The First Step to Solving the Duck Curve: Energy Efficiency

Written by Michael Fink

VEIC

Toll-free: (800) 639-6069

veic.org

Table of Contents

Abstract

In his 10-step proposed solution[1] for solving the "Duck Curve" distributed solar generation problem, Jim Lazar's first step is to implement energy efficiency measures that are "time-targeted" to peak hours. Here we present findings from retroactively examining groups of Vermont residential accounts that have implemented VEIC's incentivized energy efficiency measures to determine whether those measures have alleviated or exacerbated the Duck Curve problem. Seven different energy efficiency measures were examined among residential accounts. Among examined efficiency measures, major LED lighting installations were the closest match to the hoped-for time-targeted energy efficiency measures in Lazar's paper. LED installations significantly lowered consumption during both the morning and evening peak periods of the wider ISO-New England grid, while reducing the off-peak baseload much less severely. The other examined efficiency measure with a profound effect on the residential loadshape was cold-climate heat pumps. While heat pumps exhibited the unwanted outcome of increased consumption during peak periods, they also increased consumption during off-peak periods dramatically. Depending on the metric used or needs of local energy distribution engineers, cold-climate heat pumps could either be very helpful or a hindrance to solving the Duck Curve problem.

[1] Lazar, J. (2016). Teaching the "Duck" to Fly, Second Edition. Montpelier, VT: The Regulatory Assistance Project. Available at: http://www.raponline.org/document/download/id/7956

Introduction

What's a Duck Curve?

The "Duck-Curve" problem refers to a specific loadshape curve observed on a power grid, usually caused by distributed solar power generation. The figure below shows a pronounced example of Vermont's Duck Curve shape on a sunny day (June 7, 2017) compared to a cloudy day with similar day length and temperature (June 6, 2017) (data from ISO-New England).

The shape of the June 7 day in the figure above is known as the "Duck Curve" because of its resemblance to the belly of a duck in flight.[2]

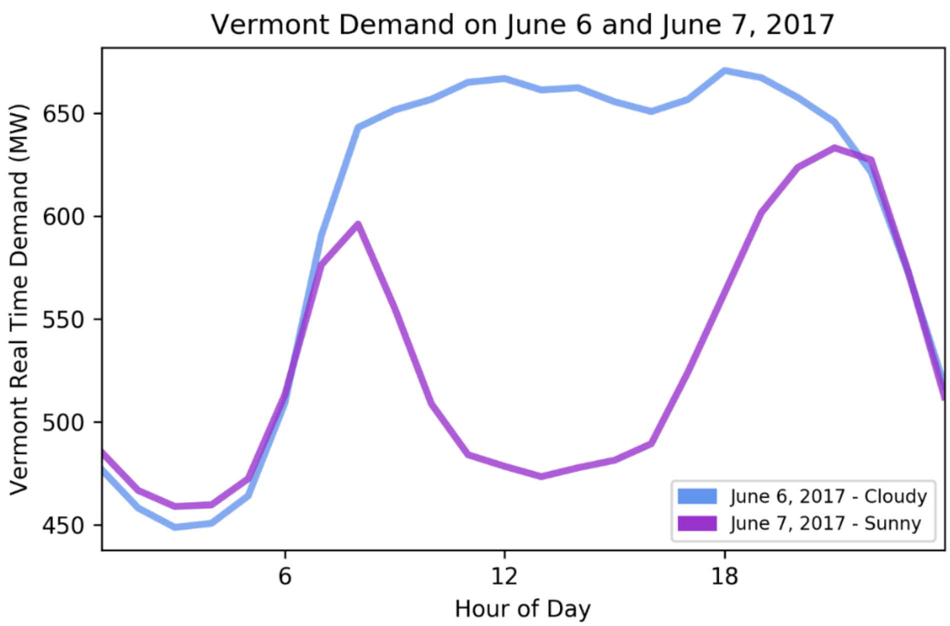

Duck Curve Problems

As the Duck Curve shape becomes more pronounced over the years, there are two problems that are likely to manifest: overproduction and steep ramp-up conditions.

The island of Oahu is experiencing overproduction problems.[3] As a relatively small, isolated grid, Oahu's distributed solar production can cause substantial voltage swings if a cloud moves across the island. The local electric company (Hawaiian Electric) claims these voltage swings could be damaging to their customers' property

---

[2] In Vermont, it has sometimes been referred to as the "Champ Curve" as an homage to the mythical lake monster rumored to inhabit Lake Champlain. It has also been suggested that the "Champ Curve is an apt moniker because it is possible that if distributed solar installations continue on their current trajectory, we may see the lake monster's belly (midday period of peak-production) dip underwater – pushing Vermont's real-time demand at peak solar production

[3] https://www.scientificamerican.com/article/analysis-clouds-over-hawaiis-roofto/

or even endanger their line workers. As a result, grid-tied residential solar installations have been curtailed in Oahu.

