Wind Integration and Carbon Taxes in ISO New England

Gabriel Gonzalez Sutil and Thomas Braun*

2020
December

Abstract

The increased integration of wind electricity into existing grids is poised to catalyze significant market changes. This study assesses the prospective impacts of wind power integration on operational costs, generator profits, consumer expenses, and market prices within the ISO New England system, employing a simplified one-bus system model. We solve the system unit commitment problem via mixed-integer linear programming, and we compute prices using the Envelope Theorem and dual variables. We use nine scenarios, ranging from 0 MW to 20,000 MW of installed wind capacity, alongside an additional scenario featuring a $50 per ton carbon tax. Scenarios were selected using agglomerative clustering. Our findings indicate that policies promoting greater wind penetration should be accompanied by a carbon tax. Without the tax, increasing wind capacity results in reduced operational costs, lower median prices, and diminished average industry profits. Conversely, the scenario with a carbon tax shows higher prices passed on to consumers, resulting in a surplus transfer from consumers to the industry but overall higher welfare. We recommend that policymakers consider using tax revenues to directly compensate for energy price increases caused by the tax.

1 Introduction

Electricity is the world's fastest-growing form of end-use energy consumed, and has been so for many decades already (Conti et al., 2016). Although electricity is a clean and relatively safe form of energy, its generation and transmission has direct consequences on the environment. Such impacts, however, largely depend on the specific types of electric power plants (US Energy Information Administration, 2019).

Over the past century, the main energy sources used for electricity generation have been fossil fuels. In 2019, thermal facilities, including coal, gas, and oil power plants, accounted for 72.16% of the total installed capacity in ISO New England. In addition, thermal generation accounted for 59.86% of total generation in the ISO New England market in the same year. Thermal facilities obtain thermal energy by burning fossil fuels, a process which releases combustion gases into the atmosphere. In 2019, thermal generation was the source of 22,424 thousand tons of carbon dioxide (CO_2), 1.6 million pounds of

*School of International and Public Affairs,Columbia University. For more information please send an email to gg2718@columbia.edu. We are thankful to Bolun Xu and students at the Energy Optimization and Economics course at Columbia University. We are also thankful to Jeffrey Shrader for valuable comments.

sulfur dioxide (SO_2), and 6.9 million pounds of Nitrogen oxides (NO_x) emissions in the ISO New England market. Some plants also produce solid residue[1] which are mixed with water creating the so-called sludge which is stored in retention ponds. Several of these ponds have burst and already caused extensive damage and pollution downstream.

Nearly all thermal plants' combustion byproducts have negative effects on the environment and human health. CO_2 is a greenhouse gas, and its emission in the atmosphere directly contributes to global warming. Another common gas emitted during the combustion of fossil fuels is SO_2, which causes acid rain and also worsens respiratory illnesses and heart diseases, particularly in children and the elderly (US Energy Information Administration, 2019). NO_x is a precursor of ground-level ozone which irritates and damages the lungs. In addition, some heavy metals such as mercury are released to the atmosphere and are hazardous to human and animal health. Finally, particulate matter (PM) results in hazy conditions in cites and, coupled with ozone, contributes to asthma and chronic bronchitis. Note that very small, or fine PM, is also believed to cause emphysema and lung cancer (US Energy Information Administration, 2019).

Concerns about air quality and climate change have raised the interest in alternative power generating solutions such as wind power. CO_2 emissions in year t can be expressed as

$$E_t = EI_{jt} \cdot S_{jt} \cdot Q_t$$

where EI_{jt} is the carbon intensity of technology j measured in $t\mathrm{CO}_2/MWh$, Q_t is the market level generation in period t measured in MW, and S_{jt} is the share generated by technology j measured in %. In this simplified set up, clean technologies are carbon neutral. Therefore, a decrease of the market share of carbon-emitting technologies relative to cleaner ones such as wind generation would decrease aggregate CO_2 emission levels.

Increasing the share of wind electricity can be expected to trigger additional changes in the market. First, it will affect the overall operating structure of the market as there will be a resource reallocation due to the fact that the integration of wind generation will displace other plants. Second, it will generate changes in the electricity equilibrium price since wind generation has a marginal cost assumed to be equal to zero. These effects are not only related to price levels but also to the price variance. Changes in average price are important since they have direct impacts on generators' profits and the price paid by consumers. Thus, an increase of prices will have distributional effects in the short-run. Moreover,the real-time price variance is an important variable because, when price risk is high, investors would be discouraged and would require higher risk premiums (Winkler, Gaio, Pfluger, & Ragwitz, 2016). Similarly, largely fluctuating day-ahead prices mean there are some critical issues in the whole system resource planning (i.e. resource inadequacy). This project will analyze the effect of wind integration on the day-ahead market in order to asses the economic ramifications wind integration may have.

The present project analyses the prospective effects of wind integration on electricity prices and operating costs in the ISO New England system, under the assumption of a one-

[1] Bottom ash includes the largest particles that collect at the bottom of the combustion chamber, and fly ash is the smaller and lighter particulates.

bus system model.[2] In our assessment, we will consider different levels of installed wind capacity under a variety of characteristic scenarios in relation to both electricity demand and wind conditions typical for New England. Ultimately, our report aims at providing a recommendation to local policy makers.

2 Theoretical Background

Policy makers have attempted to support an increase in the share of wind electricity generation, which directly affects the unit commitment solution and the day-ahead market clearing conditions. While electricity demand is inelastic in the short-term, any change in the supply curve will necessarily affect the generation of each unit, the market price level, and the profitability of different units. However, the magnitude of these effects is not clearly defined theoretically, since the aggregated effect will ultimately depend on the specific generation mix, the precise market structure as well as other characteristics. Indeed, many different effects could be working in opposite directions, and the final net effect will depend on the relative strength of each of these effects. In this sense, studying the effect of wind power integration on the day-ahead market structure is a non-trivial question.

2.1 Operating Costs and Social Welfare

If we start by fixing a instantaneous period t, the objective of a system operator is to maximize social welfare in this period t. Note that this implies the system operator should not favor one party over the other. Assume there is a representative consumer which has a benefit from consuming an electricity quantity D_t. Let $B_t(D_t)$ be the related benefit. By definition,

$$B_t(D_t) = \int_0^{D_t} p_t(D_t) dD_t$$

where $p_t(D_t)$ is the demand function. Now, let $C_{i,t}(g_{i,t})$ be the cost of generating quantity $g_{i,t}$ with generator i. By definition, and assuming a quadratic functional form, the cost function can be expressed as

$$C_{i,t}(g_{i,t}) = b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t}$$

where a_i and b_i are dispatch cost coefficients for generator i, NLC_i is the no-load cost for generator i, SUC_i is the start-up cost for generator i. This set up assumes that the parameters of the cost function are the same for different time periods. Moreover, $u_{i,t}$ is a binary variable which takes value 1 if generator i is on during time period t and otherwise zero. Finally, $v_{i,t}$ is a binary variable which takes value 1 if generator i is turned on at the start of period t, otherwise zero.

In this context, total cost in the fixed period t is given by $C_t(g_t) = \sum_{i=1}^N C_{i,t}(g_{i,t})$ where $g_t = \sum_{i=1}^N g_{i,t}$. Thus, for the fixed period t the system operator's objective is to maximize

[2] This latter assumption is supported by the fact that ISO New England has minor congestion issues on its network.

$$\int_0^{D_t} p_t(D_t)dD_t - \sum_{i=1}^{N}(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t})$$

which represents the difference of the consumers' benefit and the operating cost. Since the demand function is assumed to be inelastic and "passive", in the sense that it is fixed for each period t, the first term is constant. So, solving the previous problem is equivalent to minimizing the operating cost function by choosing $g_{i,t}$, $u_{i,t}$, and $v_{i,t}$. Given that there are $t \in \{1, 2, ..., T\}$ periods, the system operator's objective is to solve

$$\min_{\{g_{i,t}, u_{i,t}, v_{i,t}\}} \sum_{t=1}^{T}\sum_{i=1}^{N}(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t})$$

Therefore, understanding the effect of wind integration on operating costs is an important variable to analyze since it will define the effect of wind integration on social welfare. Note that in the previous analysis and contrary to below sections, the choice of reserve has not yet been taken into account.

2.2 Profit and Consumer Cost

In the previous section, we explained that demand is inelastic and constant for each instantaneous time period. Hence, analyzing the effect of wind integration on consumer costs is relevant since the electricity cost will have distributional consequences in the short-term. The cost to consumers is given by the electricity and reserve prices in this market, and the amount of reserve and electricity consumed.

Note that the analysis of the price is not only relevant for consumers but also for generators. Indeed, generator's i profit is defined as:

$$\Omega_i = \sum_{t=1}^{T}[\pi_t g_{i,t} - C_{i,t}(g_{i,t})]$$

where π_t is the price in period t. Therefore, firms' profits not only depend on the level of generation for each period t but also on the market price for that period. Analyzing the profit is important to understand overall the industry dynamic. In fact, firms with negative profits will sooner or later exit the market, while new technologies with positive profit perspectives will be attracted to the market.

2.3 Changes in Prices

Wind power can generate a partial displacement of fossil fuels in the market, which could generate downward pressure on prices.[3] Most power systems dispatch electricity such that

[3] The effect of wind integration on average prices also includes the effect investments have on prices. Wind integration requires the construction of wind facilities in regions where demand is supplied by other types of facilities (Tra, 2016). Hence, firms inside the market will have to install and operate renewable technologies. Then, prices might tend to increase inside the market due to investment costs (Fischer, 2010). For the purpose of the present analysis, we will consider that wind power is already installed, and we will abstract from any investment costs.

units with high marginal costs operate only when demand is very high. In general, wind energy has very competitive marginal costs (virtually equal to zero) and can bid at very low prices displacing fossil generation (Gelabert, Labandeira, & Linares, 2011). As a consequence, there will be a reduction of the average cost on the market(Tra, 2016), and the average electricity price should decrease in a competitive generation market (Woo, Horowitz, Moore, & Pacheco, 2011).

Kaufmann and others (2016) call this effect the "Demand Reduction Induced Price" (DRIP). In particular, the displacement of non-renewable electricity sources will reduce demand for fossil fuels. Thus, non-renewable generation average costs of operation decrease, and electricity prices subsequently fall (Fischer, 2010). The DRIP effect is closely related to what some authors call the merit-order effect (Winkler et al., 2016).

The magnitude of the merit-order effect is conditional on the possibility to substitute costly technologies with cheaper ones. In fact, when fossil electricity is operating at high marginal costs, the introduction of cheaper renewable sources might generate a strong reduction in prices. However, when the difference of costs is smaller, the reduction of marginal costs won't be significative. This project will quantify the average price reduction in the ISO New England market.

Moreover, the price variance is an important variable to analyze. Electricity demand shows strong seasonality behavior, and prices typically exhibit strong daily, weekly and monthly fluctuations (Wozabal, Graf, & Hirschmann, 2016). As a consequence, units with high costs operate only when demand is very high (Kaufmann & Vaid, 2016). In this context, the integration of wind power may increase price variance since wind is volatile and it is hardly adjustable to electricity demand, so when wind is not available we have to ramp-up more expensive units (Ketterer, 2014). Indeed, a large variance of the day-ahead market represents a resource inadequacy and critical issues in the resource planning of the system.

The most convenient way authors found to analyze the price variance into a price model is by using the concept of residual demand (Wozabal et al., 2016), that it is the difference, in our case, between wind variable production and demand level. Then, we can define a distribution function of the residual demand, with a positive density on the real numbers. This distribution can be associated with a distribution function of prices. When wind generation decreases by any decrease of wind flow, the residual demand has to be covered by conventional facilities in very short period of time (Winkler et al., 2016). Given that these sources usually have higher marginal costs, prices will have to adjust, affecting the day-ahead price variance.

In this context, a higher share of wind source would have two direct effects on the distribution function of the residual demand. First, the residual demand would be shifted to a zone where prices are smaller (zone where the supply curve is flat) as wind electricity usually has very low marginal costs. Second, the variance of the density of the residual demand will change, and the price variance will consequently be affected. This means that the effect of wind integration on the price variance will depend on the shapes of both the density functions and the supply curve, as well as on the effect higher share of wind electricity will have on these two (Wozabal et al., 2016). The latter relationship will strongly depend on the location, the natural environment and the specific technology used to generate wind electricity (Kaufmann & Vaid, 2016).

3 Formal Model

In this project, our aim is to to assess the impact of wind capacity integration into the ISO New England network on operating costs, generators' profits, consumer costs, and prices. For this, we program a unit commitment problem where the objective is to determine which units of a given system have to be started and stopped for each period of time, as well as how much the started units have to generate given exogenous demand, reserve requirements, and wind profiles under a variety of installed capacity scenarios.

Let $t \in \{1, 2, ..., T\}$ be the time period index and $i \in \{1, 2, .., N\}$ be the generator index. Since ISO New England has only minor congestion issues, we will consider that there are no transmission limits within the network and that our problem can therefore be solved as a single-bus model. In other words, and following section 2.1, our target is to minimize the following objective function:

$$\sum_{t=1}^{T}\sum_{i=1}^{N}(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t}) \tag{1}$$

which represents the total operating cost of the power system. Recall that $u_{i,t}$ is a binary variable which takes value 1 if unit i is on during period t. We assume that all the units are initially off. In addition, $v_{i,t}$ is a binary variable which takes value 1 if unit i is turned on at the start of period t and otherwise 0. So, $v_{i,t}$ is defined as

$$v_{i,t} - z_{i,t} = u_{i,t} - u_{i,t-1} \tag{2}$$

where $z_{i,t} \in \{0, 1\}$ takes value 1 if generator i is turned off at the start of period t and otherwise 0. Note that a generator cannot simultaneously be on *and* off, which is imposed by condition (3):

$$v_{i,t} + z_{i,t} \leq 1 \tag{3}$$

Constraints (2) and (3) represent the generator's start-up and shut-down logic. In addition, we know that, once a unit is running, it may not be shut down immediately. The minimal time it needs to remain on is the so-called "minimum up time". In mathematical terms, if $u_{i,t} = 1$ and $t_i^{up} < t_i^{up,min}$, then $u_{i,t+1} = 1$. Similarly, once a unit has just been shut down, it may not be immediately re-started. In mathematical terms, if $u_{i,t} = 0$ and $t_i^{down} < t_i^{down,min}$, then $u_{i,t+1} = 0$. In order to introduce these dynamics into the problem we will impose the generator minimum up and minimum down time constraints as follow:

$$\sum_{\tau=\max\{t-Tup_i+1,1\}}^{t} v_{i,\tau} \leq u_{i,t} \tag{4}$$

$$\sum_{\tau=\max\{t-Tdn_i+1,1\}}^{t} z_{i,\tau} \leq 1 - u_{i,t} \tag{5}$$

where Tup_i is the minimum up time of generator i and Tdn_i is the minimum down time of generator i. While the initial version of the instructions indicated a lower summation bound of "$\min\{t - Tup_i + 1, 1\}$", our team figured out that this would yield incorrect results, especially since the *min* function would generate negative indices for the summation and correctly suggested to introduce a *max* function instead.

A generator must also satisfy its generation limit constraints:

$$G_i^{min} \cdot u_{i,t} \leq g_{i,t} \leq G_i^{max} \cdot u_{i,t} \tag{6}$$

where G_i^{min} and G_i^{max} are the minimum and maximum generation limits for generator i. Note that we only need to satisfy this constraint for those units which are on in period t. Thus, $g_{i,t} \in \{0, [G_i^{min}, G_i^{max}]\}$.

In addition, we know that variations in electrical output of any given unit are technically constrained. In fact, changes in generation for each individual unit have to remain below a certain ramp rate constraint. Let RR_i be the ramp rate constraint of generator i. Then, we must always have

$$-RR_i \leq g_{i,t} - g_{i,t-1} \leq RR_i + G_i^{min} \cdot v_{i,t} \tag{7}$$

Note that the previous condition allows to override the ramp rate limits when the unit has just been turned on. In our model, we further assume that wind generation has zero marginal cost and can be curtailed. More specifically, let w_t be the wind power integrated to the grid in moment t, W the installed wind capacity, and α_t the time-variant wind capacity factor at time t. For any t, wind generation must satisfy

$$w_t \leq \alpha_t W \tag{8}$$

In this context, the power balance has to be satisfied meaning that the amount generated from wind capacity and the amount generated using all the conventional generators has to be equal to the demand in each period. In other words, at each point in time, we must have:

$$\sum_{i=1}^{N} g_{i,t} + w_t = D_t \tag{9}$$

In addition, the system we model will be subject to the so-called "3+5" reserve rule, meaning that the reserve amount in each period t must at least be equal to 3% of the demand in period t and 5% of integrated wind. Since the system reserve will be provided by facilities that are online, the amount generated and used for reserve in each unit i cannot exceed the total available capacity at those units. Finally, the reserve cannot exceed the ramp rate of the units. All in all, the system reserve constraints are given by

$$\sum_{i=1}^{N} r_{i,t} \geq 3\% \cdot D_t + 5\% \cdot w_t \tag{10}$$

$$r_{i,t} \leq G_i^{max} \cdot u_{i,t} - g_{i,t} \tag{11}$$

$$r_{i,t} \leq RR_i \tag{12}$$

In this problem, the continuous decision variables are the generation of unit i during period t, the hourly reserve generator provided by generator i during time period t, and the wind power generation during time period t (i.e. $g_{i,t}$, $r_{i,t}$, and w_t respectively). The binary decision variables are $u_{i,t}$, $v_{i,t}$, and $z_{i,t}$. Therefore, the mixed-interger linear programming problem we solve in this project is explicitly defined by:

$$
\begin{aligned}
\min_{u_{i,t},v_{i,t},z_{i,t},g_{i,t},r_{i,t},w_t} \; & \sum_{t=1}^{T}\sum_{i=1}^{N}(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t}) \\
\text{s.t. } & G_i^{min}\cdot u_{i,t} \le g_{i,t} \le G_i^{max}\cdot u_{i,t}, \;\; -RR_i \le g_{i,t} - g_{i,t-1} \le RR_i + G_i^{min}\cdot v_{i,t} \\
& v_{i,t} - z_{i,t} = u_{i,t} - u_{i,t-1}, \;\; v_{i,t} + z_{i,t} \le 1 \\
& \sum_{\tau=\tau_1}^{t} v_{i,\tau} \le u_{i,t}, \;\; \sum_{\tau=\tau_2}^{t} z_{i,\tau} \le 1 - u_{i,t} \\
& 0 \le w_t \le \alpha_t W \\
& \sum_{i=1}^{N} g_{i,t} + w_t = D_t \\
& \sum_{i=1}^{N} r_{i,t} \ge 3\% \cdot D_t + 5\% \cdot w_t, \;\; r_{i,t} \le G_i^{max}\cdot u_{i,t} - g_{i,t}, \;\; 0 \le r_{i,t} \le RR_i \\
& u_{i,t} \in \{0,1\}, \;\; v_{i,t} \in \{0,1\}, \;\; z_{i,t} \in \{0,1\}
\end{aligned}
$$

where $\tau_1 = \max\{t - Tup_i + 1, 1\}$ and $\tau_2 = \max\{t - Tdn_i + 1, 1\}$. Note that, in this formulation, the constraint $g_{i,t} \le G_i^{max}\cdot u_{i,t}$ is not necessary as it is automatically embedded within $r_{i,t} \le G_i^{max}\cdot u_{i,t} - g_{i,t}$ (given that $r_{i,t} \ge 0$). However, we are going to be as explicit as possible in the formulation in this case since adding it changes trivially the time to achieve an optimal value. The reason, as we explain later, is that Gurobi does a presolve reduction procedure in which it already simplifies the problem as much as possible. So, the set of constraints will be simplified by Gurobi.

