59.883**	-34.422**+	-27.900**+	-30.145**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
West				
Wind	-8.444**	-0.065	-1.411**+	-3.282**+
NGprice	5.116**	4.195**+	5.120**	5.992**+
Load	2.717**	1.776**+	1.662**+	1.793**+
Nuclear	-4.681**	-2.830**+	-2.039**+	-2.055**+
Intercept	-44.462**	-37.027**+	-32.251**+	-32.015**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A3. *Regression Results* – Dependent Variables are Balancing Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
Wind	-0.487**	-0.624**+	-0.440**+	-0.356**+
NGprice	6.467**	4.912**+	6.054**+	7.289**+
Load	1.934**	1.159**+	1.214**+	1.229**+
Nuclear	-3.052**	-1.159**+	-1.951**+	-2.384**+
Intercept	-41.165**	-24.228**+	-24.934**+	-23.824**+
Month*YR FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
Houston				
Wind	0.100*	-0.556**+	-0.346**+	-0.247**+
NGprice	7.750**	5.241**+	6.457**+	7.734**
Load	2.172**	1.174**+	1.218**+	1.235**+
Nuclear	-4.332**	-1.174**+	-2.159**+	-2.648**+
Intercept	-51.390**	-26.033**+	-26.564**+	-25.436**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
South				
Wind	0.394**	-0.444**+	-0.290**+	-0.158**+
NGprice	7.637**	5.175**+	6.342**+	7.676**
Load	2.206**	1.136**+	1.183**+	1.202**+
Nuclear	-3.706**	-1.800**+	-2.105**+	-2.463**+
Intercept	-55.886**	-24.742**+	-25.470**+	-25.586**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
West				
Wind	-1.523**	-1.333**+	-0.720**+	-0.518**+
NGprice	5.593**	4.675**+	5.930**+	7.253**+
Load	1.967**	1.272**+	1.250**+	1.247**+
Nuclear	-3.227**	-1.269**+	-1.913**+	-2.348**+
Intercept	-35.428**	-27.212**+	-25.194**+	-23.969**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A4. *Wind and Price Volatility* (β_i when Dependent Variables are Volatility Measure A, no month by year fixed effects)

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
σ_A^2	570.271**	-0.450+	2.521+	38.993+
σ_A^3	0.109**	0.057**+	0.052**+	0.117**
σ_A^4	0.618**	-0.037**+	0.032+	0.660**
Houston				
σ_A^2	1192.389**	1.712+	9.847*+	81.098**+
σ_A^3	0.161**	0.073**+	0.090**+	0.189**
σ_A^4	0.979**	-0.032**+	0.094+	0.943**
South				
σ_A^2	1742.594**	1.219+	9.813**+	73.682**+
σ_A^3	0.183**	0.082**+	0.102**+	0.195**
σ_A^4	1.003**	-0.039**+	0.081	1.014**
West				
σ_A^2	507.986**	9.417**+	37.141**+	96.974**+
σ_A^3	0.032	-0.007+	0.009	0.066*
σ_A^4	0.288*	-0.032*+	0.002+	0.412*

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A5. *Wind and Price Volatility* (β_i when Dependent Variables are Volatility Measure B, no month by year fixed effects)

	OLS	Q(.25)	Q(.5)	Q(.75)
North				
σ_B^2	694.159**	-0.559+	0.734+	64.769**+
σ_B^3	0.006	0.002	0.004	0.010
σ_B^4	0.573**	-0.011+	0.345**+	1.220**+
Houston				
σ_B^2	1154.316**	-0.380+	2.282+	106.489**+
σ_B^3	0.010	0.009	0.004	0.009
σ_B^4	0.797**	0.020	0.492**+	1.748**+
South				
σ_B^2	1485.572**	-0.308+	1.689+	107.868**+
σ_B^3	0.006	0.005	-0.005	0.004
σ_B^4	0.818**	0.037	0.515	1.735**+
West				
σ_B^2	471.999**	2.203**+	10.047**+	73.056**+
σ_B^3	-0.025*	-0.049**+	-0.020*	-0.012
σ_B^4	0.644**	0.230**+	0.584**	1.203**+