For now, an Oahu-type overproduction problem is not a concern for Vermont's power grid. Although Vermont and Oahu have similar total electricity consumption, Vermont benefits from being tied into a much larger grid on to which Vermont can sink excess production. If the larger New England grid (that Vermont is part of) starts to have a "Duck-Curve" loadshape that resembles Vermont's, overproduction may become a more serious problem.

California already suffers from the "ramp-rate" problem. In the case of California, it is predicted that by 2020, the state's demand in the three hours surrounding sunset will increase by 13,000 MW.[4] A ramp of this steepness requires flexible peaking power plants to meet demand. Given a ramp of this magnitude, California has two possible solutions:

1. Meet this demand by leaving baseload and intermediate plants idling during the day, even though they are not necessary because solar is handling the load.

2. Close baseload plants and some intermediate plants and replace them with flexible plants that can be used to satiate the sudden post-sunset demand.

We have one major problem with solution (1): the primary benefit of distributed solar power is that it has a very low carbon footprint. If it becomes necessary to leave natural gas plants or other carbon-intensive plants idling during the day in preparation for the sudden sunset-related demand, much of this low carbon footprint benefit is wiped out.

The problem with solution (2) is simply that it would be very expensive to replace our baseload fleet of generation plants with clean peaking plants. It would be difficult to convince either taxpayers or ratepayers that a large upfront investment in peaking plants or energy storage (either batteries or pumped-storage) is worth the benefit.

Of course, a third possible solution is to tolerate brownouts at sunset while electricity generation from plants comes online, but it is unlikely that Americans would be willing to tolerate brownouts or rolling blackouts during their dinners.

State of the Duck Curve in Vermont

Currently in Vermont, the Duck Curve is most pronounced on the very hottest and very coldest weekdays. The characteristic Duck Curve shape has become more severe each year for nearly two decades.

[4] Lazar, J. (2016). Teaching the "Duck" to Fly, Second Edition. Montpelier, VT: The Regulatory Assistance Project.

On the hottest days of summer. On the hottest days of the Vermont summer, the peak period of the day is usually at 7 or 8 PM. The following plot shows the average demand of the fifteen hottest days of Vermont's 2017 summer. The gray lines show the loadshape on each of the individual fifteen days.

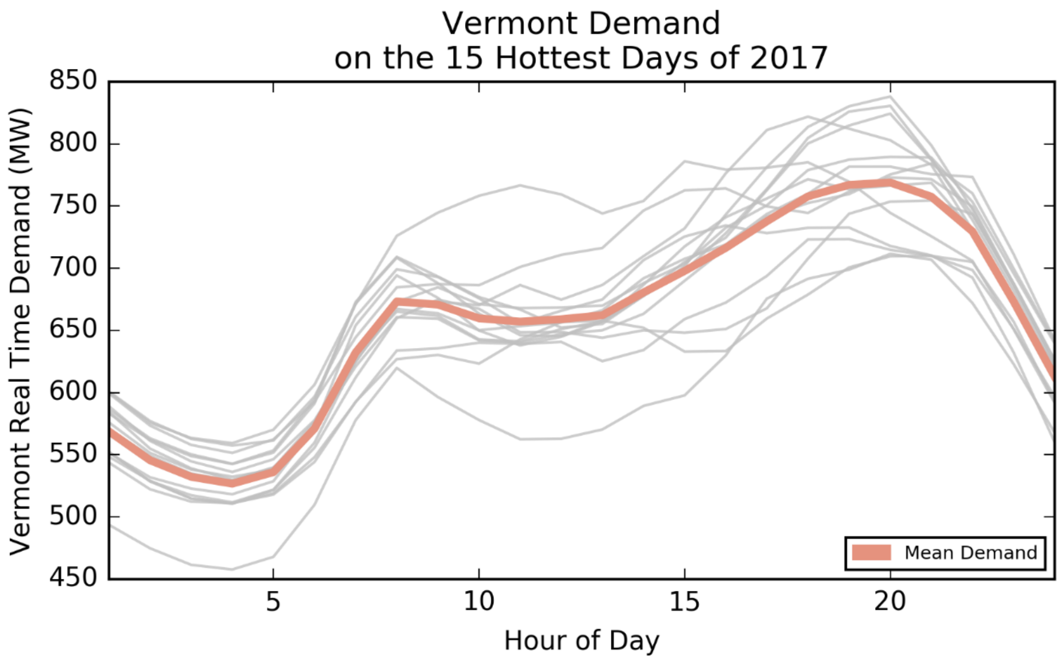

Three notes on this plot[5]:

1. There is a steep ramp-up in the morning hours.
2. There is a less steep ramp, but higher amplitude peak in the evening.
3. Even among the hottest fifteen days of 2017, the loadshape varies considerably.