3.1 Prices and Dual Variables

As previously explained, the price of energy is a key variable for our analysis. The price of reserve is also an important variable we need to consider. The prices are obtained by computing the dual variables associated with the corresponding constraints. In order to understand this statement, we need to rely on the Envelope Theorem. Let

$$V(g_{i,t}^*, u_{i,t}^*, v_{i,t}^*, z_{i,t}^*, r_{i,t}^*, w_t^*) = \min \sum_{t=1}^{T}\sum_{i=1}^{N}[b_i(g_{i,t}^*)^2 + a_i g_{i,t}^* + NLC_i u_{i,t}^* + SUC_i v_{i,t}^*]$$

be the value function and λ be the dual variable associated with the power balance constraint (9). In addition, assume that μ_R is the dual variable associated with the minimum reserve constraint (10). If Re_t denotes the minimum reserve requirements in the latter constraint, we may rewrite this constraint as:

$$\sum_{i=1}^{N} r_{i,t} \ge Re_t$$

So, by the Envelope theorem:

$$\frac{\partial V(\cdot)}{\partial D_t} = \frac{\partial L(\cdot)}{\partial D_t} = \lambda \quad \text{and} \quad \frac{\partial V(\cdot)}{\partial Re_t} = \frac{\partial L(\cdot)}{\partial Re_t} = \mu_R$$

where $L(\cdot)$ is the Lagrangian of the problem. The dual variable λ represents the least additional cost of servicing one more MW of load. Similarly, the dual variable μ_R represents the least additional cost of one additional MW of reserve requirement. Thus, the price of energy and reserve are respectively given by the dual variables λ and μ_R according to:

$$\sum_{i=1}^{N} g_{i,t} + w_t = D_t : \lambda \quad \text{and} \quad \sum_{i=1}^{N} r_{i,t} \geq 3\% \cdot D_t + 5\% \cdot w_t : \mu_R$$

where $\lambda \geq 0$ and $\mu_R \geq 0$. Note that $\mu_R > 0$ only when the inequality constraint is binding (due to complementary slackness conditions).

One challenge in this set up is that we are solving a mixed-integer problem. Since the dual variables are associated with the gradient of the Lagragian function, we won't be able to recover the dual variables from the mixed-integer problem. In the implementation section we explain how we calculate these dual variables in a two-step approach, by relying on both an economic dispatch problem and a classical (i.e. continuous) linear optimization problem.

4 Implementation

In this section, we briefly describe the data and explain our proposed method to choose a representative set of scenarios for demand and wind. In addition, we explain how we defined wind scenarios, the way the model was formulated in Matlab, the choice of the solver, and the method to calculate the dual variables.

4.1 Input Data

Our model relies on three datasets which correspond to the data used in the paper "An 8-Zone Test System Based on ISO New England Data: Development and Application" (Krishnamurthy, Li, & Tesfatsion, 2016). These files characterize conventional energy generation (supply-side characteristics), energy demand over a 4-day period (demand-side characteristics) and wind capacity profiles over a 4-day period.

4.1.1 Generators

The generator level data includes the location, maximum and minimum capacity in MW, fuel type, minimum up- and down-time in hours, ramp rate in MW/h, and three coefficients calibrating the dispatch cost.[4]

The energy mix of our power system is provided in Figure 1. As expected, the energy mix is in agreement with the "Test System" described in Figure 5 from (Krishnamurthy et al., 2016).

[4]For a detailed description, please refer to the appendix at the end of the document.

Figure 1: Installed capacity for ISO New England by fuel type

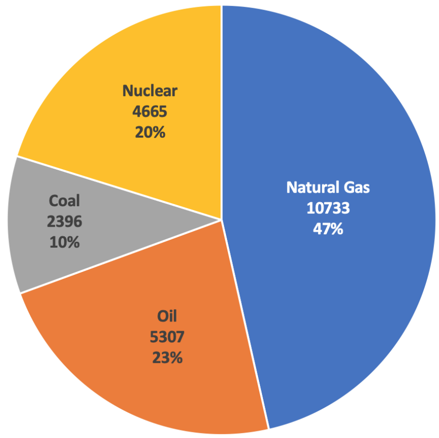

Note: Fuel oil includes RFO, coal includes bituminous and sub-bituminous, and natural gas is made of NGA1, NGA4, NGIR, NGLN, NGMN, NGPN, NGT1, NGT2, NGT4, and NGTN. This plot does not take the NEMASSBOST area into account.

It appears that almost half (47%) of the installed capacity corresponds to natural gas power generation. By contrast, oil (23%), nuclear (20%) and coal (10%) account for the remainder of the market.

The fuel type used by a particular facility has significant consequences on its technical flexibility and cost structure. For instance, while nuclear facilities display significant capacities, very large large start-up costs, long minimum up- and down-time characteristics, natural gas facilities are much smaller and more flexible. Conversely, once operating, nuclear facilities display a significant advantage compared to natural gas plants in terms of variable costs. Compared to any other fuel type, natural gas further displays the unique advantage of typically allowing for important ramp rates. More generally, we summarize the key characteristics of generators in ISO New England classified by conventional fuel type in Table 1, where the median values of our generator data characteristics have been computed.[5]

[5] Given the presence of outlier values in the dataset, average values would mis-represent our dataset and we therefore propose to look at median values. Moreover, we observe a strong coherence with values for cost coefficients as presented in Table III of (Krishnamurthy et al., 2016). The only inconsistency with Table III is related to the b coefficient for fuel oil, where (Krishnamurthy et al., 2016) claim an upper limit of 0.0342 which is inferior to our median value. In fact, each of generator 18, 19, 21, 26, 27, 28, 29 and 30 has a b coefficient strictly above 0.0342.

Table 1: Median characteristics of the 76 generators considered in the present report

	(*unit*)	Natural gas	Coal	Nuclear	Fuel oil
Capacity	(MW)	188	150	881	400
Minimum generation	(MW)	38	60	705	140
Minimum up-time	(h)	6	16	24	16
Minimum down-time	(h)	5	13	48	11
Ramp rate	MW/h	400	120	120	120
Start-up cost	(\$)	2500	13873	900000	82500
No-load cost	(\$/h)	250	745	1250	6000
Dispatch cost coefficient a	(\$/MWh)	57	18	10	176
Dispatch cost coefficient b	(\$/(MW)^{2}h)	0.0046	0.0001	0.0002	0.0344

As already explained, our assessment does not consider investment costs and will consider that these 76 generators are all available to use as of the first period of our unit commitment problem. No scenarios are further considered for the above conventional generators. Conversely, as explained in the following two subsections, the datasets for both demand and wind may strongly fluctuate according to a variety of scenarios.

4.1.2 Demand

The demand data at our disposition covers demand in each of the eight zones for ISO New England, over a 4-day period (96 hours) and under 90 different scenarios. The time resolution of the demand data set is one hour, and the demand displays strong daily seasonality patterns, achieving its minimum during the night and its maximum during the day. Among all scenarios and time-steps, the extreme demand values range from only 9,019 MW to as much as 17,765 MW. Figure 2 illustrates the entire demand dataset by representing a superposition of all 90 time series scenarios.

Figure 2: Demand profiles over the 4-day time window

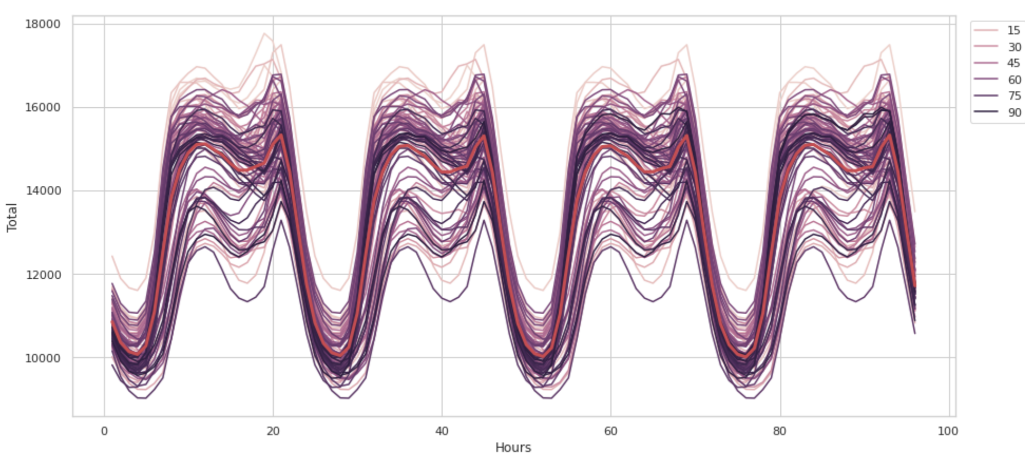

Note: Superposition of the 90 demand profiles (see colored labels in the legend) over the 4-day time window under investigation. The average demand profile is plotted in red.

4.1.3 Wind

Similarly to the previous dataset for demand, the wind capacity dataset corresponds to hourly wind capacity factors (with values between 0 and 1), over a 4-day time-frame, under 90 different scenarios.

The distribution of wind capacity factor data shows that, for 27% of all observations in the dataset, the wind capacity factor is smaller than 0.05 (see Figure 3). The amount of hours with higher wind capacity factors typically decreases with the wind capacity factor, meaning that strong wind is typically less frequent than milder wind. The dataset, however, displays one notable exception to this general rule, namely for extremely strong wind. In fact, 5% of all hourly observations for wind capacity appear to have a value above 0.95. This is likely due to the truncation, at a threshold, for any type of extreme wind situation.

Figure 3: Distribution of hourly wind capacity factors across all scenarios

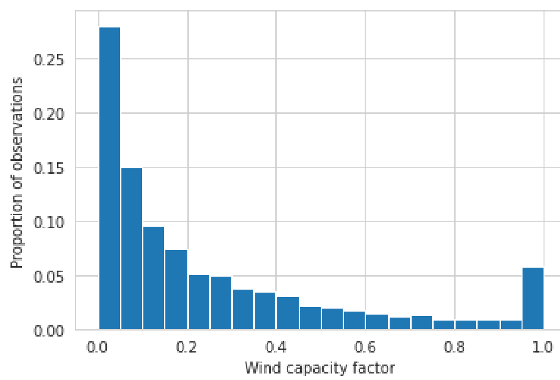

The average wind capacity factor in our data is 0.25, while the median is 0.14 due to the left-skewness of the distribution. Figure 4 presents the plot of the scenarios.

Figure 4: Wind capacity factor profiles over the 4-day time window

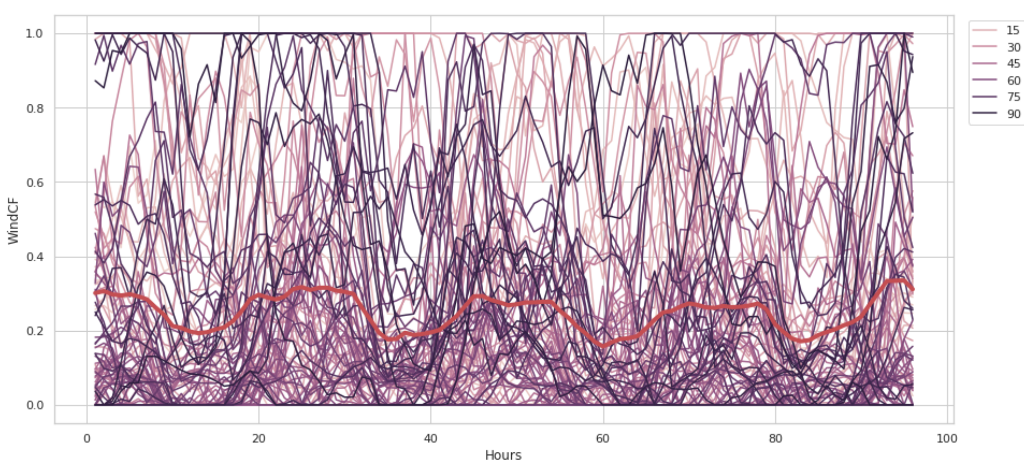

Note: Superposition of the 90 wind capacity factor profiles (see colored labels in the legend) over the 4-day time window under investigation. The average wind profile is plotted in blue.

While individual wind scenarios do not visually display as strong regularity over the 4-day period as the demand profiles do, we observe from the above chart that, on average, the wind capacity factor is typically higher during night time, and slightly drops around noon. This situation is quite unfortunate as it is, on average, in opposition to the previously identified pattern characterizing demand. Figure 5 illustrates this finding.

Figure 5: Average demand (in red) and wind capacity factor (in blue) over 4 days

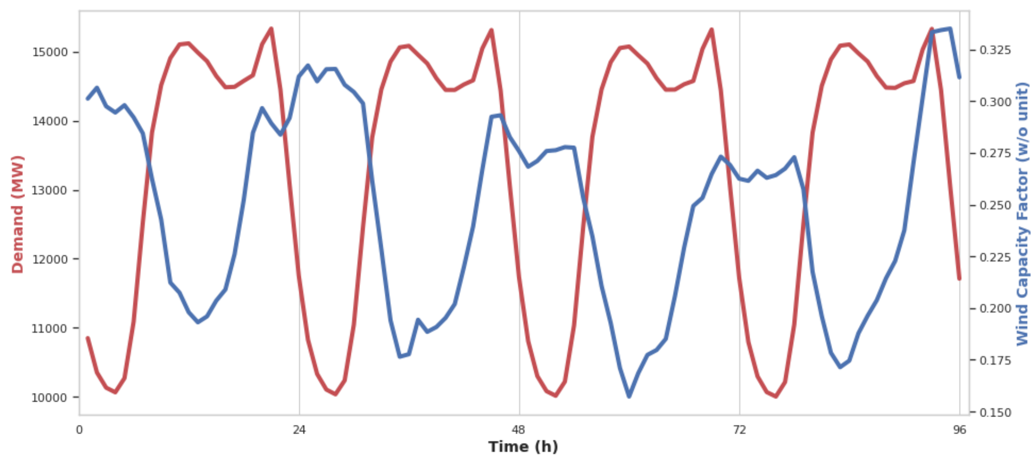

Figure 5 shows that there is an imbalance between the average demand peaks and the average wind capacity factor. This will impose large rampage requirements as we introduce wind into the system. This means that the variance of the residual demand is expected to increase, and the effect of this on the day-ahead price will depend on how flexible are the units and the generation cost of these facilities.

4.2 Selection of Representative Scenarios

Given the large number of different scenarios for wind capacity and demand data, we propose to restrict our assessment to a subset of representative scenarios. More specifically, using Python's *scipy* environment, we use a machine learning algorithm for agglomerative clustering in order to classify both our demand and wind capacity databases into three clusters each.[6] Finally, we separate the 90 scenarios into nine different clusters, by confronting the demand and wind classifications and performing a cross-classification. The classifications of demand and wind profiles are respectively summarized by dendrograms depicted in Figure 6 and 7.

[6]For this purpose, we use a complete linkage approach with euclidean affinity metric.