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A6. *Regression Results* – Dependent Variables are Ancillary Service Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
Reg.-Up				
Wind	0.190**	0.008+	0.029+	0.208**
NGprice	2.169**	0.786**+	1.511**+	2.841**+
Load	0.507**	0.193**+	0.308**+	0.425**+
Nuclear	-0.761**	-0.707**	-0.528**+	0.030+
Intercept	-14.874**	-2.157**+	-7.586**+	-16.881**+
Month*YR FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
Reg.-Down				
Wind	0.398**	0.303**+	0.339**+	0.387**
NGprice	2.219**	1.211**+	1.726**+	2.628**+
Load	-0.270**	-0.064**+	-0.119**+	-0.207**+
Nuclear	0.340**	-0.184**+	0.046	0.344**
Intercept	3.496**	0.230+	0.488**+	0.810*+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
Resp. Res.				
Wind	0.718**	0.261**+	0.352**+	0.598**+
NGprice	2.665**	0.882**+	1.791**+	3.183**+
Load	0.589**	0.176**+	0.330**+	0.494**+
Nuclear	-0.648**	-0.596**	-0.436**+	-0.266**+
Intercept	-24.211**	-4.475**+	-12.288**+	-21.778**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO
Non-Spinning				
Wind	3.080**			0.102**
NGprice	0.060			0.138**
Load	1.309**			0.354**
Nuclear	-4.913**			-0.908**
Intercept	-47.792**			-5.451**
Month*Yr. FE	NO	NO	NO	NO
Hour FE	NO	NO	NO	NO

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A7. *Regression Results* – Dependent Variables are Ancillary Service Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
Reg.-Up				
Wind	-3.511**	-0.945**+	-1.964**+	-3.085**
NGprice	0.310**	-0.021+	-0.037+	0.091
Load	0.847**	0.345**+	0.535**+	0.734**+
Nuclear	-1.429**	-0.636**+	-0.816**+	-0.939**+
Intercept	-7.409**	0.375**+	-1.348**+	-3.011**+
Month*YR FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
Reg.-Down				
Wind	-0.750	-0.663**	-0.294	-0.279
NGprice	0.543**	0.236**+	0.408**+	0.808**+
Load	-0.247**	-0.030**+	-0.082**+	-0.159**+
Nuclear	0.563**	-0.169+	0.085+	0.508**
Intercept	8.612**	5.113**+	5.681**+	5.811**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
Resp. Res.				
Wind	-4.175**	-0.616*+	-2.404**+	-3.592**
NGprice	1.894**	0.105*+	0.309**+	0.801**+
Load	0.927**	0.360**+	0.554**+	0.753**+
Nuclear	-1.442**	-0.806**+	-0.933**+	-1.182**
Intercept	-19.549**	-3.152**+	-5.127**+	-7.055**+
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO
Non-Spinning				
Wind	-2.195			0.063
NGprice	2.061**			0.034
Load	1.957**			0.339**
Nuclear	-1.489**			-0.167
Intercept	-66.843**			-4.870**
Month*Yr. FE	YES	YES	YES	YES
Hour FE	NO	NO	NO	NO

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

Table A8. *Regression Results* – Dependent Variables are Ancillary Service Prices

	OLS	Q(.25)	Q(.5)	Q(.75)
Reg.-Up				
Wind	0.219**	0.044**+	0.088**+	0.210**
NGprice	2.188**	0.773**+	1.482**+	2.540**+
Load	0.413**	0.073**+	0.146**+	0.259**+
Nuclear	-0.568**	-0.504**	-0.262**+	0.138+
Intercept	-13.723**	-0.304**+	-5.625**+	-11.865**+
Month*YR FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
Reg.-Down				
Wind	0.378**	0.285**+	0.338**+	0.352**
NGprice	2.206**	1.192**+	1.700**+	2.415**+
Load	-0.208**	-0.074**+	-0.108**+	-0.129**+
Nuclear	0.214**	-0.121**+	0.038+	0.257**
Intercept	1.097*	-0.329*+	-0.877**	-1.416**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
Resp. Res.				
Wind	0.754**	0.287**+	0.415**+	0.633**+
NGprice	2.688**	0.896**+	1.837**+	3.182**+
Load	0.473**	0.061**+	0.173**+	0.265**+
Nuclear	-0.411**	-0.385**	-0.193**+	0.046+
Intercept	-22.432**	-2.778**+	-9.981**+	-17.751**+
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES
Non-Spinning				
Wind	3.268**			0.110**
NGprice	0.100			0.087**
Load	0.939**			0.245**
Nuclear	-3.981**			-0.562**
Intercept	-47.679**			-4.289**
Month*Yr. FE	NO	NO	NO	NO
Hour FE	YES	YES	YES	YES