On the coldest days of 2017. Vermont's 2017 winter was substantially warmer than average, but still had enough 'cold' days to gather fifteen that exhibit the characteristic 'cold' Duck Curve shape. The following plot shows how Vermont's real-time demand behaved on those days:

[5] An important point to note: for this entire paper and appendices, when a specific hour or a value from a specific hour is plotted on a graph, the value refers to the hour ***ending*** at the indicated time. For example, if it is indicated a house used 1.8 kWh at 20:00, that would mean that 1.8 kWh of energy was used between 19:00 and 20:00

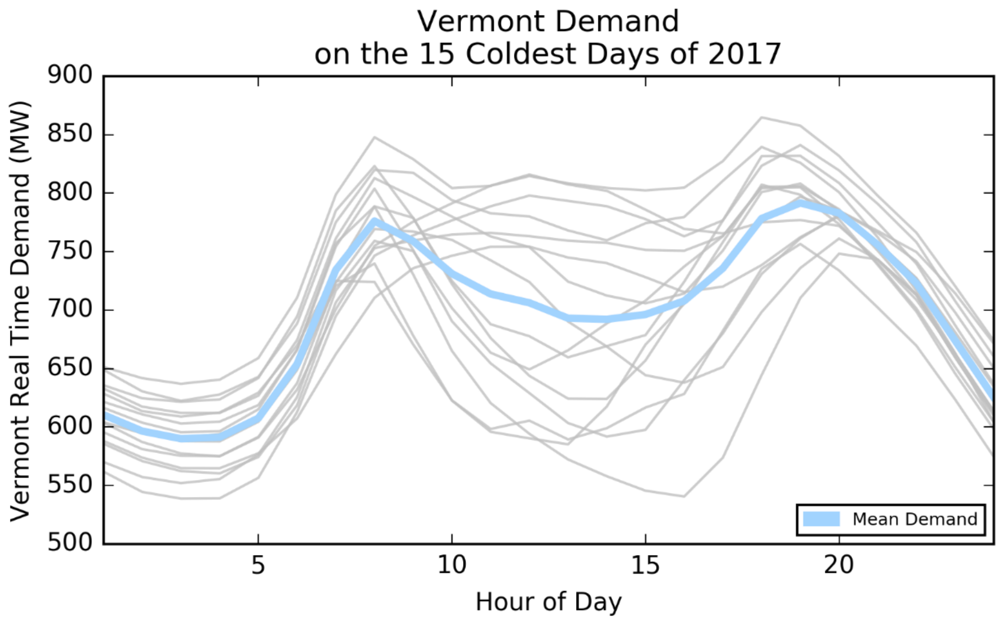

Evident here is how much the loadshape changes depending on the cloud cover. While Vermont's coldest days are typically clear, the cloudy days included among the fifteen are obvious – they have peaks during the midday hours. Also worth noting here: the average evening peak demand is slightly earlier in the winter compared to the summer. One possible explanation is that the need for lighting occurs earlier in the evening in winter than it does during the summer.

During the shoulder season. For the sake of comparison, here is the behavior of Vermont's grid during the fifteen days that required the least climate control (hereafter referred to as the 'shoulderiest' days).