Figure 6: Classification of the 90 demand profiles into 3 clusters (green, red, blue)

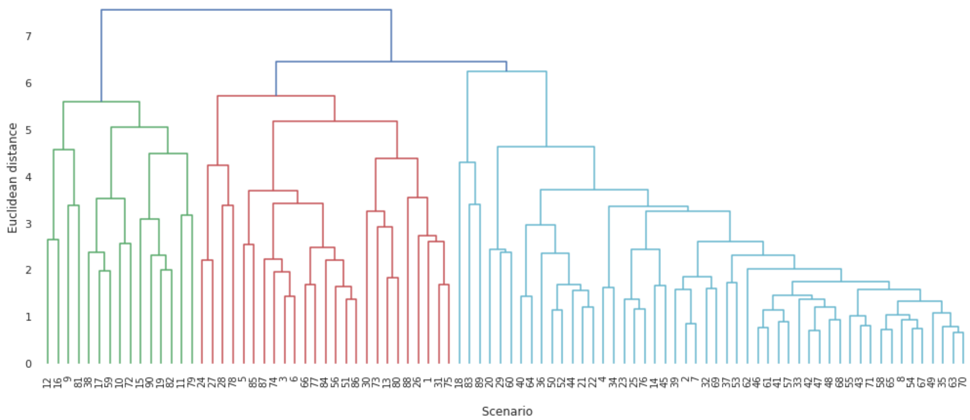

Figure 7: Classification of the 90 wind profiles into 3 clusters (green, red, blue)

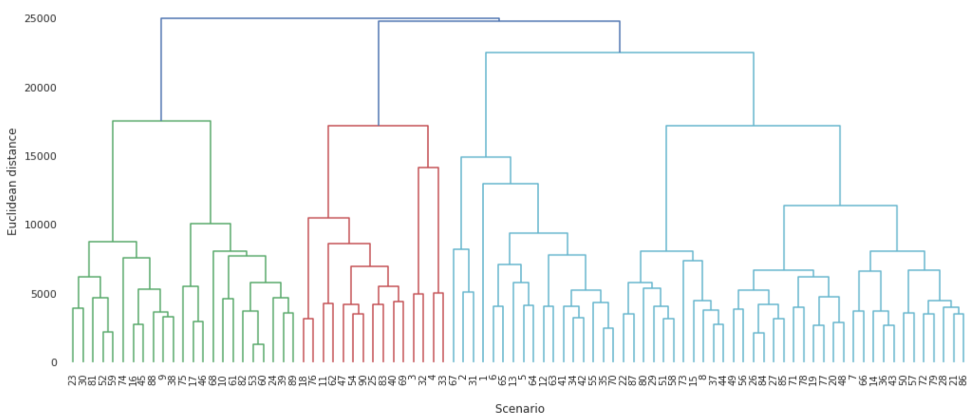

Using these dendrograms, we classify demand and wind scenarios into nine clusters as summarized in Table 2. Instead of manipulating all 90 scenarios, our classification approach suggests that we may restrict our attention to "only" nine representative scenarios, provided each of these nine scenarios is taken from a different cluster. For the purpose of our assessment, we arbitrarily propose to select the smallest scenario number within each of the nine clusters (see scenario numbers in bold). Thus, while scenario 3 appears quite unique in the dataset due to its particular combination of demand and wind data, most scenarios fall under the "blue-blue" cluster and can be effectively represented by a single scenario such as scenario 2. All in all, we thus propose to restrict our assessment to scenarios 1, 2, 3, 4, 9, 11, 12, 23 and 24.

Table 2: Cross-classification of 90 scenarios into 9 clusters

		Wind clusters as per Figure 7		
		1 (green)	2 (red)	3 (blue)
Demand clusters as per Figure 6	1 (green)	**9**, 10, 16, 17, 38, 59, 81, 82	**24**, 30, 74, 75, 88	**23**, 39, 45, 46, 52, 53, 60, 61, 68, 89
	2 (red)	**11**, 90	**3**	**4**, 18, 25, 32, 33, 40, 47, 54, 62, 69, 76, 83
	3 (blue)	**12**, 15, 19, 72, 79	**1**, 5, 6, 13, 26 − 28, 31, 51, 56, 66, 73, 77, 78, 80, 84 − 87	**2**, 7, 8, 14, 20 − 22, 29, 34 − 37, 41 − 44, 48 − 50, 55, 57, 58, 63 − 65, 67, 70, 71

Note: Cross-classification of scenarios according to both wind capacity *and* demand profiles. The 9 proposed "representative" scenarios (one per box) are highlighted in bold and (arbitrarily) correspond to the scenario with the smallest numeric identifier.

Figure 8 provides a graphical representation of the descriptive statistics (mean and standard deviation) associated to the 9 chosen scenarios. For instance, it appears from this plot that scenario 2 represents (on average) a high-demand / low-wind scenario while scenario 9 corresponds to a wind-friendly / low-demand scenario. Furthermore, it appears that in terms of variance, scenario 2 is characterized by relatively higher variability in demand while scenario 9 sees relatively more fluctuations in terms of wind capacity factor.

Figure 8: Descriptive Statistics for the nine representative scenarios

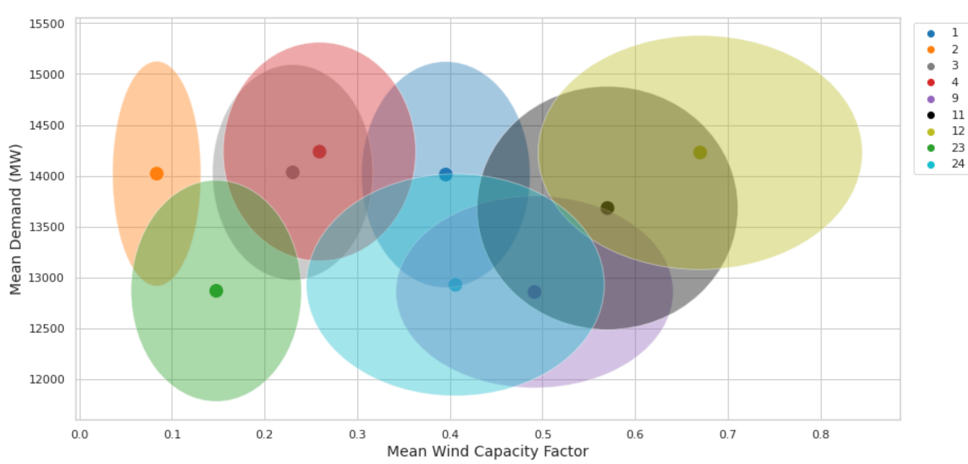

Notes: Each of the nine chosen representative scenarios is characterized by an ellipse, the centre of which has the scenario's mean wind capacity factor (on the x-axis) and mean demand (on the y-axis) as coordinates. The horizontal axis of each ellipse measures the standard deviation of corresponding scenario's wind capacity factor while the vertical axis measures the standard deviation of the corresponding scenario's demand.

Figure 8 further illustrates that, in absolute terms, all scenarios have roughly the same standard deviation in demand. By contrast, the standard deviation in wind capacity factor can vary to a significant extent from one scenario to another. For instance, the standard deviation in wind for scenario 12 appears to be 3.7 times that of scenario 2.

In order to validate our previous choice, we can verify whether our 9 chosen scenarios roughly cover the whole support of the demand and wind capacity factor. This verification step is easily performed either by observing that the nine ellipses of Figure 8 tend to cover different spaces of the "demand x wind" space, or by visually comparing Figures 9 and 10 with Figures 2 and 4 respectively. All in all, we conclude that our choice of nine scenarios seems to accurately represent the entire dataset of 90 scenarios.

Figure 9: Demand profiles for the 9 representative scenarios

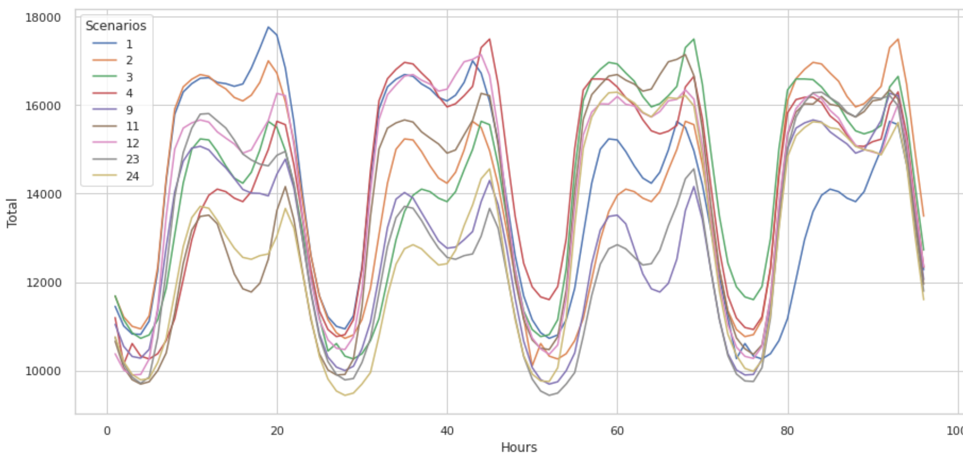

Notes: Superposition of demand profiles over the 4-day time window for the 9 representative scenarios as identified in Table 2

Figure 10: Wind capacity factor profiles for 9 representative scenarios

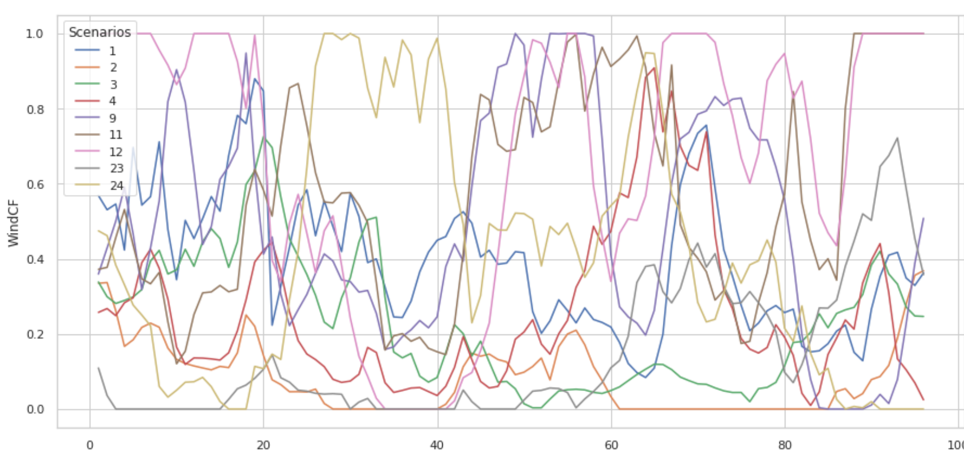

Notes: Superposition of wind capacity profiles over the 4-day time window for 9 representative scenarios as identified in Table 2

4.3 Formulation Strategy in Matlab

Our unit commitment problem described in section 3 is entirely coded in Matlab. The code uses two main functions. The first one is used to load the data, as well as to shape it for the optimization algorithm. The second one implements a two-step unit commitment problem described in details in the present section.

The optimization problem is coded following the CVX syntax. In our formulation, the objective function is coded in a matrix format which allows to include a tax by defining a specific parameter in the function. This provides us with greater flexibility for the extensions conducted in later parts of our analysis. Moreover, the constraints use vectors and loops over both time and generators whenever possible in order to keep the code as concise and readable as possible. Many of the constraints are implemented directly as described in section 3.[7]

Some of the constraints, however, require special attention. First, we divide the ramp rate constraint into two different constraints. Recall that each individual unit has to satisfy its own ramp rate constraint $-RR_i \leq g_{i,t} - g_{i,t-1} \leq RR_i + G_i^{min} \cdot v_{i,t}$ where RR_i is the ramp rate of generator i. This constraint can be separated into the maximum ramp-up rate and the maximum ramp-down rate constraints. So, for $t > 1$

$$g_{i,t} - g_{i,t-1} \leq RR_i + G_i^{min} \cdot v_{i,t} \quad \text{and} \quad g_{i,t-1} - g_{i,t} \leq RR_i$$

For $t = 1$, the relevant condition to use in the code is $g_{i,1} \leq RR_i + G_i^{min} \cdot v_{i,1}$. In a similar fashion, recall that constraint (2) establishes that $v_{i,t} - z_{i,t} = u_{i,t} - u_{i,t-1}$. Therefore, as explained in the previous paragraph, for $t = 1$, the relevant condition to use in the code is $v_{i,1} = u_{i,1}$ since $z_{i,1} = 0$.

Similarly, the generation limit constraint $G_i^{min} \cdot u_{i,t} \leq g_{i,t} \leq G_i^{max} \cdot u_{i,t}$ is divided into two conditions:

$$G_i^{min} \cdot u_{i,t} \leq g_{i,t} \quad \text{and} \quad g_{i,t} \leq G_i^{max} \cdot u_{i,t}$$

In our formulation, we will use $G_i^{min} = 0$. We analyzed the computation time of the problem when we define positive lower bounds to generation. However, the data has very tight bounds, and solving this problem is computational intensive. Since the computation time is significantly larger than when we use $G_i^{min} = 0$, we chose to use $G_i^{min} = 0$ allowing the system to be very flexible.

The problem is coded for each of our 9 representative scenarios (deterministic approach), and for different wind integration scenarios. This is performed via an inner and outer loop over both scenarios and wind integration. The outer loops define the wind integration level, while the inner loop defines a given scenario. So, for each combination of inner and outer loops, the unit commitment problem is solved. We use wind integration levels from 0 MW (i.e. no installed wind capacity) to 20,000 MW (i.e. substantial installed wind capacity) with a used 1,000 MW incremental step.

In this setup, it is important to store the variables which are obtained from solving the problem in each combination of inner and outer loops. In order to be consistent and for ease of manipulation, we store the variables in arrays of four dimensions. Indeed, each of our

[7]For more details, please refer to the code.

output matrices can be defined according to *generator's ID* x *time period* x *scenario number* x *installed wind capacity*. For quantities that do not vary along one of these dimensions (e.g. energy or reserve price does not depend upon generator), our output is a sparse matrix where only the first entry along that dimension is populated. Similarly, profits are calculated at the generating unit level, so it only populates the first column. Finally, the optimized objective function value populates the first element (i.e. scalar).

As explained before, we are solving multiple mixed-integer problems. Therefore, and in order to be able to compute the dual variables, we implemente a two-step procedure. In the first step, we solve the unit commitment problem, the solution of which determines the optimal values for all of our six decision variables, in particular for the three binary variables $u^*_{i,t}$, $v^*_{i,t}$, and $z^*_{i,t}$. These values can then be used in a second step where we solve the same problem while fixing all binary variables to their values given by the solution to the first step. In other terms, while the decision variables in our first step are $\{g_{i,t}, r_{i,t}, w_t, u_{i,t}, v_{i,t}, z_{i,t}\}$, they are restricted to $\{g_{i,t}, r_{i,t}, w_t\}$ in the second step as these correspond to continuous variables. Using this problem, we are able to obtain the dual variables for the price of energy and reserve.

4.4 Solver Selection and Solution Method

The problems we solve in practice have linear objective functions (as opposed to quadratic ones as suggested in section 3; we explain this choice in the following section), no quadratic constraints and involve both continuous *and* integer decision variables. Such problems belong to the family of so-called "mixed-integer linear programming problems"[8]. For the purpose of solving our problem, we propose to rely on the solver created by Gurobi[9] and which relies on linear-programming based on a branch-and-bound algorithm.

Below we explain some of the techniques used by Gurobi to solve the problem at hand. Note that this is a high-level description using the information provided by Gurobi on its own webpage[10]. In fact, the model can include sophisticated branch variable selection techniques. The goal of these methods is to limit the size of the branch-and-bound tree that must be explored in order to reduce time requirements for solving.

The solver starts by doing a linear programming presolve reduction of the problem. Presolve refers to a collection of problem reductions. These reductions are intended to reduce the size of the problem and to tighten its formulation. The idea behind this method is to determine if multiple restrictions can be reduced to a smaller number of constraints. By applying this reduction method, the solver does not consider the integrability restrictions (i.e. capture the discrete nature of some decisions). For example, assume we have the constraint $2x_1 + 2x_2 \leq 1$ and assume that x_1 and x_2 are binary non-negative variable. The constraint can be re-written as $x_1 + x_2 \leq \frac{1}{2}$. The only way to satisfy these conditions is with $x_1 = x_2 = 0$ given the non-negativity. Then, both of these variables and this constraint can be removed from the formulation, which is thus simplified.

After the reduction is done, the solver proceeds to a linear-programming relaxation.

[8]If we allow for a quadratic objective function, the problem would become a mixed-integer quadratic linear programming problem which is more computational demanding.

[9]https://www.gurobi.com/

[10]https://www.gurobi.com/resource/mip-basics/

The problem is afterwards solved. Note that if the solution to the relaxed problem happens to satisfy all of the integrality restrictions, the solution is an optimal solution of the original mixed-integer programming. If not, the solver uses a branch-and-bound algorithm. The variable which is used to define the branching becomes a so-called "branching variable", and the branching methods allows to replace the original problem with a series of simplified problems by dividing the original problem using the "branching variables".

The branching algorithm used by Gurobi does not differ from the usual branch-and-bound algorithm. Indeed, the "branch' part of the algorithm is named after the idea that we can solve the relaxed problem at each node of a search tree, and if we get an optimal solution which does not satisfy the integrability condition, we should consider two "branches". The "bound" part relies on the fact that an integer solution to a subproblem in the branch-and-bound method leads to a bound on the optimal objective value of the original problem. Appendix II presents some details of the algorithm used by Gurobi.

In this procedure used by Gurobi, the fact to have "strong" incumbents[11] and finding them quickly can be extremely valuable. First, it may not be possible to solve a problem to optimality in a restricted amount of time. So, we want to have the best possible feasible solution at termination. Also, the better the objective value of the incumbent node, the more likely it is that the value of a linear programming relaxation will exceed it and hence lead to a node being fathomed.

In this context, the solver checks if a good integer feasible solution can be extracted in the node, even though integrality has not yet resulted due to the branching. For example, it may be that many of the variables under integer constraint have values quite close to integers. So, the solver would consider rounding some of these variables to their nearby values, fixing them to these values, solving the resulting linear programming relaxation, and repeating this procedure several times in the hopes that all integer variables will satisfy their integer constraints. If they do, and if the resulting feasible solution has a better objective value than the current incumbent, the solver replaces that incumbent and proceeds. This is an heuristic method used by Gurobi.