*Significant at the 5% level; **Significant at the 1% level

+Quantile coefficient is significantly different than the corresponding OLS estimate at the 5% level

APPENDIX B

Table B1. *2SLS Results – Product Positioning* (Sister Variables are Dummy Variables)

	$A \leq 12$		$13 \leq A \leq 19$		$A \geq 20$		NAS	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Intercept	1.059***	0.698***	0.664***	0.344	0.991***	0.595***	1.318***	0.858***
	(0.222)	(0.127)	(0.209)	(0.429)	(0.230)	(0.164)	(0.203)	(0.138)
$Capacity$	0.014	0.002	0.075***	0.082**	0.055**	0.044	0.015	0.012
	(0.029)	(0.022)	(0.027)	(0.035)	(0.028)	(0.030)	(0.027)	(0.024)
$Speyside$	-0.088	0.091	-0.353***	-0.280*	-0.252**	0.071	-0.307***	-0.110
	(0.109)	(0.127)	(0.107)	(0.150)	(0.112)	(0.180)	(0.114)	(0.141)
$Sister$	-0.599*		-0.338		-0.755**		-0.811***	
	(0.328)		(0.287)		(0.299)		(0.298)	
$SisterRegion$		-0.502***		-0.458		-0.747***		-0.630***
		(0.169)		(0.623)		(0.262)		(0.195)
$SisterNon$		0.154		0.323		0.005		0.155
		(0.153)		(0.632)		(0.170)		(0.177)
Cragg-Donald F-statistic	10.167	11.801	10.445	0.298	10.730	6.948	11.397	10.036

***Significant at the 1% level; **5%; *10%
Standard errors are reported in parentheses

WORKS CITED

Alexander, P.J., 1997. Product variety and market structure: A new measure and a simple test. Journal of Economic Behavior & Organization 32: 207-214.

Angrist, J.D., Pischke, J.-S., *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton: Princeton University Press, 2009. Print.

Arroyo, J.M., Conejo, A.J., 2000. Optimal response of a thermal unit to an electricity spot market. IEEE Transactions on Power Systems 15, 1098-1104.

Ashenfelter, O., 2008. Predicting the Quality and Prices of Bordeaux Wine. *The Economic Journal* 118: F174-F184.

Ashenfelter, O., Hosens, D., 2010. The effect of mergers on consumer prices: evidence from five mergers on the enforcement margin. Journal of Law and Economics 53(3), 417-466.

Benitez, L.E., Benitez, P.C., van Kooten, G.C., 2008. The economics of wind power with energy storage. Energy Economics 30, 1973-1989.

Berry, S.T., Waldfogel, J., 2001. Do Mergers Increase Product Variety? Evidence from Radio Broadcasting. The Quarterly Journal of Economics 116(3): 1009-1025.

Bird, L., Bolinger, M., Gagliano, T., Wiser, R., Brown, M., Parsons, B., 2005. Policies and market factors driving wind power development in the United States. Energy Policy 33, 1397-1407.

Carlton, D.W., Dana, Jr., J.D., 2008. Product Variety and Demand Uncertainty: Why Markups Vary with Quality. *The Journal of Industrial Economics* 56(3): 535-552.

Chang, J., Ummels, B.C., van Sark, W.G.J.H.M., den Rooijen, H.P.G.M., Kling, W.L., 2009. Economic evaluation of offshore wind power in the liberalized Dutch power market. Wind Energy 12, 507-523.

Chen, C.-L., Lee, T.-Y., Jan, R.-M., 2006. Optimal wind-thermal coordination dispatch in isolated power systems with large integration of wind capacity. Energy Conversion and Management 47, 3456-3472.

Chen, C.-M., Yeh, C.-Y., Hu, J.-L., 2011. Influence on uncertain demand on product variety: evidence from the international tourist hotel industry in Taiwan. *Tourism Economics* 17(6): 1275-1285.

Chernozhukov, V., Hong, H., 2002. Three-step censored quantile regression and extramarital affairs. Journal of American Statistical Association 97, 872-882.

Clò, S., Cataldi, A., Zoppoli, P., 2015. The merit-order effect in the Italian power market: The impact of solar and wind generation on national wholesale electricity prices. Energy Policy 77, 79-88.