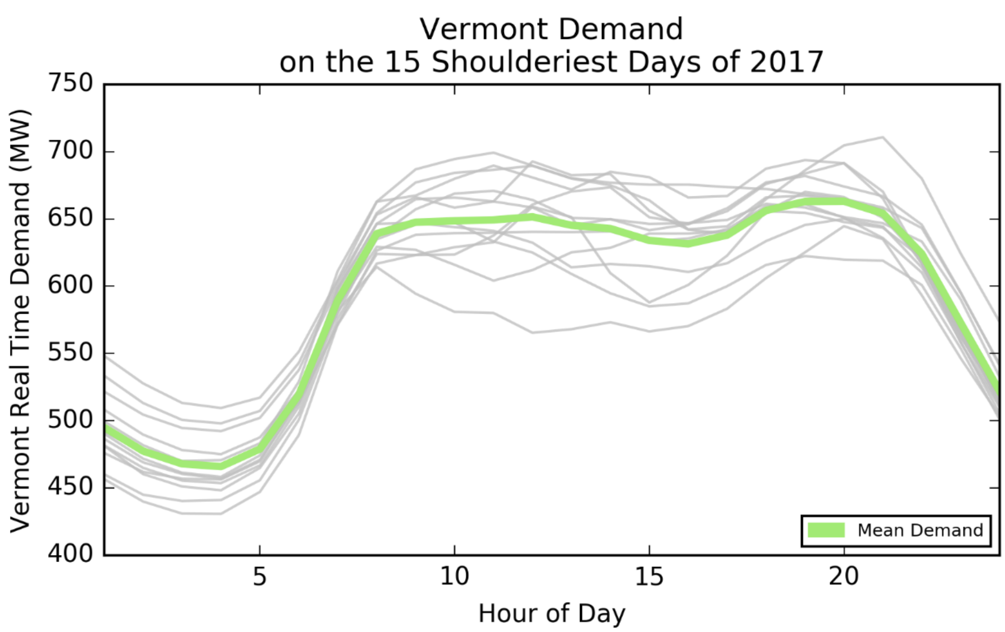

Notice that the demand is nearly flat during the normal waking hours of your average Vermont human. Those days among these fifteen shoulder days that were sunny still exhibit a mild Duck Curve shape, but because there is no need for residential climate control, the morning and evening peaks are much less pronounced than in both the “hot duck” and “cold duck” cases.

Vermont’s growing duck. Vermont’s Duck Curve is becoming more accentuated. Every year, more distributed solar generation is installed on our grid and midday in-state generation grows, but Vermont’s air-conditioners, heating systems, and other primary residential energy consumers remain mostly unchanged, creating similar demand each year. The result is a chunkier duck. Here are plots showing the growth of Vermont’s hot and cold Duck Curves over the past 14 years:

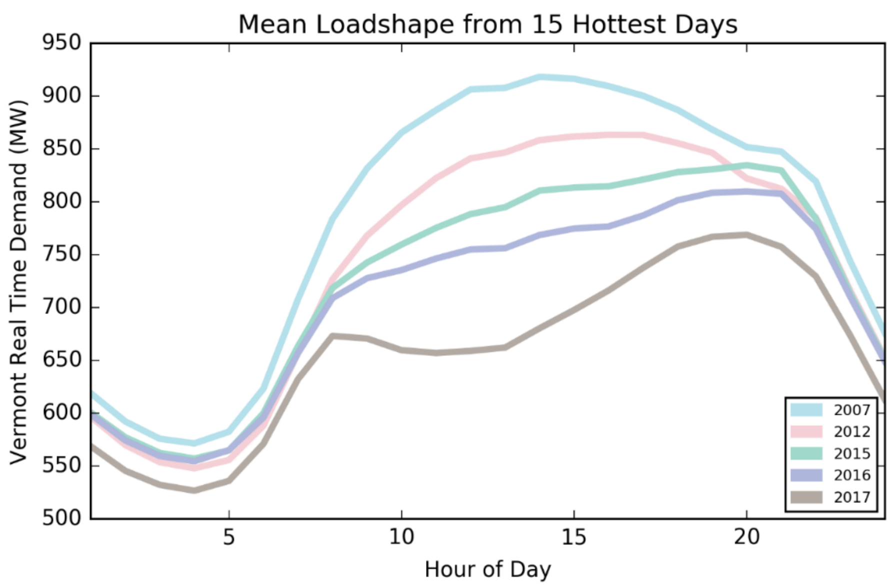

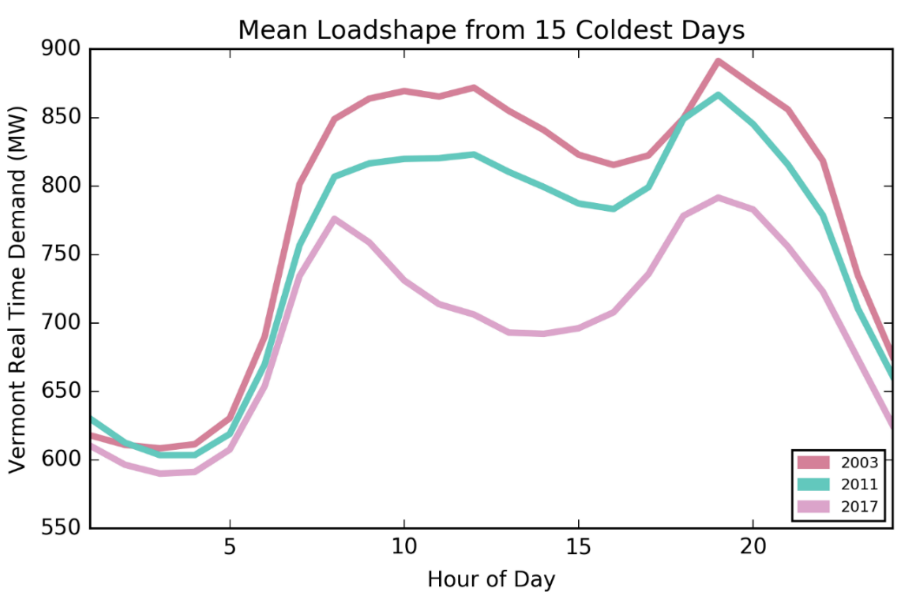

The most obvious feature of these plots is a good one for Vermont's inhabitants: Vermont's real-time demand, as seen by ISO-New England, is decreasing. Unfortunately, the reduced demand has also brought with it steeper ramps than Vermont has ever experienced before. As distributed solar generation becomes more widespread, Vermont should expect the ramps to continue their trajectory and become steeper.

How does Vermont's situation compare to others? While the Duck Curve problem is definitely getting worse, Vermont has a couple advantages that make its situation easier to deal with than Hawaii's overgeneration problem or California's ramp/peaking problem.

With both problems, Vermont has one big advantage over Oahu and California: Vermont is a very small part of a much larger grid. Vermont's grid is similar in size to Oahu's grid, but because Oahu is an island (electrically), overproduction results in voltage swings. If Vermont were to overproduce slightly, it could likely push the current back on to the larger ISO-New England grid which, as a whole, does not suffer from distributed overproduction problems.

California's grid ramps are concerning because 1) the ramps are steep and 2) California is the biggest fish in the pool. That is, if California has a steep demand increase, they will have to supply the electricity for that demand increase themselves because it is unlikely that there will be sufficient reserves elsewhere on the grid for purchase to satisfy their demand. Fortunately for Vermont, Vermont represents only 4% of the larger ISO-New England grid. If Vermont's demand increases by 20% or 40%, that's still just a tiny blip in the overall grid that is easily satiated[6].

Intention of this Examination. This is intended only as an examination of how employed residential efficiency measures in Vermont are currently affecting Vermont's Duck Curve. For reasons that are discussed in the Methods section of this paper, this should not be considered a proper experiment; rather, we urge the reader to consider this an examination of existing efficiency measures already in the wild.

Methods

As mentioned in the previous section, this is not a proper experiment with randomly selected residences comprising the experimental and control groups. Instead, a review of electricity use change among those who employed efficiency measures is presented here. Of course, this increases uncertainty and introduces systematic error.

Experimental Group. The group referred in this paper as the 'experimental' group is comprised of all Green Mountain Power residential customers that employed the examined efficiency measure (eg. refrigerator replacement or LED lighting) during

[6] Personal interview. Kim Jones & Dan Belarmino, GMP Distribution Engineers.

2016. This group size varied among efficiency measure groups, but usually numbered between 200 and 1000.

Control Group. The control group was selected using a custom-written computer program to select 1000 residential accounts from Green Mountain Power at random. The same control group was used to compare to each of the efficiency measure experimental groups. Of course, it is possible (or even probable) that members of the control group may have the efficiency measure being examined; despite this possibility, the overlap should have minimal effect on the outcome – GMP's residential service points outnumber the largest experimental group by a ratio of nearly 200 to 1.

Quality Check. Some accounts in both the control and experimental groups were discarded for poor data quality. For a residential account to be retained and included in the study, it must meet the following criteria:

1. The account must be active before January 1, 2015.
2. The account must be active as of December 31, 2017.
3. The account must have "good" AMI data for at least 90% of the days between January 1, 2015 and December 31, 2017.
4. For a day to qualify as 'good', at least 87 of 96 daily AMI energy readings must be present.

Method for Comparison. A difference-in-differences (DID) technique was used to compare the efficacy of different efficiency measures against the control population. Here's how the energy savings or loss was calculated, with a quick example of use of the DID method:

Suppose we'd like to calculate the "difference in differences" of average energy change for homes with heat pump dryer installations. The formula for the calculation is:

$$\Delta E = \left(E_{exp:post} - E_{exp:pre}\right) - \left(E_{control:post} - E_{control:pre}\right)$$

So if the heat pump experimental group changed from 19.3 kWh of consumption in the pre-period to 16.2 kWh of consumption in the post-period, while the control group changed from 24.1 kWh to 23.8 kWh. In this case, our calculation would be:

$$\Delta E = (16.2\,kWh - 19.3\,kWh) - (23.8\,kWh - 24.1\,kWh)$$

$$\Delta E = (-3.1\,kWh) - (-0.3\,kWh)$$

$$\Delta E = -2.8\,kWh$$

In this case, our heat pump experimental group is calculated to have saved 2.8 kWh of electricity.