The last method used by the solver involves the theory of cutting planes. Without having to resort to branching, cutting planes enable to anticipate and discard fractional solutions, thereby significantly reducing the space of solutions to be explored by the solver.

The behaviors of most of the strategies and techniques described here can be adjusted using Gurobi parameters. As explained in the solver webpage, the goal in designing and building the Gurobi Optimizer has been to make the default settings be as robust as possible across a broad range of models.

4.5 Time Scale and Problem Complexity

As announced in the previous section, we opt for a linear objective function by using only the linear term of the generation cost curve. The reason for this resides in the computation capabilities and time requirements to optimize a mixed-integer quadratic program. Figure

[11] Suppose the solver solves a linear-programming relaxed problem for some node in the search tree. If it happens that all of the integrality restrictions in the original problem, the solver assigns this node the category of a permanent leaf, or "fathomed", where it does not has to branch. In addition, the solver denotes this integer solution as the "incumbent".

11 illustrates how the optimization time scales up when a linear or quadratic objective functions are respectively used as we increase the periods we are optimizing over, for 1 scenario and 1 wind integration level.

Figure 11: Computation time for our algorithm as a function of the number of periods considered

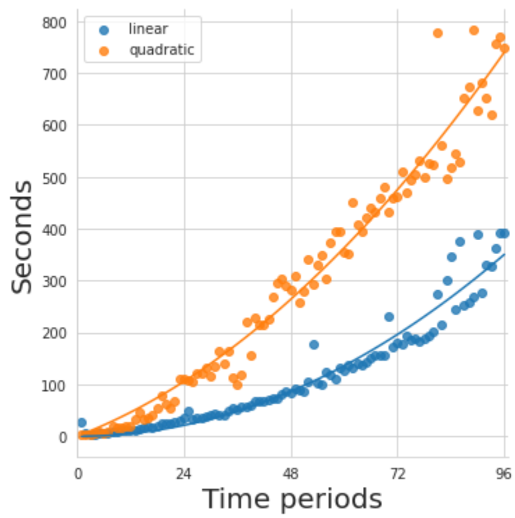

Notes: Computation time on the y-axis (in seconds) for a single scenario and a single installed wind capacity level under linear (blue scatter) or quadratic (orange scatter) objective cost function for various possible values of T between 1 and 96 periods (x-axis). Quadratic trends (lines) are fitted through ordinary least squares method. Running Ramsey's reset test on the null hypothesis that a second order polynomial would *not* fit the data better than a linear trend provides an F-value of 151 with p-value lower than 10^{-20} for the linear objective function and an F-value of 36 with p-value lower than 10^{-7} for the quadratic objective function. We may therefore strongly reject the null hypothesis and conclude that our computation time displays a quadratic complexity irrespective of the linear or quadratic character of our objective cost function.

Figure 11 shows that each optimized problem with linear objective function takes around 6 minutes for 96 periods. If we optimize 9 scenarios for 21 installed wind capacity levels, we have to solve 189 optimization problems. This implies 19 hours in total using the linear objective function. However, if we use a quadratic function the time requirement increases to 12 minutes by run and 38 hours in total.

The question is what we are losing by using a linear objective function rather than a quadratic one. First note the coefficients of the quadratic term in the objective function are small (see Table 1). Indeed, if we assume all the units are operating and are producing at their maximum capacity, the contribution of the quadratic part to the cumulated marginal cost is only 2.5% on average, and the median is 1.6%. If the plants are operating at their minimum capacity, the contribution falls down to 0.8% on average with a median of 0.3%. In this context, we check the ranking of marginal costs with and without the term associated with the quadratic term of the cost function. The ranking is different under each situation, but the order is not too much affected if we drop the quadratic term from the cost function.

In order to continue the analysis of the consequences of a linear objective function, we use a simplified version of our program described in section 4.3. Using our code, it is easy to compute the "theoretical" / "idealized" supply curve of our 76-generator network. For this purpose, we ignore all dynamic constraints (in relation to minimum up/down-time, ramp rates and start-up costs) and exclusively focus on dispatch cost parameters assuming all the units are already online. The resulting supply curve for the installed capacity, deprived from any wind capacity, is provided in Figure 12.

Figure 12: Supply curve for the 76 generators with and without the quadratic term in the dispatch cost

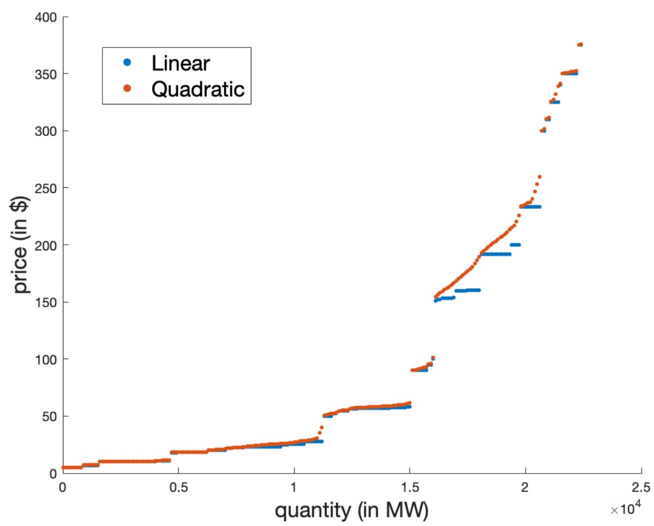

Note: Supply curve for the 76 generators in our dataset for ISO New England (excluding NEMASS-BOST) with and without the quadratic term in the dispatch cost.

Figure 12 clearly shows that dropping the quadratic term directly affects the continuity of the supply function and thereby the continuity of equilibrium prices. In fact, the two supply curves present very similar shapes but the supply curve obtained using the linear objective function is not continuous due to the constant marginal cost assumption. Therefore, we will be losing smoothness by using a linear objective function which will translate, mainly, into discrete jumps in equilibrium prices.

5 Base Case Results

The present section describes the results obtained for the base case, i.e. ignoring any carbon tax. For this analysis, we explore each of the installed wind capacity levels illustrated on the x-axis of Figure 13. Overall, we assess a complete range of installed wind capacities from 0 MW to 20,000 MW, which represents as much as 86.5% of the total installed capacity.

Figure 13: Wind Capacity

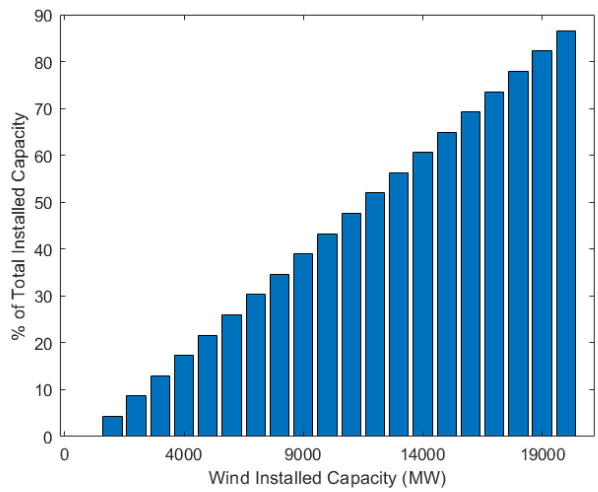

5.1 Wind Generation and Integration

The first variable of interest is wind generation which directly depends on both the wind capacity factor and the installed capacity. While the former factor is dependent on any given scenario, the latter is fixed by the chosen installed wind capacity. Indeed, wind energy generation is given by the multiplication of the capacity factor multiplied by the installed capacity. Thus, a higher capacity level will necessarily increase the wind energy generation level, but the precise level of wind integration will further depend on market considerations such as demand levels. Figure 14 illustrates our main results.

Figure 14: Wind generation - Base Case

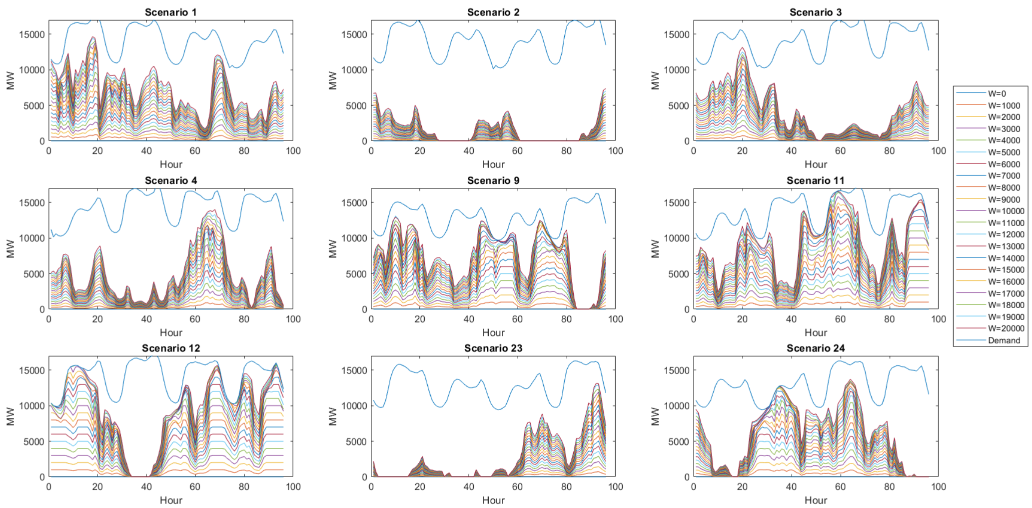

Note: Evolution of wind generation (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

It is clear from the above figure, that scenarios 2, 3, 4 and 23 are far less "wind friendly" than for instance scenarios 9 or 11. This result is consistent with Figure 8, where we were already in the position to predict which scenarios had high wind capacity factors. More generally, we observe from the above figure that, for each scenario, the generation of wind energy tends to converge point-wise towards the demand profile when subject to rising installed wind capacity levels, except when there is no wind. Indeed, even under an assumption of infinite installed wind capacity, the network cannot satisfy demand when there is no wind and will always continue to rely on conventional energy generation. This is a direct illustrations of the limits of renewable energies due to the fluctuations of their sources. The fact that higher installed wind capacity levels do not simply translate into a mere affinity transformation of a previous level is due to the dynamic interaction with conventional units. For instance, ramp-rate of minimum up-time constraints may impose the activity of some units over certain periods of time, therefore affecting the general shape of the wind generation profile.

One key indicator of interest is the penetration rate of wind energy. More specifically, the wind penetration rate measures the proportion of demand covered by wind and is formally expressed by:

$$\text{Penetration rate} = \frac{\text{Total Amount of Wind Energy Produced}}{\text{Total Demand}} \qquad \text{(in \%)}$$

This rate can be calculated either at the aggregate level (see upper plot in Figure 15) and for each scenario, or at each individual point in time (at the hourly step) and across scenarios. This latter approach ultimately provides access to the statistical distribution of wind integration for each installed wind capacity level (see bottom plot in Figure 15) and could be particularly informative to select a desirable level of wind capacity.

Figure 15: Wind Integration - Base Case

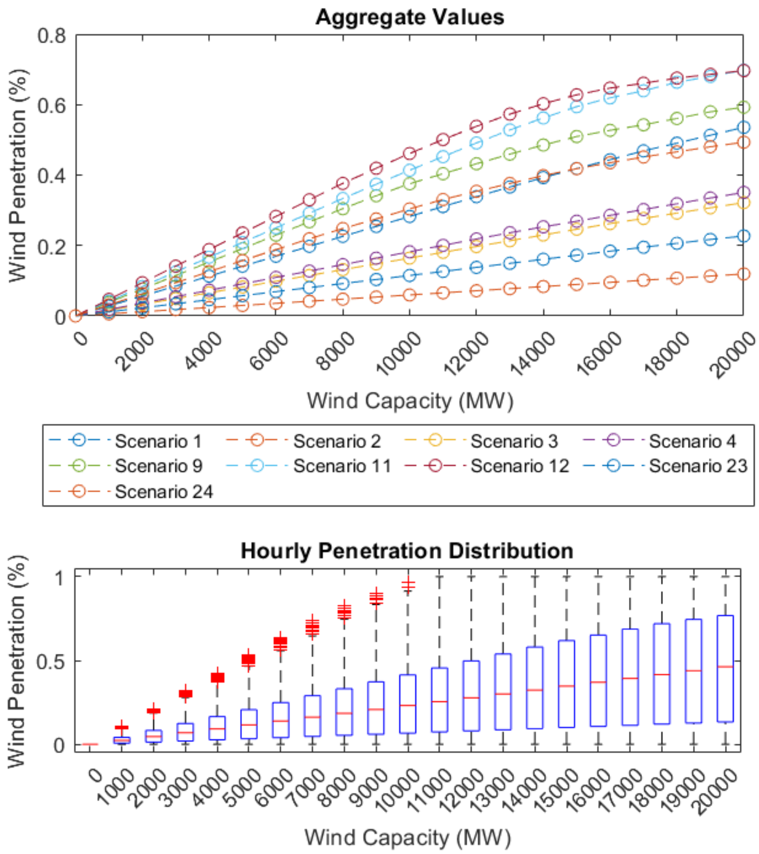

Note: The upper plot illustrates the dependency of the total wind penetration rate with installed wind capacity levels for each of the 9 representative scenarios separately. The bottom plot illustrates the distribution of hourly wind penetration rates across all nine scenarios for various installed wind capacity levels in the form of boxplots. The median (second quartile) of the distribution is represented by the red horizontal line within the blue box, which indicates the first and the third quartiles. Black whiskers depict extreme points not yet considered as being outliers while outliers are individually shown as red crosses.

It is clear from the first plot that the penetration rate of wind power is highly dependent on the chosen scenario of demand and wind. In fact, highest penetration rate is achieved for scenario 12. We further note that, as we increase the installed capacity level, the slope of some of the lines tends to flatten. This implies that, at very high installed capacity levels, the addition of more wind capacity would only marginally enable to integrate more wind into the network, leading to an inefficiency. This is an interesting result, since it suggests that there exists an optimal level of wind capacity. Such level would need to be determined by the comparison between the expected profits from an additional unit of wind capacity and the fixed costs involved for such a decision.

From the second plot, we conclude that, as wind capacity is increased, it becomes - occasionally - possible to supply the whole demand with wind for a specific hour. Yet, most of the distribution mass remains at levels strictly lower than 100% and each distribution in our dataset includes at least one point in time[12] when none of the demand is supplied by wind (0% points). As a consequence the median of the distribution does not increase as rapidly as the first quartile. Since installing new wind capacity has a cost, this second plot again shows that there exists an optimal wind capacity for the overall market. In the following section, we are going to conduct additional analyses which allow us to determine such optimal wind capacity level for the ISO New England network.

5.1.1 Conventional Generation

As explained in the previous section, the level of wind generation generally increases with the installed wind capacity level. Since demand is given exogenously, other generation sources need to be displaced to accommodate for more wind power. This section analyses the effects of wind penetration rate on the conventional sources such as coal, nuclear, natural gas and fuel oil.

Coal is not as flexible as other fuels, and the operation cost is lower than for natural gas or oil. As a consequence, as illustrated by Figure 16, coal is easily displaced in wind-friendly scenarios such as 9, 11, 12 and 24. By contrast, scenarios 1, 3, 4 and 23 require very high levels of installed wind capacity levels in order to start pushing some coal out of the market. Finally, scenario 2 never manages to displace any coal, which will be of particular interest when we will introduce a carbon tax into the system.

Figure 16: Coal Generation as Percentage of Demand - Base Case

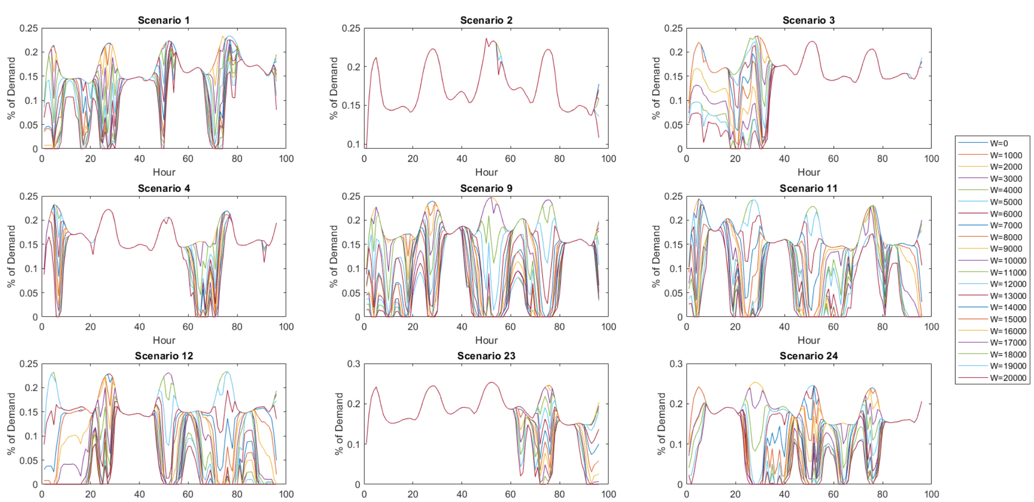

Note: Evolution of coal generation (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

[12] This is consistent with Figure 14, where it appears for instance that scenario 2 or 23 have a couple of hours with no wind generation at all.

Nuclear is the most inflexible sources due to its very high start-up costs. As a consequence, displacing nuclear requires both high wind levels *and* very large installed wind capacity levels (see Figure 17). Note that, in scenarios 9, 11, 12 and 24, it is possible to completely displace nuclear at certain points in time. This results is consistent with the very high penetration levels depicted in Figure 15.