Cormio, C., Dicorato, M., Minoia, A., Trovato, M., 2003. A regional energy planning methodology including renewable energy sources and environmental constraints. Renewable and Sustainable Energy Reviews 7, 99-130.

Cutler, N.J., Boerema, N.D., MacGill, I.F., Outhred, H.R., High penetration wind generation impacts on spot prices in the Australian national electricity market. Energy Policy 39, 5939-5949.

Delarue, E., De Jonghe, C., Belmans, R., D'haeseleer, W., 2011. Applying portfolio theory to the electricity sector: energy versus power. Energy Economics 33, 12-23.

Delarue, E., Luickx, P., D'haeseleer, W., 2009. The actual effect of wind power on overall electricity generation costs and CO2 emissions. Energy Conversion and Management 50, 1450-1456.

Denny, E., O'Malley, M., 2006. Wind generation, power system operation, and emissions reduction. IEEE Transactions on Power Systems 21, 341-347.

Di Cosmo, V., Malaguzzi Valeri, L., 2012. Relation between wind and electricity prices in a deregulated market: the case of Ireland. Working Paper, 31st USAEE/IAEE North American Conference.

Di Cosmo, V., Malaguzzi Valeri, L., 2014. The incentive to invest in thermal plants in the presence of wind generation. Energy Economics 43, 306-315.

Draganska, M., Mazzeo, M., Seim, K., 2009. Beyond plain vanilla: Modeling joint product assortment and pricing decisions. Quantitative Marketing and Economics 7(2): 105-146.

Feldstein, M.S., 1971. The error forecast in econometric models when the forecast-period exogenous variables are stochastic. Econometrica 39, 55-60.

Gandhi, A., Froeb, L., Tschantz, S., Werden, G.J., 2008. Post-Merger Product Repositioning. The Journal of Industrial Economics 56(1): 49-67.

Garciá-González, J., Barquín, J., 2000. Self-unit commitment of thermal units in a competitive electricity market. IEEE Power Engineering Society Summer Meeting.

Goodhue, R.E., LaFrance, J.T., Simon, L.K., 2009. Wine Taxes, Production, Aging and Quality. *Journal of Wine Economics* 4(1): 65-83.

Green, R., Vasilakos, N., 2010. Market behavior with large amounts of intermittent generation. Energy Policy 38, 3211-3220.

Holttinen, H., Tuhkanen, S., 2004. The effect of wind power on CO2 abatement in the Nordic Countries. Energy Policy 32, 1639-1652.

Jacobsen, H.K., Zvingilaite, E., 2010. Reducing the market impact of large shares of intermittent energy in Denmark. Energy Policy 38, 3403-3413.

Jaeger, E., 1981. To Save or Savor: The Rate of Return to Storing Wine. *Journal of Political Economy* 89(3): 584-592.

Jónsson, T., Pinson, P., Madsen, H., 2010. On the market impact of wind energy forecasts. Energy Economics 32, 313-320.

Karki, R., Billinton, R., 2001. Reliability/Cost implications of PV and Wind Energy Utilization in Small Isolated Power Systems. IEEE Transactions on Energy Conversion 16, 368-373.

Karki, R., Billinton, R., 2004. Cost-effective wind energy utilization for reliable power supply. IEEE Transactions on Energy Conversion 19, 435-440.

Katzenstein, W., Apt, J., 2009. Air emissions due to wind and solar power. Environmental Science & Technology 43, 253-258.

Ketterer, J.C., 2014. The impact of wind power generation on the electricity price in Germany. Energy Economics 44, 270-280.

Krasker, W.S., 1979. The Rate of Return to Storing Wines. *Journal of Political Economy* 87(6): 1363-1367.

Lehtilä, A., Pirilä, P., 1996. Using an energy systems optimization model to support policy planning in Finland. Energy Policy 24, 805-819.

Li, F., Kuri, B., 2005. Generation scheduling in a system with wind power. IEEE/PES Transmission and Distribution Conference & Exhibition: Asia and Pacific (Dalian, China)

Liik, O., Oidram, R., Keel, M., 2003. Estimation of real emissions reduction caused by wind generators. International Energy Workshop, 24-26 June 2003, IIASA, Laxenburg, Austria.

Lowery, P.G., 1996. Generating unit commitment by dynamic programming. IEEE Transactions of Power Apparatus and Systems 85, 422-426.