The same technique is also used to create entire residual loadshapes, where the procedure outlined above is repeated for each hour of the day. These loadshapes are used extensively to determine whether a measure is saving energy or expending more energy than would be expected if the measure had not been installed. Here is an example of a plot with residual loadshapes to describe the increase/decrease in energy usage over the day, as a result of the efficiency measure installation:

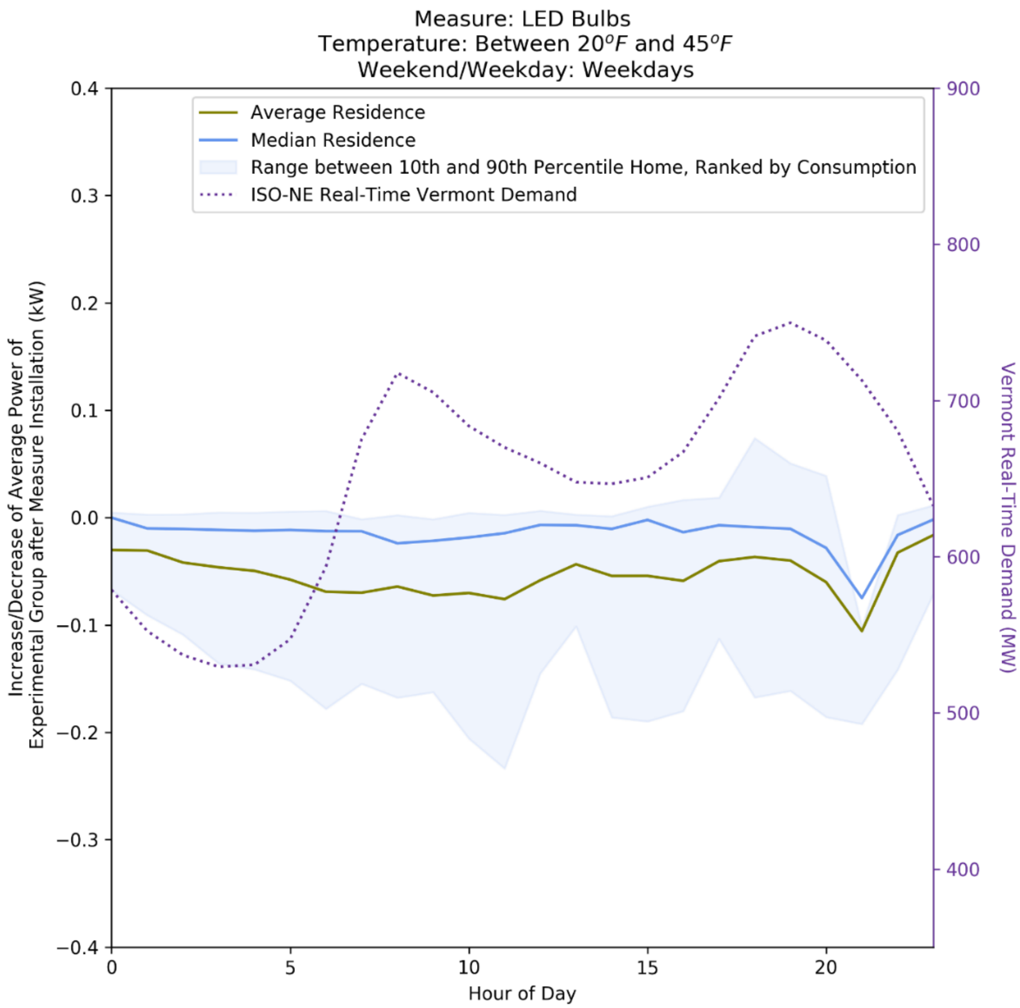

Shown here are data relating to LED lighting replacements. The purple-dotted line uses the right axis. This is the average ISO-New England data real-time demand in 2017 for all days that meet this bin's temperature and weekday criteria (in this case, weekdays with an average temperature between 20oF and 45oF. This ISO-New England data is here only for quick reference, so it is easy to see how the efficiency measure might affect the Duck Curve shape.

The green trace uses the left axis (as do all of the remaining graph features). The green trace uses the average energy usage of the experimental and control groups and, using the DID method, computes a “loadshape residual” for this temperature/weekday/efficiency measure combination. For example, suppose that at hour 6, the green line has a value of -0.07 kW; this means that on weekdays with an average temperature between 20oF and 45oF, the average home in the experimental group is using 70 W less power between the hours of 5:00 and 6:00.