Figure 17: Nuclear Generation as Percentage of Demand - Base Case

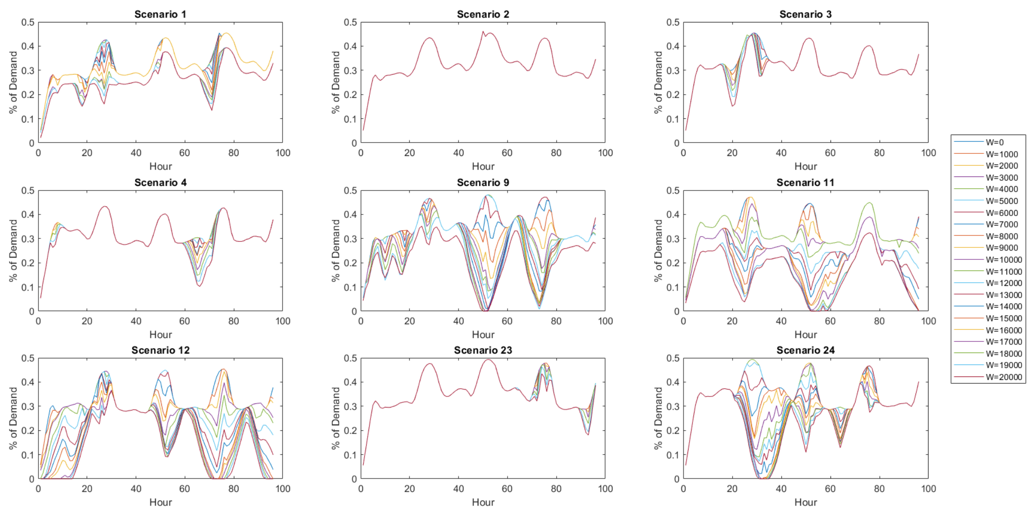

Note: Evolution of nuclear generation (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

Natural gas and fuel oil are the most expensive conventional sources in our database. Consequently, in the ISO New England, it can be expected that wind power is likely to have largest impacts on those two forms of conventional energy. Figures 18 and 19 confirm this prediction for each of those two sources. Most strikingly and contrary to the case of coal and nuclear, even scenario 2 now sees a drop in natural gas and oil.

Figure 18: Natural Gas Generation as Percentage of Demand - Base Case

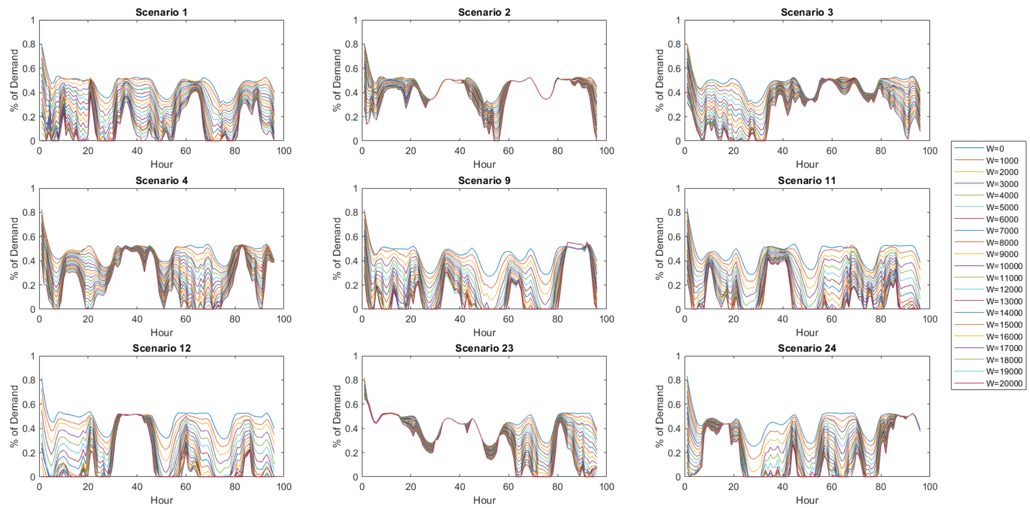

Note: Evolution of natural gas generation (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

Figure 19: Oil Generation as Percentage of Demand - Base Case

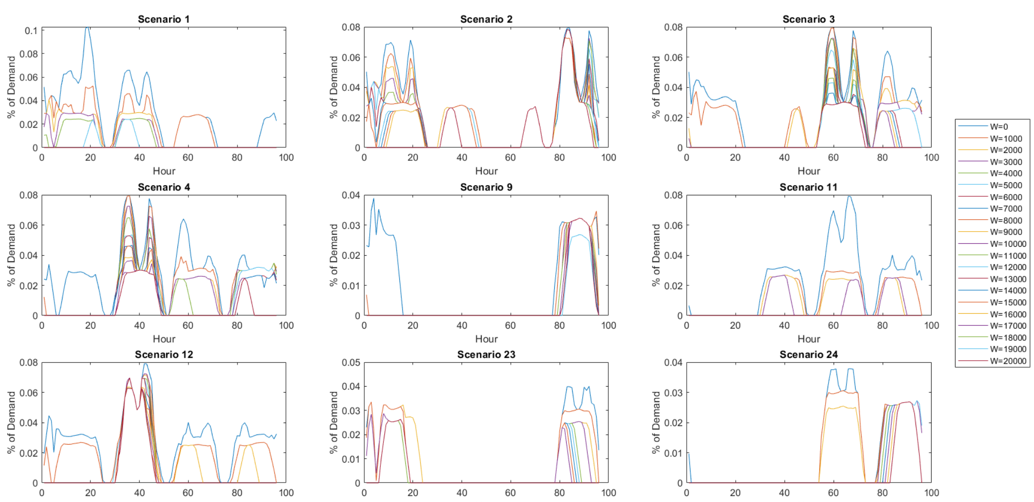

Note: Evolution of fuel oil generation (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

It is important to note that oil is only used as a peaker energy source and exclusively during specific hours of the day. Therefore, the percentage of the demand supplied by oil remains generally very low (less than 10% of total load in any of our chosen scenarios). The effects of wind power on oil integration therefore appear to be limited in absolute terms.

5.1.2 Energy Prices

In the previous section, we have shown that wind is highly adequate to displacing oil and natural gas. As explained in section 2, this will affect the average price as well as the variance of the price distribution since wind power depends on wind availability. This section analyses the effect of wind integration on the energy price.

Figure 20 presents the evolution of prices under our nine selected scenarios and for various installed wind capacity levels.

Figure 20: Evolution of Energy Prices - Base Case

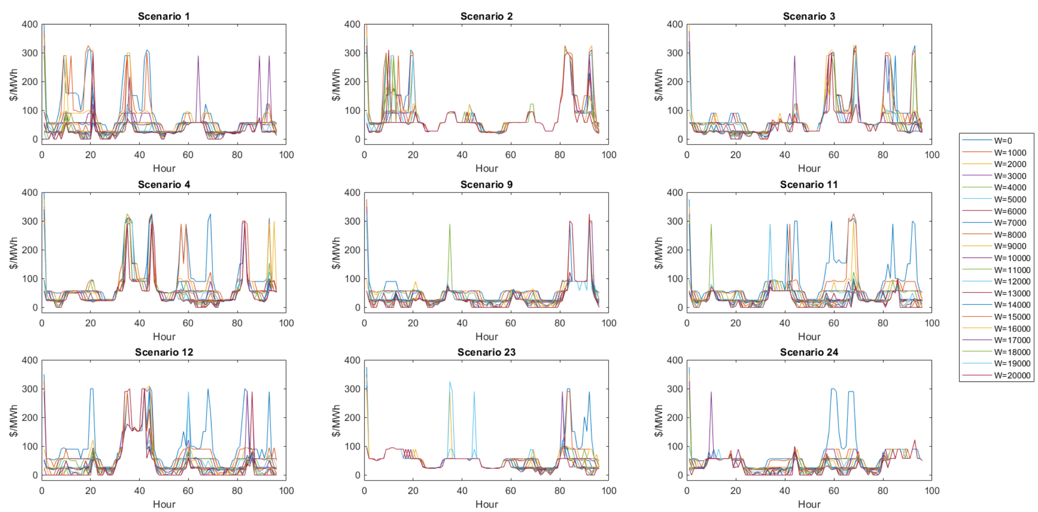

Note: Evolution of energy prices (y-axis) over the 4-day time period (x-axis) for each representative scenario (one plot per scenario) and under various installed wind capacity levels (see legend).

The reader may be wondering why all the price profiles shown in Figure 20 display an initial market price of about 400\$/MWh in the very first hour of the time window under investigation. The reason behind this very large price is that our model assumes all units are initially turned off. Therefore, the most flexible but also most expensive units are required to meet overall demand. Indeed, more cost-friendly units are subject to stringent ramp-rate constraints that seriously limit their ability to cover the initial demand.

According to the above plots, market price levels seem to jump between a limited number of possible levels. On top of being directly dependent on the evolution of demand, energy prices appear to be very sensitive to the installed wind capacity level, with certain spikes at the (approximately) 300\$/MWh level appearing only for given installed wind capacity levels. Depending on the chosen scenario, such spikes appear more or less frequently. For instance, only two price spikes appear in the worst-case situation with no installed wind capacity (i.e. $W = 0$) for the wind-friendly scenario 24, while scenario 2 sees - with certainty

- at least one price spike irrespective of the chosen level of installed wind capacity.

The effect explained on the previous paragraph has been anticipated in section 4.5 and can be expected to be less pronounced if we were to run our code using a quadratic objective cost function. Moreover, this is expected since the introduction of wind generation has zero marginal cost, and hence, when wind is not available we need to turn on very expensive and flexible units.

With the exception of such sudden spikes, the evolution of prices is roughly consistent from one wind capacity level to the other. More precisely, we can observe that, as higher levels of wind capacity are reached, prices seem to decrease on average. This is intuitive as wind power is assumed to have a zero marginal cost. In order to substantiate this observation more clearly, we propose to inspect different boxplots describing the hourly distribution of prices for the different wind capacity levels.

Figure 21: Distribution of Energy Prices as a Function of Installed Wind Capacity Levels - Base Case

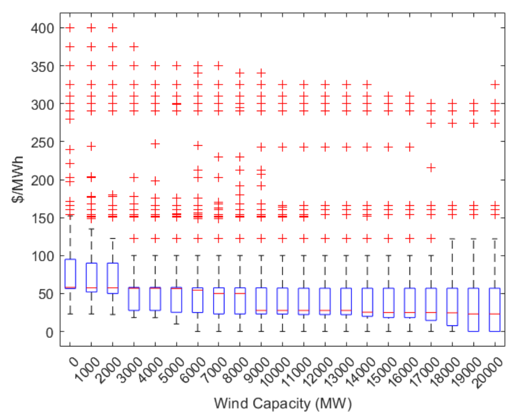

Note: The median (second quartile) of the price distribution is represented by the red horizontal line within the blue box, which indicates the first and the third quartiles. Black whiskers depict extreme points not yet considered as being outliers while outliers are individually shown as red crosses.

Figure 21 shows that, as we move to higher levels of wind capacity, the median shifts to lower values. A very interesting result from this plot is that outlier values (in red) tend to become smaller with increased installed wind capacity. By contrast, we note that the interquartile range tends to increase with higher wind capacity levels. Whether the overall variance of prices decreases or increases with installed wind capacity needs to be separately calculated as it cannot be clearly visualized on the boxplots. Figure 22 presents the evolution of the standard deviation of energy prices for different levels of installed wind capacity.

Figure 22: Energy Price Standard Deviation - Base Case

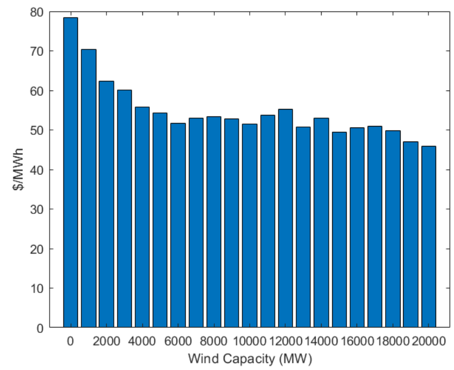

In the figure we can see that the decrease in the standard deviation of prices is faster for low levels of wind capacity. As we increase the capacity, the slope of the decrease of the standard deviation reduces, meaning that the gain in terms of the reduction in the price variance becomes smaller. This result can respond to two factors. First, if we include too much wind into the system, the distribution of the residual demand will be changed in such a manner that the decrease in the price standard deviation might not only stop, but also, it might eventually increase. Second, this fact can respond to the actual plants operating into the market. If wind integration is very high, and demand cannot be served by more wind, then the flexible units have to be used. If we already have online flexible units as gas (which are not as expensive as oil), the price variance can be expected to decline. However, if wind also displaces gas, the most expensive units may have to be turned on since those will be the most flexible ones.

Figure 23 takes a closer look at the changes in the distribution of residual demand under increasing installed wind capacity levels and for each of our nine selected scenarios. For this purpose, we fitted a normal distribution. The residual demand is defined as the difference of the demand and the wind generation. This means that the residual demand has to be supplied with conventional electricity sources. We can see in these series of nine plots that, as wind integration increases (lighter-gray colors) the distribution mean value of the residual demand decreases, while the variance increases.

Figure 23: Distribution of Residual Demand - Base Case

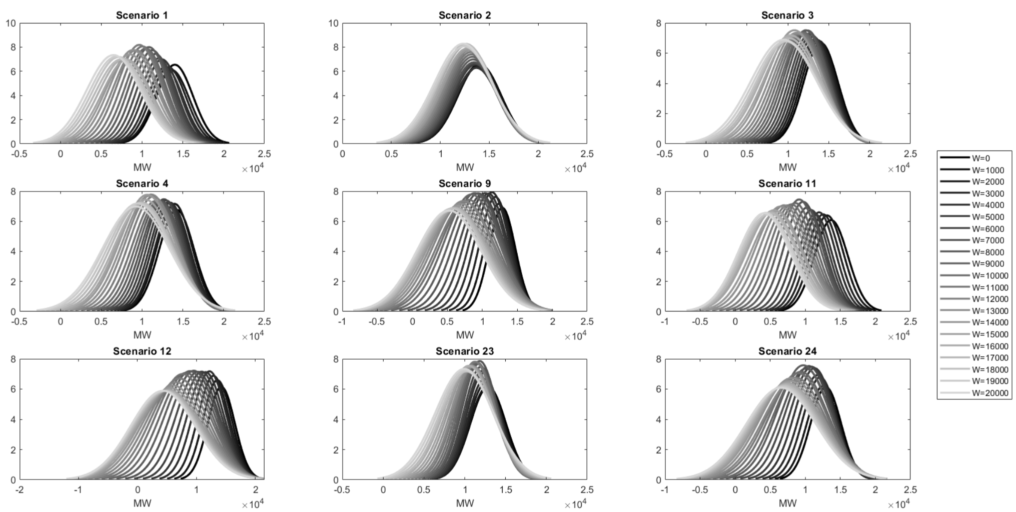

Note: Evolution of the residual demand distribution, fitted as a normal, with increased installed wind capacity levels.

The standard deviation analysis presents an interesting result. First, we can see that the variance of the residual demand increases as the system has a larger share of wind capacity. Yet, the standard deviation of the energy price decreases as the wind installed capacity increases. We can explain this situation with the combination of two factors: (1) the displacement of very expensive units by wind, (2) there are still flexible units (natural gas) on which can be ramped-up without facing large costs when wind generation is limited or not available. This observation is partly driven by the fact that we assume $G^{min} = 0$, as this allows for a particularly flexible response.

5.1.3 Profit and Market Industry

Analyzing the effect of wind integration on the aggregate profit and the average profit is important to understand the producer surplus in this market as well as the overall market dynamics. The description we have provided so far, has a direct connection with firms' profit. In fact, if prices are smaller, revenue levels necessarily decrease for conventional producers. In addition, if some specific fuels are displaced, the average profit of those specific firms will further decrease.

Figure 24 presents the aggregate profit in the ISO New England market after the introduction of wind power in the network.

Figure 24: Dependence of Profits on Installed Wind Capacity Levels - Base Case

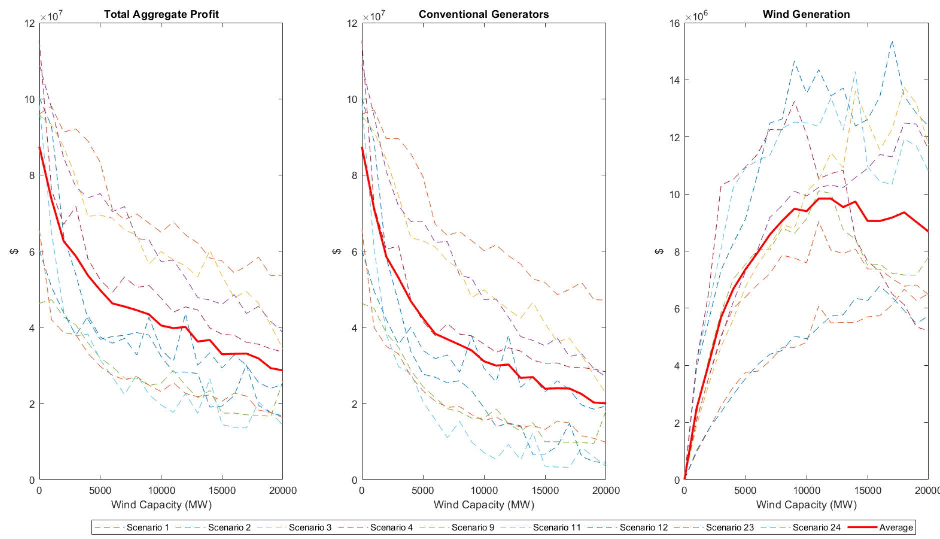

Note: Producer profit as a function of installed wind capacity levels for each of the nine representative scenarios as well as the average thereof (in red) for all producers in aggregate (left chart), for conventional producers only (middle chart) and for producers of wind power (right chart).