Luickx, P.J., Delarue, E.D., D'haeseleer, W.D., 2010. Impacts of large amounts of wind power on the operation of an electricity generation system: Belgian case study. Renewable and Sustainable Energy Reviews 14, 2019-2028.

Lund, H., 2006. Large-scale integration of optimal combinations of PV, wind and wave power into the electricity supply. Renewable Energy 31, 503-515.

Lund, H., Münster, E., 2003a. Management of surplus electricity-production from a fluctuating renewable-energy source. Applied Energy 76, 65-74.

Lund, H., Münster, E., 2003b. Modelling of energy systems with a high percentage of CHP and wind power. Renewable Energy 28, 2179-2193.

Mazzeo, M.J., 2002. Product Choice and Oligopoly Market Structure. The RAND Journal of Economics 33(2), 221-242.

Milstein, I., Tishler, A., 2011. Intermittently renewable energy, optimal capacity mix and prices in a deregulated market. Energy Policy 39, 3922-3927.

Morthorst, P.E., 2003. Wind power and the conditions at a liberalized power market. Wind Energy 6, 297-308.

Parsons, B., Milligan, M., Zavadil, B., Kirby, B., Dragoon, K., Caldwell, J., 2004. Grid Impacts of Wind Power: A Summary of Recent Studies in the United States. Wind Energy 7, 87-108.

Richter, C.W., Sheblém, G.B., 2000. A profit-based unit commitment GA for the competitive environment. IEEE Transactions on Power Systems 15, 715-721.

Salo, S., Tahvonen, O., 2002. On Equilibrium Cycles and Normal Forests in Optimal Harvesting of Tree Vintages. *Journal of Environmental Economics and Management* 44: 1-22.

Salo, S., Tahvonen, O., 2003. On the economics of forest vintages. *Journal of Environmental Dynamics & Control* 27: 1411-1435.

Shcherbakova, A., Kleit, A., Blumsack, S., Cho, J., Lee, W., 2014. Effect of increased wind penetration on system prices in Korea's electricity markets. Wind Energy 14, 1469-1482.

Sensfuß, R., Ragwitz, M., Genoese, M., 2008. The merit-order effect: a detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany. Energy Policy 36, 2086-3094.

Sioshansi, R., Hurlbut, D., 2010. Market protocols in ERCOT and their effect on wind generation. Energy Policy 38, 3192-3197.

Stoft, S., 2002. *Power system economics: designing markets for electricity*. Piscataway, NJ: IEEE Press; New York, NY: Wiley-Interscience, 2002.

Sweeting, A., 2010. The effects of mergers on product positioning: evidence from the music radio industry. RAND Journal of Economics 41(2): 372-397.

Ummels, B.C., Gibescu, M., Pelgrum, E., Kling, W.L., Brand, A.J., 2007. Impacts of wind power on thermal generation unit commitment and dispatch. IEEE Transactions on Energy Conversion 22, 44-51.

Uusivuori, J., Kuuluvainen, J., 2005. The Harvesting Decisions When a Standing Forest with Multiple Age-Classes Has Value. *American Journal of Agricultural Economics* 87(1): 61-76.

Valenzuela, J., Mazumdar, M., 2001. Making unit commitment decisions when electricity is traded at spot market prices. IEEE Power Engineering Society Winter Meeting.

Voorspools, K., D'haeseleer, W., 2006. Impact assessment of using wind power. Solar Energy 80, 1165-1178.

Watson, R., 2009. Product Variety and Competition in the Retail Market for Eyeglasses. The Journal of Industrial Economics 57(2): 217-251.

Wohlgenant, M.K., 1982. Inventory Adjustment and Dynamic Winery Behavior. *American Journal of Agricultural Economics* 62(2): 222-231.

Woo, C.K., Zarnikau, J., Moore, J., Horowitz, I., 2011a. Wind generation and zonal-market price divergence: Evidence from Texas. Energy Policy 39, 3928-3938.

Woo, C.K., Horowitz, I., Moore, J., Pacheco, A., 2011b. The impacts of wind generation on the electricity spot-market price level and variance: The Texas experience. Energy Policy 39, 3939-3944.

Wood, A.J. Wollenberg, B.F., 1984. *Power Generation, Operation, and Control*. New York, NY: John Wiley & Sons.

Zarnikau, J., Hallett, I., 2008. Aggregate industrial energy consumer response to wholesale prices in the restructured Texas electricity market. Energy Economics 30, 1798-1808.