As we can see in the above figure, the overall producer surplus decreases as more and more wind capacity is installed. The reason for this is that the increase in the profit of wind generators is too small to compensate the drop in the profit of conventional producers. Note that, as wind penetration rises, it reaches a point (around 12,000 MW of installed wind capacity) at which the average profit achieves a maximum. A priori, this suggests that there is a capacity level at which wind generators would be maximizing their profit. Since our chosen scenarios cover the entire scope of possible demand x wind scenarios but are not weighted by any frequency of occurrence, we may not yet conclude that the level identified on Figure 24 corresponds to the actual optimum of the problem.

Note that the dynamics seen in the previous graph respond mainly to the changes in prices described previously. Moreover, conventional generators are producing less and this decreases the conventional profit. Now, these dynamics may be fundamentally different depending on the underlying fuel type. Figure 25 thus investigates the profit by fuel type as a function of installed wind capacity.

Figure 25: Decomposition of the Scenario-Averaged Producer Profit by Conventional Fuel Type as a function of installed wind capacity - Base Case

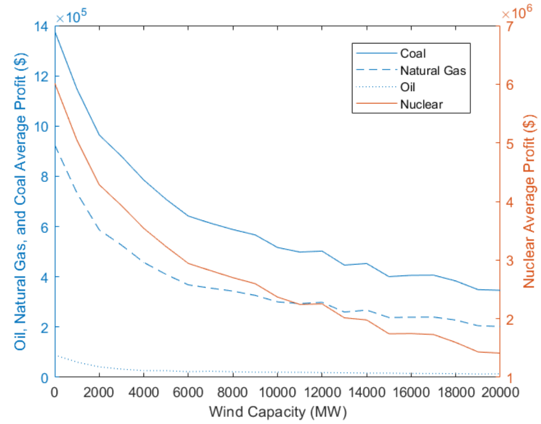

The entry of wind clearly affects the profit of conventional plants. Figure 25 shows that the introduction of the very first wind units have the highest (negative) impact on the profit of conventional plants when looking at averages across our nine selected scenarios. Moreover, note that oil reaches profit levels which are close to zero as early as wind capacity reaches 4,000 MW. This may appear to be problematic since those flexible plants are likely to be pushed out of the market in the long-run, leading to potential reliability issues as soon as intermittency of wind stikes.

From this analysis we can conclude that conventional plants are, on average over our nine selected scenarios, worse off given that the energy price tends to decrease due to the introduction of wind power, and the generation level is also smaller. The next section presents an analysis of producer revenues which, at the same time, represents the cost to consumers.

5.1.4 Consumers' Cost

Consumers pay for the electricity as well as the reserve. This amount is equivalent to the revenues collected by generators. This section presents an analysis of these two variables.

The total consumer cost is given by the energy and reserve payment. This means that

$$\text{Consumers' Cost} = \sum_{t=1}^{T} P_t^e w_t + \sum_{t=1}^{T}\sum_{i=1}^{N} P_t^e g_{i,t} + \sum_{t=1}^{T} P_t^r \sum_{i=1}^{N} r_{i,t}$$

where P_t^e is the energy price in period t, P_t^r is the reserve price in period t, and $g_{i,t}$ and $r_{i,t}$ are the generation and reserve in unit i at period t. Finally, w_t is the wind generation in period t.

Figure 26 presents the consumers' cost as well as its decomposition into conventional energy, wind power, and reserve costs as a function of installed wind capacity.

Figure 26: Consumer Costs (in Total as well as Decomposed into Conventional Energy Cost, Wind Power and Reserve Costs) as a Function of Installed Wind Capacity - Base Case

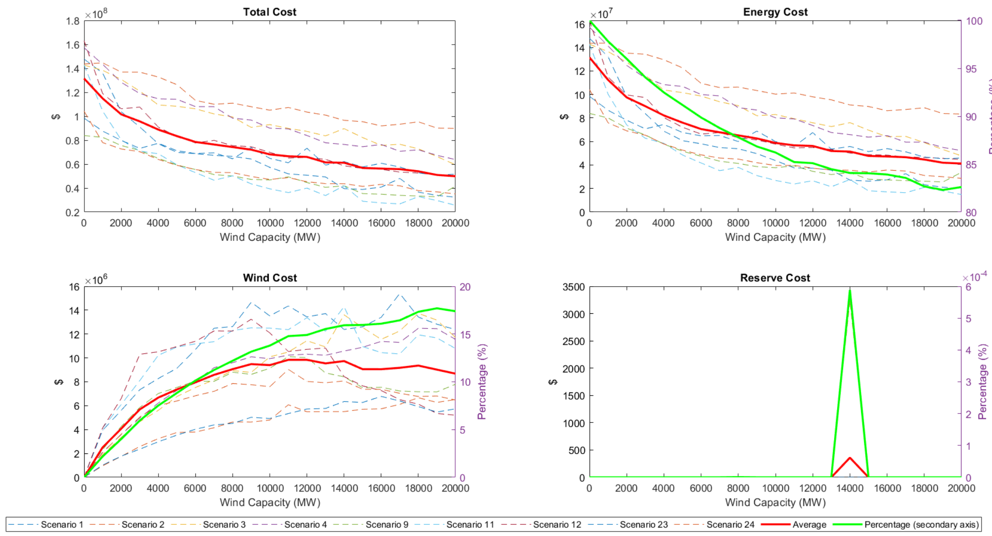

Note: The red line is the average cost across the nine representative scenarios. The green line depicts the average proportion of each of the components in total costs. This value is on the secondary axis.

As we can see on the above chart, consumers benefit from a lower overall cost thanks to the introduction of wind generation. The more wind power is generated, the lower the consumer cost tend to be. This holds true in general for each of our nine representative scenarios and in particular for the average across scenarios. While the shape of the reserve cost component appears to be surprising, it is important to note that, in magnitude the effect of reserve cost on the overall cost is *de minimis*. Indeed, the maximal proportion of total costs due to reserve costs remains negligible and below $6{\cdot}10^{-4}$ %. On average across our nine scenarios, wind costs account up to 12% of total costs and seem to achieve a maximum at around 11,000 MW. By contrast, the average trend over both total and conventional energy costs displays a strictly decreasing trend without converging to any particular level at the range of wind levels considered here (namely up to 20,000 MW).

The ISO New England webpage establishes that 14.8 million people live in the area of the market.[13] Using this information in conjunction with the results presented in Figure 26, we can retrieve the actual per capita consumer cost (see Figure 27).

[13]https://www.iso-ne.com/about/key-stats/

Figure 27: Per Capita Consumer Cost - Base Case

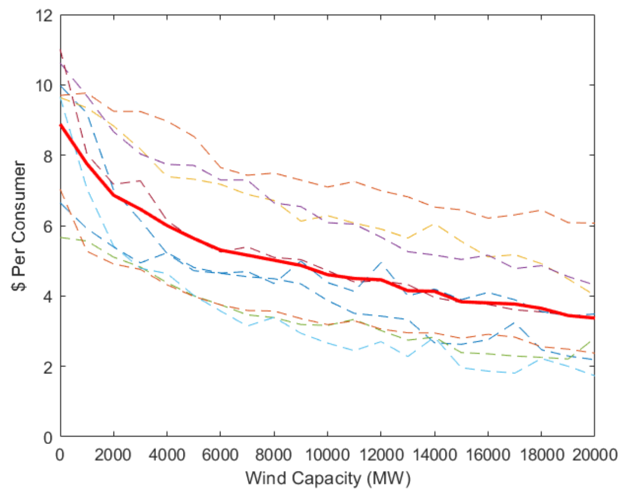

Note: The red line is the average cost per consumer across the nine representative scenarios. The cost per capita was calculated using the number of people living in the IS New England area (i.e. 14.8 million people).

As we can see on the above chart and when considering the average cost value across the nine scenarios, the per capita consumer cost substantially decreases from approximately $9.0 to $3.5 when wind capacity reaches 20,000 MW. This represents a percentage decrease of 61%. Yet, note that the slope of this decrease becomes flatter as we move to higher levels of wind capacity, following the overall shape of total consumer costs depicted in Figure 26. Moreover, it is important to note that the differences between individual scenarios is substantial. The difference between the best- and worse case envelope of our nine scenarios appears to remain more or less independent on the installed capacity level and amounts to approximately $5.

5.1.5 Operation Cost

In an earlier section, we have shown that the introduction of wind is displacing mainly oil and natural gas. Yet, coal can also be displaced when the scenario is wind-friendly and when we have high levels of installed wind capacity. Since we assumed that wind has zero marginal cost, this situation will affect the operation costs of generators.

Figure 28 shows the operating costs for our nine representative scenarios as a function of installed wind capacity. Irrespective of any given scenario, operating costs appear to strictly decrease with higher wind capacity levels. This observation is quite intuitive as the most expensive units are progressively being pushed out of the market as soon as more wind power is integrated into the network.

Figure 28: Operating Costs - Base Case

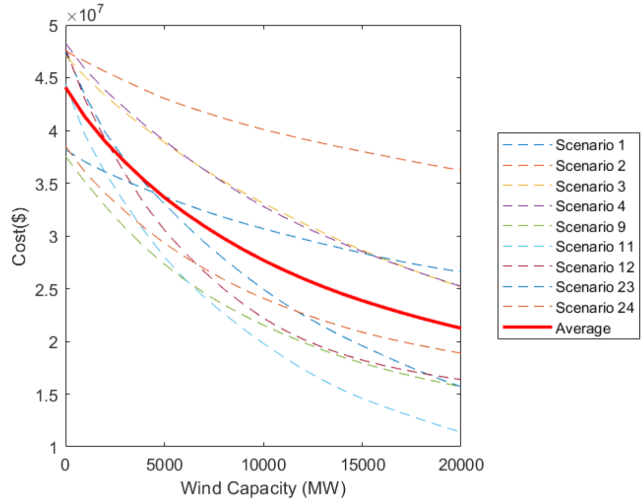

Note: The red line is the average cost per consumer across the nine representative scenarios.

In order to appreciate the benefit of each individual MW of additional installed wind capacity unit, we can propose to compute the marginal decrease of installed wind capacity. This marginal gain represents the reduction in the operating costs as we increase the installed capacity of wind[14] (see Figure 30).

Figure 29: Marginal Benefit - Base Case

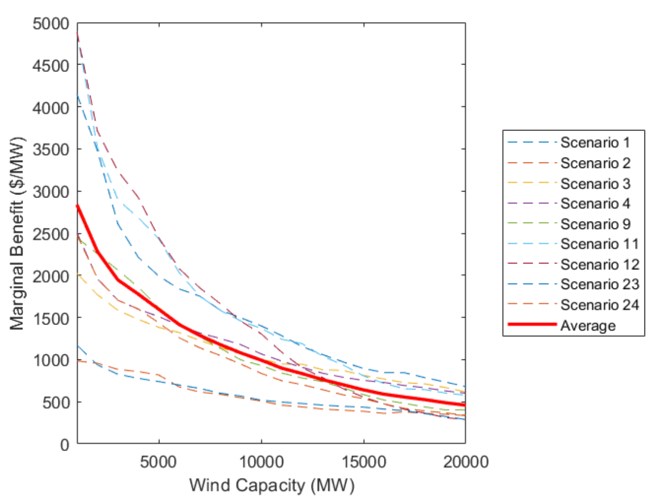

Note: The marginal benefit is the reduction of operating cost per installed MW of wind. The red line is the average marginal benefit across the nine representative scenarios.

As expected, the first units appear to have a large marginal benefit. However, as we increase the level of wind capacity, such benefits in terms of operating cost reductions per

[14] This marginal curve was approximated linearly given the different observations we have.

unit decrease.

From a cost-benefit perspective, investment in new MW of wind has to be done as long as the benefits are greater than the marginal costs of that unit, which is given by the investment cost, land cost, among others. The U.S. Energy Information Agency (EIA) estimates[15] that the construction cost of wind power facilities from 25 to more than 200 MW ranges from 1,676 to 1,268 dollars per killowat in 2018, depending on the size of the facility. These values are presented in Table 3. The same agency[16] establishes that wind turbines are designed with lifespans of between 20 and 25 years. Assuming that the capacity factor is constant throughout the whole life of the turbine, the annualized investment cost is given by:

Table 3: Construction Cost of Wind Farms (2018) per MW

	$(unit)$	25-100 MW	100-200 MW	+200 MW
Investment	$(\$/KW)$	1,676	1,435	1,268
Lifespan	(Years)	[20,25]	[20,25]	[20,25]
Annualized Investment	$(k\$/MW)$	[67.0,83.8]	[57.1,71.8]	[50.7,63.4]
Average Annual Investment	$(k\$/MW)$	75,420	64,425	57,060

Note that this analysis does not consider land-cost and the impact of wind on birds and other animals. Hence, the costs values presented in the previous table are a lower bound to the true total costs.

Following Figure 30, the following plot presents the range of annual expected marginal-benefit as well as the investment costs. This plot assumes that the marginal benefit obtained for a 4-day period can be extrapolated to the whole year.

[15]EIA (2020). Average U.S. construction costs for solar and wind generation continue to fall.
[16]EIA (2017). Repowering wind turbines adds generating capacity at existing sites.

Figure 30: Cost-Benefit Analysis - Base Case

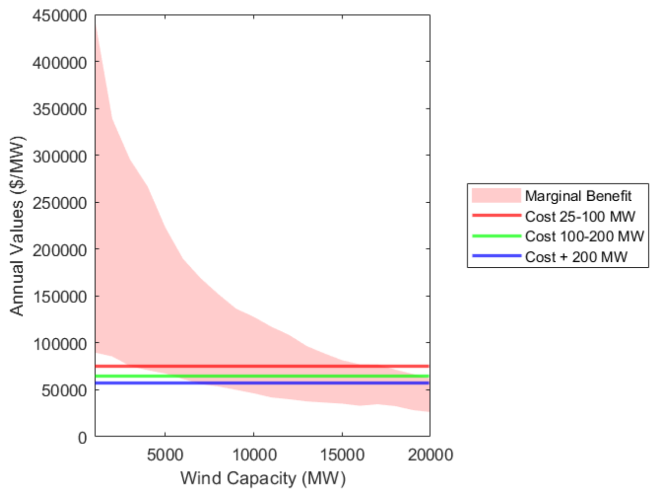

Note: The marginal benefit is the reduction of operating cost per installed MW of wind. The shaded area represents the range of marginal benefits for our representative scenarios.

Figure 30 shows a range for optimal investment levels depending on the size of the facilities. This figure shows that for a capacity level below 3,000 MW, any type of facility will have a greater annual marginal benefit than the annualized per unit investment cost in any of the 9 scenarios. Moreover, if we aim to build small size facilities, the lower bound for the optimal level investment level is reached at 3,000 MW. For the other sizes this lower bound is reached at 5,500 MW for 100-200 MW facilities and 7,500 for the larger facilities. Note that when we reach the lower bounds mentioned before, we would need to assign probabilities to each of the different scenarios in order to define the optimal wind capacity level.

6 Wind Integration and Carbon Taxes

Carbon emissions generate a welfare loss due to their role in climate change. Let $E_{i,t}(g_{i,t})$ be the emissions of generating quantity $g_{i,t}$ by generator i in period t. Let $C_{e,t}(E_t)$ be the emissions damage function where $E_t = \sum_{i=1}^{N} E_{i,t}$. In this context, the social welfare function for a given period t is given by:

$$B_t(D_t) - \left(C_{e,t}(E_t) + \sum_{i=1}^{N} (b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t}) \right)$$

Recall that the demand is inelastic and endogenously given. Then, maximizing the social welfare function for all the period is equivalent to minimize the social welfare cost function given by

$$\sum_{t=1}^{T} \left(C_{e,t}(E_t) + \sum_{i=1}^{N} (b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t}) \right)$$

Assume the damage function is linear in emissions and the marginal damage is given by $c_{e,t}$ so that $C_{e,t}(E_t) = c_{e,t} \sum_{i=1}^{N} E_{i,t}(g_{i,t})$ is the externality generated by generator i. In order to internalize this effect, we can establish a carbon tax such that its rate $\tau_t = c_{e,t}$. Thus, our social welfare cost function becomes:

$$\sum_{t=1}^{T} \sum_{i=1}^{N} \left(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t} + \tau_t E_{i,t}(g_{i,t}) \right)$$

The marginal damage can be represented by the social cost of carbon (SSC_t). So, under the introduction of a carbon tax, the revised objective function to minimize becomes:

$$\sum_{t=1}^{T} \sum_{i=1}^{N} \left(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t} + SSC_t E_{i,t}(g_{i,t}) \right)$$

The carbon tax is thus expected to change the optimal commitment solution as more polluting plants will generate larger emissions. In this section, we aim at determining the effects of a carbon tax on top of the introduction of wind power.

6.1 Emission factors

For the assessment of the effect of a tax on CO_2 emissions, we use additional information about the quantity of CO_2 that is emitted per kWh of electricity. Such emission factors primarily depend on the fuel type and the facility's efficiency. As we do not have access to more precise information about the generators in our database, we will assume, for the purpose of the present assessment, that emission factors only vary by fuel type. While nuclear and wind power do not emit any CO_2 and will be considered to have a zero emission factor, we will use average 2018 emission data for US electricity generation as provided by the US Energy Information Administration (*Frequently Asked Questions (FAQs) - U.S. Energy Information Administration (EIA)*, n.d.) (see Table 4).

Table 4: USA Electricity generation and CO_2 emissions 2018

	electricity production (GWh)	emissions ($MtCO_2$)	Emission factor (tCO_2/MWh)	Tax (\$/MWh)
coal	1 124 638	1 127	1.002	50.11
natural gas	1 246 847	523	0.419	20.97
fuel oil	21 860	21	0.961	48.03

Notes: This data includes electricity generation and CO_2 emissions according to the EIA, for the entire United States in 2018, as well as the resulting emission factors and tax rates per MWh (under the assumption of a 50\$/tCO_2 carbon tax level) for each of coal, natural gas and fuel oil.

Following Table 4 and our simplifying assumptions, the emissions levels can thus be defined as:

$$E_{i,t}(g_{i,t}) = \eta_j g_{i,t}$$

where η_j is the emission factor of fuel j. Then, the objective function to minimize becomes:

$$\sum_{t=1}^{T}\sum_{i=1}^{N}\left(b_i g_{i,t}^2 + a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t} + \tau\eta_j g_{i,t}\right)$$

As explained in the previous subsection, we propose to use a carbon tax rate set at the social cost of carbon, which reflects the economic harm due to the emission of one additional ton of CO_2 today. According to EDF, [t]*he current central estimate of the social cost of carbon is over $50 per ton* [of CO_2] *in today's dollars. While this is the most robust and credible figure available, it does not yet include all of the widely recognized and accepted scientific and economic impacts of climate change. For that reason, many experts agree this is far lower than the true costs of carbon pollution.*[17] While the precise value of the social cost of carbon is currently subject to a wide range of estimates, the value proposed by EDF appears to be a fair ballpark figure.

Irrespective of the precise value for the social cost of carbon, it is clear that coal and fuel oil for combustion is more than twice as polluting than the combustion of natural gas. The adoption of a carbon tax should therefore mechanically advantage natural gas units, as coal and fuel oil units will be penalized. By contrast, nuclear power is not subject at all to the carbon tax and is unlikely to benefit from the adoption of a carbon tax given its inherent cost advantage over other conventional fuels. Such effects can be expected to increase with the chosen carbon tax level.

6.2 Results: Tax and Wind Integration

This section presents the results for the case when we include a carbon tax as described in the previous section. It is important to note that the amount of wind generated is independent of the tax and only depends on the wind capacity factor and the chosen installed capacity. Moreover, since demand is given and assumed to be inelastic, the tax will not affect the demand levels either. As a consequence, the integration levels presented in the base computation also apply to this alternative case.

By contrast, the introduction of a carbon tax will have a direct impact on the fuel mix between conventional power generation. As we have already shown that the level of installed wind capacity also has a direct impact on the fuel mix (as it enables to displace some of the expensive units), the sensitivity to the tax can be expected to be higher for significant installed wind capacity levels. Figure 31 presents the average generation values across our nine chosen scenarios for the different fuels and wind capacity levels. In this plot, we only present coal, gas and oil since these are the only fuels affected by the tax.

[17] `https://www.edf.org/true-cost-carbon-pollution`

Figure 31: Evolution of Coal, Natural Gas and Fuel Oil Power Generation - With and Without Tax

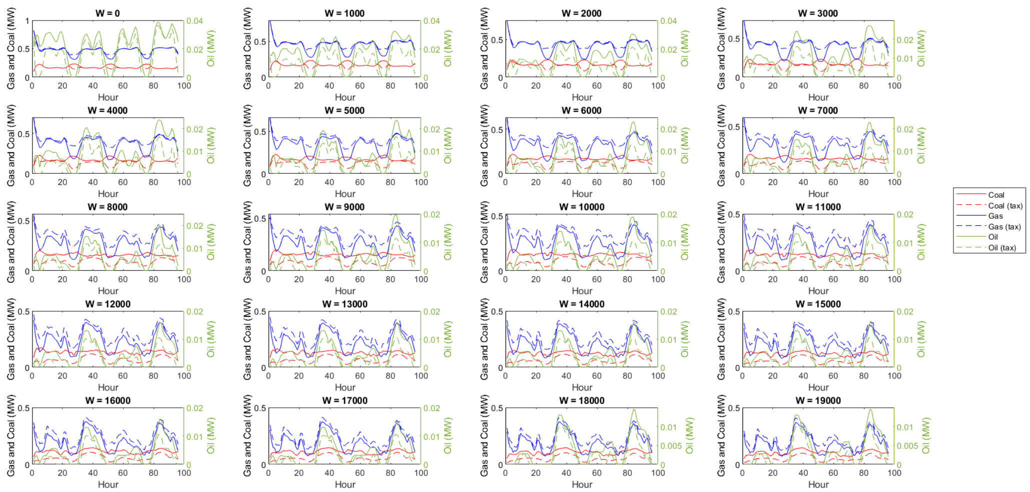

Note: Evolution of coal (red), natural gas (blue) and fuel oil (green) power generation over the 4-day timeframe with (dashed) and without (plain) a carbon tax of 50$/t$CO_2$ and for 20 different installed wind capacity levels. The output corresponds to the average across our nine selected scenarios.

From the plot, it appears that the introduction of a carbon tax favors natural gas generators over coal and fuel oil generators. Indeed, the results show that the output of natural gas generators increases under the tax scenario, while the output of each of coal and oil decreases. Moreover, the displacement of coal is larger than oil since the emissions factor of coal is larger. In the following section we discuss the implications of these changes in the equilibrium generation on the other variables of interest.

6.2.1 Effect of Taxes on Prices and Profits

The situation presented in the previous paragraphs will have a direct effect on prices. Figure 32 presents the distribution of prices with and without a carbon tax.

Figure 32: Price Distribution - With and without Tax

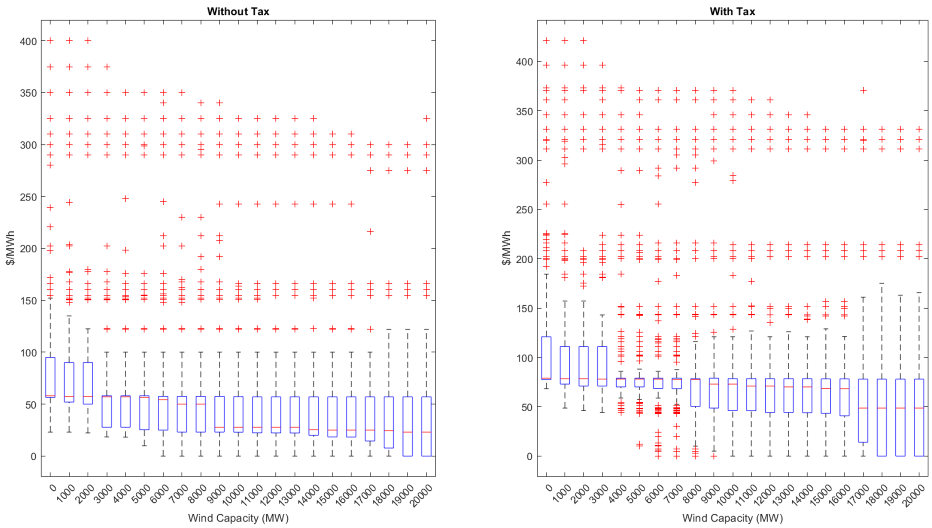

Note: Boxplots summarizing the distribution of the market price without (left plot) and with (right plot) a carbon tax of 50$/t$CO_2$ for various levels of installed wind capacity levels across our nine selected scenarios. The median (second quartile) of the distribution is represented by the red horizontal line within the blue box, which indicates the first and the third quartiles. Black whiskers depict extreme points not yet considered as being outliers while outliers are individually shown as red crosses.

From the previous figure, we conclude that, for all the wind capacity levels, the price is higher under the tax regime. This is a relevant result we must keep in mind since the tax will have distributional effects since electricity is an essential good, and poor peoples will be affected by the higher prices. While median energy prices will reach higher levels under the tax, the distribution of prices under the tax regime still reaches similarly low values. In this sense, the next plot presents the comparison of the energy prices' standard deviation:

Figure 33: Wind Generation - With Tax

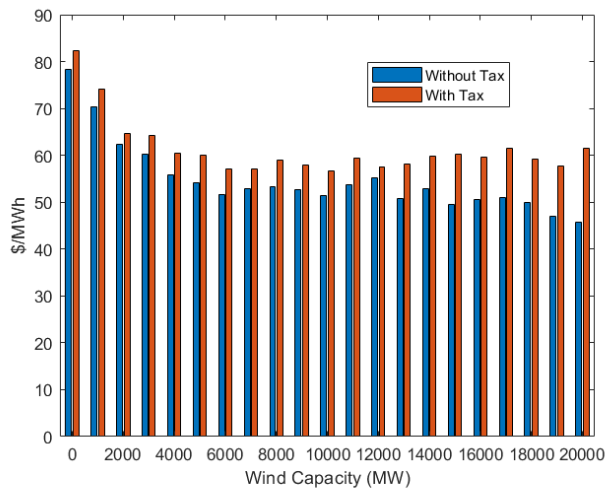

As we can see, the adoption of a carbon tax increases the standard deviation of energy prices. As wind generation is independent of the inclusion of the tax and the residual demand remains unchanged, this observation is primarily due to the effect of the carbon tax on peaker generators relying on fuel oil. As these generators are still highly necessary to adjust generation, their usage translates in a greater fluctuation of the market price.

We already know from the case without a carbon tax, that the displacement of conventional sources will decrease the average profit of those sources. The following figure presents the impact of installed wind capacity on the average profit across our nine scenarios, decomposed by conventional fuel type.

Figure 34: Profit of Coal, Natural Gas and Fuel Oil Power Generators as a function of installed wind capacity - With and Without Tax

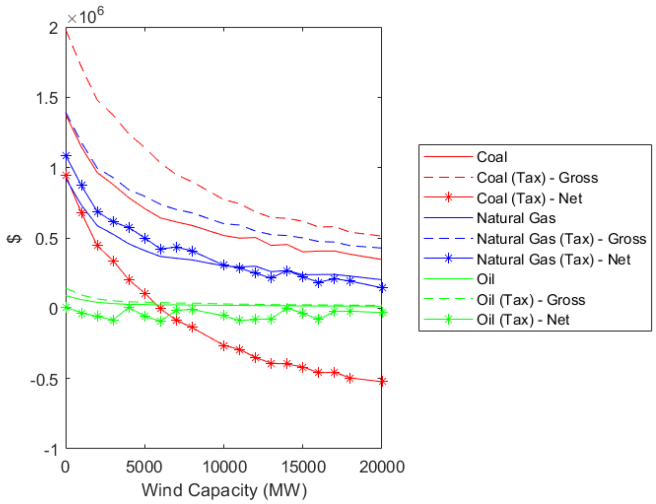

Note: The gross profit is the profit before the payment of the tax. The net profit is the profit after the payment of the tax.

Figure 34 illustrates the effect of a carbon tax on coal, natural gas and fuel oil power generators as a function of installed wind capacity. As we can see, in all the cases, the average profits across our nine scenarios decrease as the wind capacity increases. The reason is that the effect of the wind capacity combines with the effect of the tax. Therefore, we should create methods to attract long-term investment as some units might leave the market.

Moreover, the figure shows that the gross profit is higher since the price increases due to the effect of the tax. However, the plants have to pay for the tax, and hence, the net profit is below the base average profit for both coal and oil. Indeed, the tax has a large effect on the coal plants on average. Different is the situation of natural gas in which the net profit is similar to the profit in the base case. This situation will provide incentives to coal plants to invest in abatement technologies or leave the market. Moreover, this might be showing that the tax we are using for this section is to aggressive since we are forcing coal plants to have losses as wind capacity increases.

6.2.2 Industry Profit and Consumer Costs

The combination of an increase in the price will favor the overall profit of the industry. The following figure presents the average total profit across our 9 scenarios, as well as the conventional and wind profit for the tax case as well as the situation without the tax.

Figure 35: Aggregate Industry Profit as a function of installed wind capacity - With and Without Tax

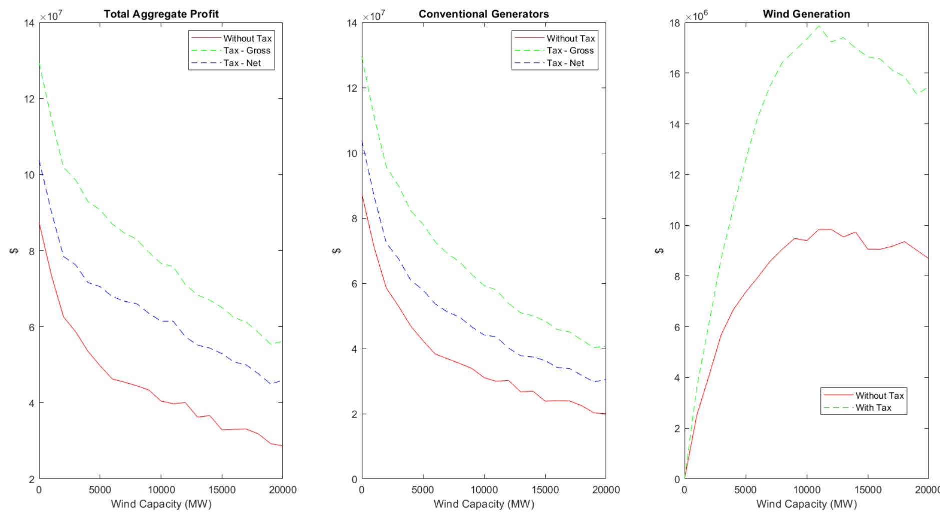

As we can see, the introduction of the tax translates into a larger average industry profit across our 9 scenarios relative to the base case. In order to reconcile Figures 34 and 35, we need to consider that figure 34 does not consider nuclear plants and wind generation. Moreover, Figure 35 does not show the individual effects by fuel, and it is an industry measure which shows that the increase of the price compensates the other effects produced by the tax. This represents a transfer of surplus from the consumers to the producers, and also, the government. Recall that the incidence of a tax depends on the elasticity of the demand and the supply. In fact, the consumers have to pay a larger price for the energy, and since the demand is inelastic, this means that consumers' have to face a larger cost. Below we present the average cost consumers face in this market, with and without the tax, across our 9 scenarios:

Figure 36: Consumers' Cost as a function of installed wind capacity - With and Without Tax

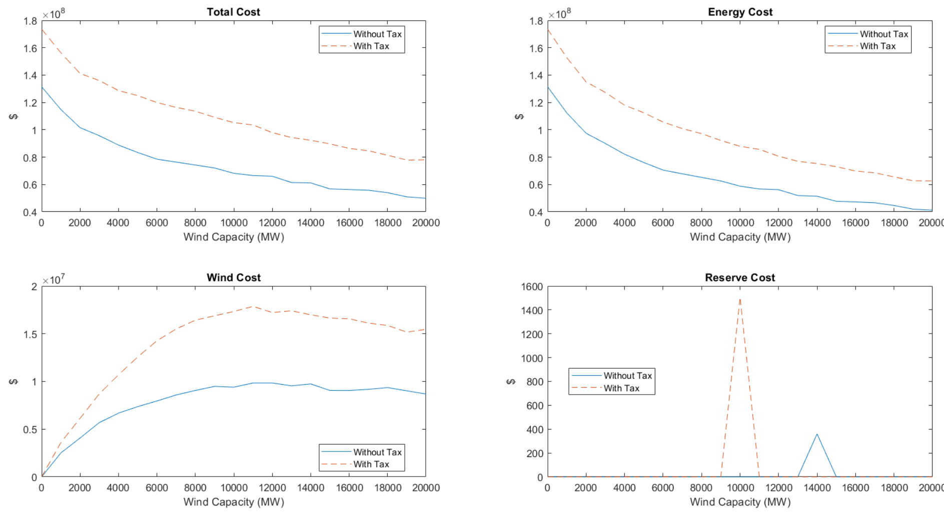

From the previous figure it is clear that there is a surplus transfer from the consumers to the producers and the government. Indeed, the larger energy price affects the consumers. Yet, we can still see that the highest the wind capacity, the cost paid by the consumers is still decreasing. In this case, the percentage decrease in consumers' cost is marginally smaller than the values presented when we did not have the tax (i.e a decrease of 50%). This shows that the introduction of the tax offsets marginally the decrease in the total cost faced by the consumers. Finally, using the same population level we presented before, we computed the per capita cost:

Figure 37: Consumers' per Capita Cost as a function of installed wind capacity - With and Without Tax

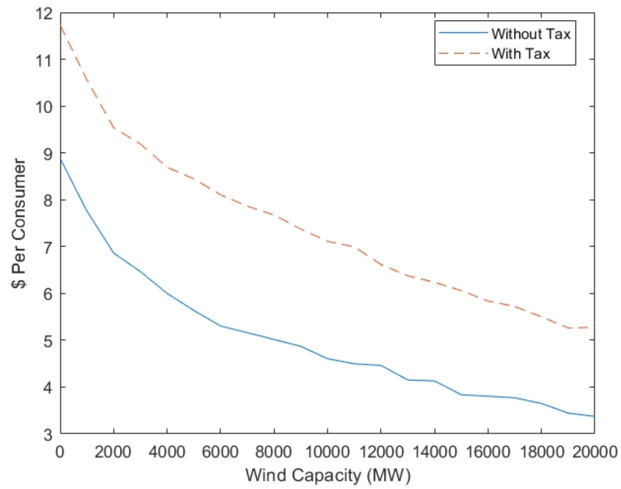

As we can see, the average cost is 3 dollars higher per consumer due to the tax. This situation will generate distributional effects on the economy. However, we could advice to the policymaker to use the tax revenue to make a direct transfer to consumers to offset the effect of the price increase.

6.2.3 Social Welfare: Having a carbon tax is good?

As explain in the theory section, the aim of the system operator is to maximize the social welfare. From the extension presented in this section, we know that the social welfare cost is defined by:

$$a_i g_{i,t} + NLC_i u_{i,t} + SUC_i v_{i,t} + SSC_t E_{i,t}(g_{i,t})$$

Recall that we are using a linear cost function given the computational advantages this function has over the quadratic one. In previous sections we have explored the effect of the tax on different variables of interest. This section explores the effect of the tax on the social welfare and compares the result with the case without the tax. In order to conduct this analysis, first note that the introduction of the carbon tax affects which fuels are producing. This will have a direct impact on total emission. Below we present the emissions under the tax regime and without the tax.

Figure 38: Total Industry Emissions as a function of installed wind capacity - With and Without Tax

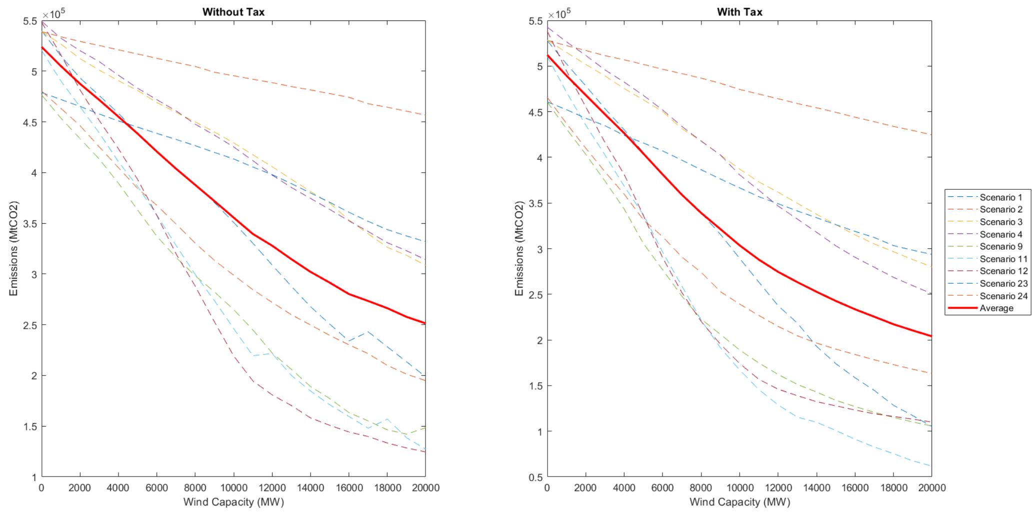

As we can see in the plots, total emissions decrease faster as we increase the wind capacity when we use a carbon tax than without the tax. Moreover, emissions are smaller with the tax than without it for all wind capacity levels. This is showing that the tax achieves its objective. Moreover, note that there is an inflection at 10,000 MW. This is showing that the decrease of emissions per MW of installed wind slows down after we reach 10,000 MW. Recall that multiplying these emissions levels by 50 (social cost of carbon) gives the economic costs of emissions.

The second component of the social welfare cost is the operating cost. The following plot presents the operating cost for both situation, with and without the tax. Note that the optimized objective function of the scenarios with tax include both the emissions damage function and the operating cost. Hence, we subtracted the damage function from the optimized objective function for those scenarios in order to identify the operating costs.

Figure 39: Operating Cost as a function of installed wind capacity - With and Without Tax

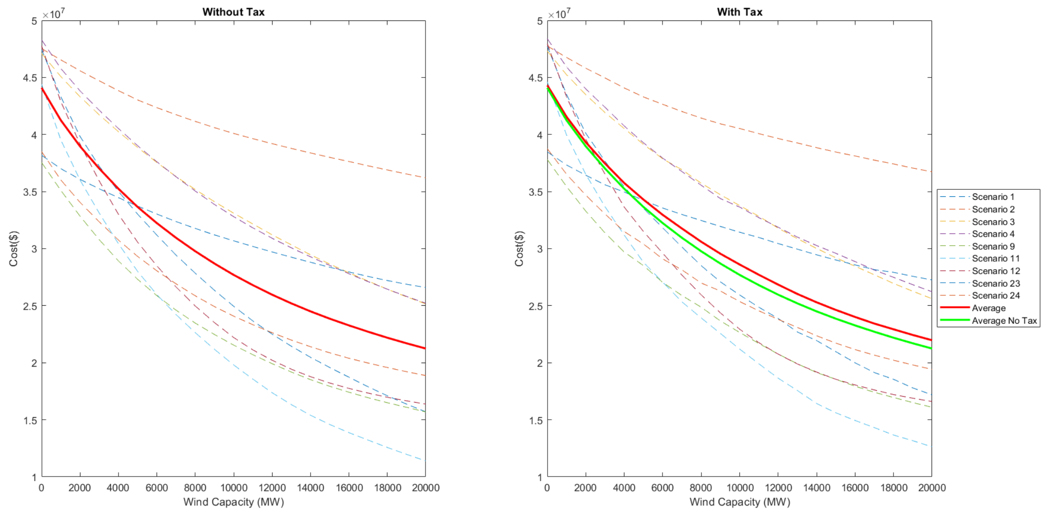

The figure shows that when we introduce the tax we have an increase in operating costs. The reason is that we are using less polluting but more costly generating units. The following figure presents the total social welfare and the marginal benefit (decrease of social cost plus pollution costs per unit of MW installed).

Figure 40: Aggregate Social Cost and Marginal Benefit as a function of installed wind capacity - With and Without Tax

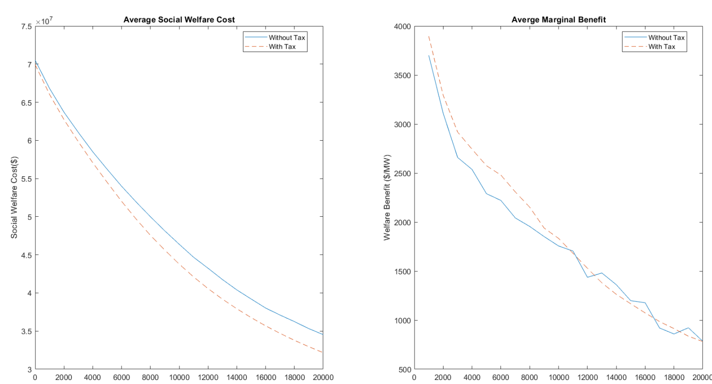

As we can see in the plot, the social cost is smaller with the tax than without the tax. Then, having a tax is optimal given that it internalizes the emissions costs. Moreover, we can see that the first units have a large marginal benefit, and this benefit is larger when we have the tax. However, the marginal benefit is decreasing in the level of wind capacity and we reach to similar marginal values with and without the tax. Similarly to what we did for the base case, we present the cost-benefit analysis for the tax scenarios. As we can see in this plot, for any value below 8,000 MW it is optimal to invest in any wind facility of any size in our 9 scenarios. Yet, this is not the optimal level, since this would require estimating expectations using the 90 scenarios.

Figure 41: Cost-Benefit Analysis - With Tax

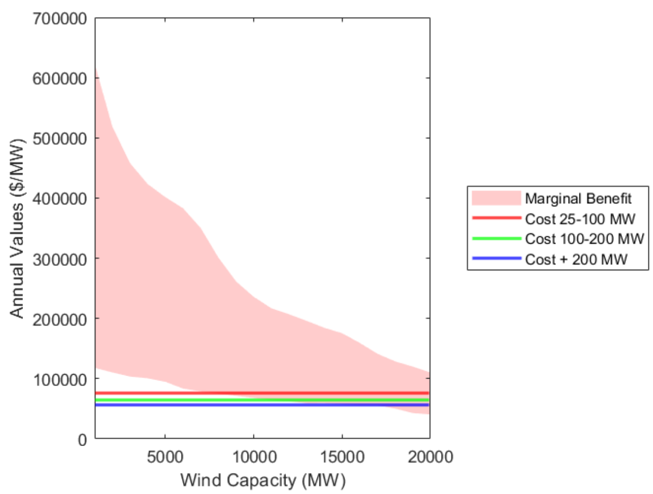

Note: The marginal benefit is the reduction of operating cost per installed MW of wind. The shaded area represents the range of marginal benefits for our representative scenarios.

7 Conclusion

In this project we have solved a unit commitment problem to determine the effect of wind penetration on the ISO New England. In the base case, we considered 21 scenarios for the installed wind capacity, from 0 MW to 20,000 MW. In the tax calculation, we incorporated a carbon tax equal to a social cost of carbon of $50 per ton. The analysis includes, but is not limited, to the change of average and industry profit, average price and price variance, consumers' costs, emissions, social welfare, and generation as we increase wind or introduce the carbon tax.

In this project, we chose 9 representative scenarios from the 90 possibilities present on the data. For this, we used a machine learning algorithm for agglomerative clustering in order to classify both our demand and wind capacity databases into three clusters each. Using these clusters we define our 9 representative scenarios. In addition, we realized a computing time analysis, concluding that using a linear objective function represents large computing time gains and minor continuity losses.

The results for the base case scenario show that the increase of wind capacity translates into lower operation. Indeed, wind will displace oil and natural gas from the market, and in less intensity coal, which are more expensive than wind. Recall that wind was modeled as zero marginal cost. This situation shifts the energy price distribution to lower median values. In addition, this price variance decreases as we allow the system to respond flexibly to the introduction of wind. The direct consequence of all these effects, is a decrease in the average and industry profit. In this sense, we recommend policy makers to plan accordingly for the long-run since it is expected the exit of controllable sources of electricity. Moreover, consumers' will face a lower cost due to the decrease in the energy cost. Indeed, we calculated that the marginal benefit of wind can be as high as \$5,000 when there is no wind, to around \$1,000 per MW when wind capacity is 20,000 MW. We conclude that if the policy maker is not planning to use a carbon tax, any investment below 3,000 MW would create social value. However, for values over this threshold a more careful calculation should be conducted.

In addition to the base case, we analyzed the situation in which we use a \$50 per ton carbon tax as we increase wind capacity. In this case, coal is the most affected fuel as it is the more polluting one. In fact, coal and oil facilities' average profit is below the profit these facilities receive in the base case. We expect that this reality will create the incentives for firms to invest in abatement technologies. Moreover, natural gas facilities will generate more than in the base case scenario, and as a consequence, the operating cost will be higher and there will be a surplus transfer from the consumers to the industry and the government. Indeed, the consumers per capita cost is 3 dollars higher than the base case due to the tax. Yet, emissions are smaller with the tax than without it, and thus, the social welfare is higher with the tax than without it. This situation shows that firms internalize the pollution cost they impose to society. Finally, the marginal benefit of wind capacity increases when we have a tax, and any wind capacity below 8,000 MW will be optimal in our 9 scenarios.

In conclusion, we propose to use a tax together with the introduction of wind facilities. Yet, this will generate a direct impact on consumers. We advice policy makers to use the tax revenue as a subsidy or direct transfer to compensate for the energy price increase due to the tax. Moreover, the government can use the tax revenue to provide cheap financing for abatement technologies so as to facilitate the access to these technologies in coal and oil facilities. If not, we fear that this facilities will exit the market in the long-run and we will loss capacity to respond to the residual demand. Finally, we conclude that for any type of facility, there is social value creation under our 9 representative scenarios if the overall installed wind capacity is below 8,000 MW. Yet, this is not the optimal level, since computing this amount would require using the 90 scenarios and estimating expectations for those scenarios.

References

Conti, J., Holtberg, P., Diefenderfer, J., LaRose, A., Turnure, J. T., & Westfall, L. (2016). *International energy outlook 2016 with projections to 2040* (Tech. Rep.). USDOE Energy Information Administration (EIA), Washington, DC (United States ....

Fischer, C. (2010). Renewable portfolio standards: when do they lower energy prices? *The Energy Journal*, *31*(1).

Frequently Asked Questions (FAQs) - U.S. Energy Information Administration (EIA). (n.d.). Retrieved 2020-11-19, from `https://www.eia.gov/tools/faqs/faq.php`

Gelabert, L., Labandeira, X., & Linares, P. (2011). An ex-post analysis of the effect of renewables and cogeneration on spanish electricity prices. *Energy economics*, *33*, S59–S65.

Kaufmann, R. K., & Vaid, D. (2016). Lower electricity prices and greenhouse gas emissions due to rooftop solar: empirical results for massachusetts. *Energy policy*, *93*, 345–352.

Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in germany. *Energy Economics*, *44*, 270–280.

Krishnamurthy, D., Li, W., & Tesfatsion, L. (2016). An 8-zone test system based on iso new england data: Development and application. *IEEE Transactions on Power Systems*, *31*(1), 234-246. doi: 10.1109/TPWRS.2015.2399171

Tra, C. I. (2016). Have renewable portfolio standards raised electricity rates? evidence from us electric utilities. *Contemporary Economic Policy*, *34*(1), 184–189.

US Energy Information Administration. (2019). *Electricity explained: Electricity and the environment.*

Winkler, J., Gaio, A., Pfluger, B., & Ragwitz, M. (2016). Impact of renewables on electricity markets–do support schemes matter? *Energy Policy*, *93*, 157–167.

Woo, C.-K., Horowitz, I., Moore, J., & Pacheco, A. (2011). The impact of wind generation on the electricity spot-market price level and variance: The texas experience. *Energy Policy*, *39*(7), 3939–3944.

Wozabal, D., Graf, C., & Hirschmann, D. (2016). The effect of intermittent renewables on the electricity price variance. *OR spectrum*, *38*(3), 687–709.

Appendix I: Generators Data Description

The data used for this project includes three files. The first dataset about generators provides key information about 76 generators located in ISO New England. Besides their location specified by New Hampshire (NH), Connecticut (CT), Southeastern Massachusetts (SEMASS), Vermong (VT), Rhode Island (RI), Maine (ME) or Western/Central Massachusetts (WCMASS)[18], our data comprises information about the following key technical parameters:

- fuel type: fuel oil (RFO), nuclear (NUC), coal (BIT and SUB) or natural gas (NGA1, NGA4, NGIR, NGLN, NGMN, NGPN, NGT1, NGT2, NGT4, NGTN). We assigned a numeric code of 1 for coal, 2 for natural gas, 3 for nuclear, and 4 for fuel oil;
- capacity: total installed capacity is 23,100 MW, with generators ranging from 43 MW to 1,244 MW (with an overall mean of 304 MW);
- minimum generation: when all 76 generators are turned on, the minimum generation is 8,694 MW and corresponds to 37.6% of the installed capacity. More specifically, the minimum generation represents 35% of capacity in the case of fuel oil plants, 80% in the case of nuclear plants, 40% in the case of coal power plants and 20% in the case of gas power plants;
- minimum up-time: from 1h to 24h, with an overall average of 10h;
- minimum down-time: from 1h to as much as 60h, with an average value of 12h;
- ramp rate: from 120 MW/h to 400 MW/h, with an overall average of 286 MW/h;
- start-up cost: from $0 to as much as $950,000, with an average of $85,000 and a median of $11,000 (i.e. distribution is heavily skewed);
- three coefficients for a dispatch cost: a no-load cost (in \$/h), a linear coefficient (in \$/MWh) and a quadratic coefficient (in $\$/(MW)^2h$)

Appendix II: Gurobi and the Branch-and-Bound Algorithm

As explained in the main text, the solver proceeds to a linear-programming relaxation and if the solution does not solve the integrability conditions, the solver uses a branch-and-bound algorithm. The variable which is used to define the branching becomes a so-called "branching variable". In a multivariate problem, we have to branch over many variables, which can be very complex. The Gurobi solver does this with a search tree. The problems generated by the search procedure are called the nodes of the tree, with the original problem constituting the root node. The leaves of the tree correspond to all the nodes at which the solver has not yet branched. In general, if the solver reaches a point at which it can solve or otherwise dispose of all leaf nodes, then it will have solved the original problem.

[18] While the entire ISO New England network is divided into eight load zones, our generator dataset only focuses on seven zones and omits the Northeastern Massachusetts/Boston (NEMASSBOST) one.

In this procedure, suppose the solver solves a linear-programming relaxed problem for some node in the search tree. If it happens that all of the integrality restrictions in the original problem are satisfied in the solution at this node, then the solver found a feasible solution to the original problem. Then, the solver assigns this node the category of a permanent leaf, or "fathomed", where it does not has to branch. In addition, the solver denotes this integer solution as the "incumbent". The incumbent solution is the best integer solution found in any part of the search tree. At the start of the search, there is no incumbent. If two integer solutions are found, the one with the best objective function value is recorded as the new incumbent.

Note that, if the linear-programming relaxation solution at the node which we are analyzing is infeasible, the node is also defined as "permanent leaf". Obviously if this node contains no feasible solution to the linear programming relaxation, then it contains no integer feasible solution. In addition, if an optimal relaxation solution is found, but its objective value is bigger (for a minimization problem) than that of the current incumbent, then this node cannot yield a better integral solution and again can be considered as a permanent leaf.

The question is how the solver defines the optimal incumbent. First, note that the objective value for the latest incumbent is a valid upper bound on the optimal solution of the given original problem given that we have a minimization problem. Second, the solver has a valid lower bound, sometimes called the best bound. This bound is obtained by taking the minimum of the optimal objective values of all of the current leaf nodes. Finally, the difference between the current upper and lower bounds is known as the "gap". When the gap is zero, the solver has effectively achieved optimality.

It is important to note that different nodes in the tree search can be processed independently. The solver takes advantage of this by running in parallel. However, the root node presents limited parallelism opportunities. Thus, models that explore large search trees can exploit cores quite effectively, while those that spend the majority of their runtime at the root node are more constrained.