The blue trace is identical to the green trace except that it is for the median residence rather than the average. Finally, the light blue shaded area bounds the 10th percentile home to the 90th percentile home (percentiles are in energy usage for that hour, not in the value of the DID; this means that the median and average values can (but rarely) fall outside of the blue shaded region).

Also used here are bar chart plots that use the DID method for other descriptive metrics. For example, consider this plot showing system coincident peak value changes on LED lighting replacements:

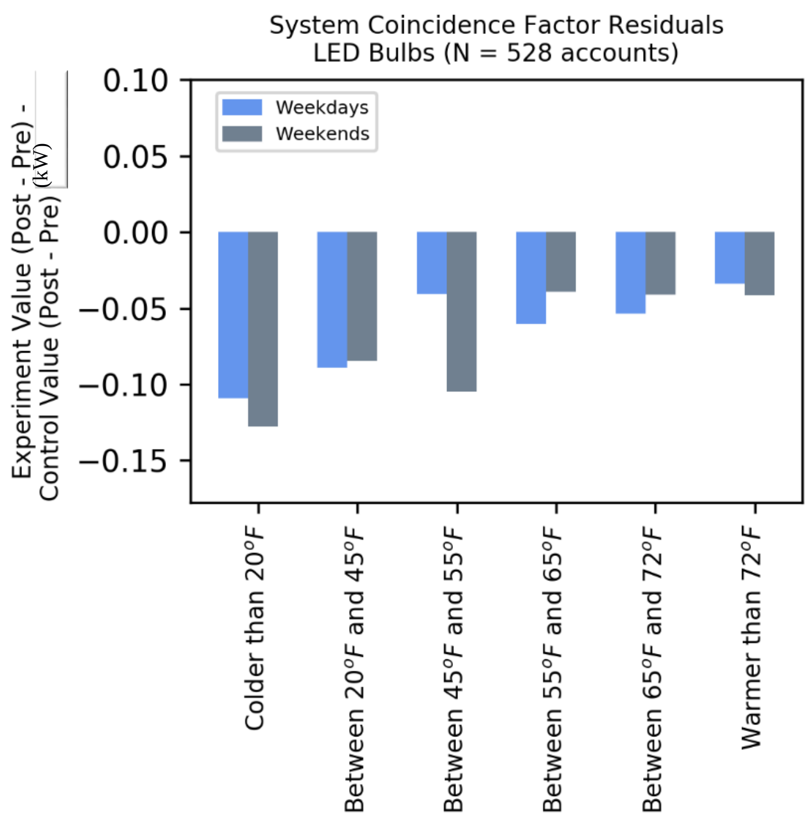

On this plot, for example, the blue bar having a value of -0.1 kW in the “Colder than 20ºF” category means that the coincident peak value for houses with LED Bulb installations dropped by about 0.1 kW on weekdays that had an average temperature of 20ºF or colder.

Results and Discussion

Among the seven efficiency measures, seven temperature bins, two weekday bins, and several different metrics calculated, there are too many individual results to include here. Instead, truncated results limited to the two most interesting efficiency measures are presented below.

Cold Climate Heat Pumps

Cold-climate heat pumps are undoubtedly the most interesting efficiency measure investigated here. These are devices that are usually used for space-heating sections of homes, but usually not relied on for whole-home heating. Their efficiency changes with outdoor (sink) temperature. Their output capacity also drops precipitously as the sink temperature drops[7]. At moderately cold temperatures, they are often one of the most economic means of space heating.

The ramifications of widespread heat pump installation on Vermont's Duck Curve are complicated. If measured by coincident peak value (the consumption of a house while the grid is peaking), heat pumps are damaging to the Duck Curve; homes that have them installed use significantly more energy when the ISO-New England grid is peaking:

[7] Most cold-climate heat pumps shut off between -5°F and -20°F

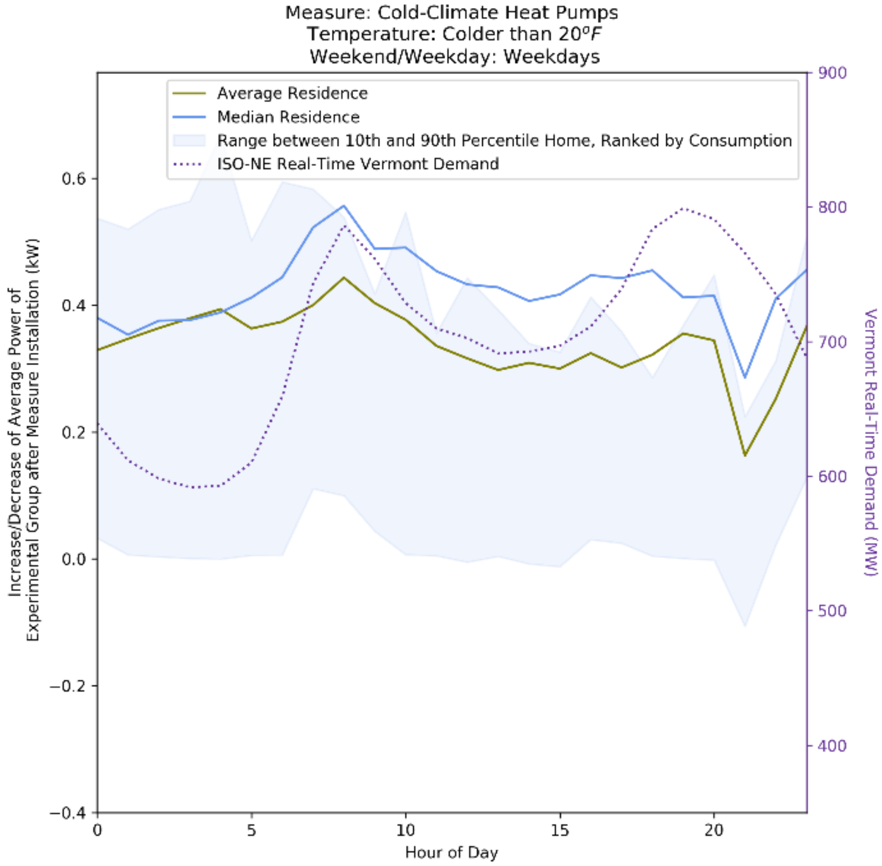
Measure: Cold-Climate Heat Pumps
Temperature: Colder than 20°F
Weekend/Weekday: Weekdays
Average Residence
Median Residence
Range between 10th and 90th Percentile Home, Ranked by Consumption
ISO-NE Real-Time Vermont Demand
Increase/Decrease of Average Power of Experimental Group after Measure Installation (kW)
Vermont Real-Time Demand (MW)
Hour of Day
0.6
0.4
0.2
0.0
−0.2
−0.4
900
800
700
600
500
400
0
5
10
15
20

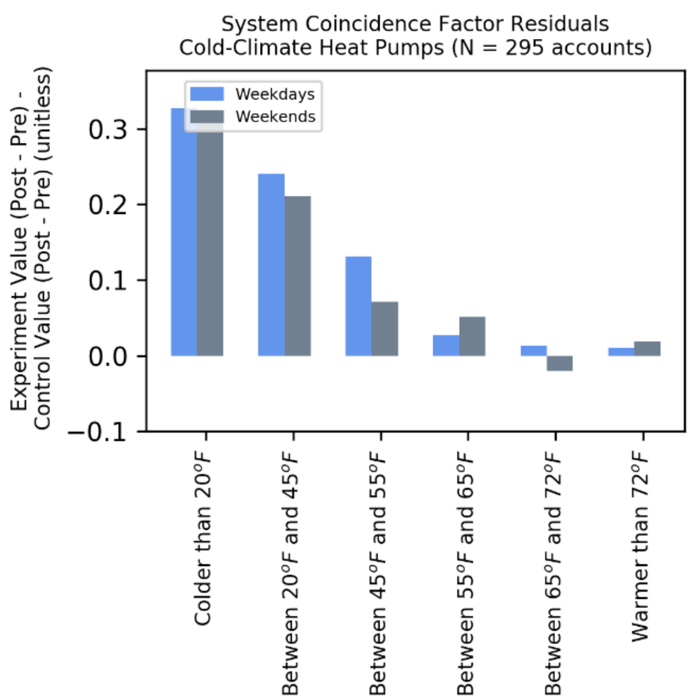
System Coincidence Factor Residuals
Cold-Climate Heat Pumps (N = 295 accounts)
Weekdays
Weekends
Experiment Value (Post - Pre) - Control Value (Post - Pre) (unitless)
0.3
0.2
0.1
0.0
−0.1
Colder than 20°F
Between 20°F and 45°F
Between 45°F and 55°F
Between 55°F and 65°F
Between 65°F and 72°F
Warmer than 72°F

Alternatively, if one were to use a metric like:

$$\frac{P_{peak}}{P_{baseload}}$$

Smaller values of this ratio of powers indicate a flatter loadshape. Because the cold-climate heat pump pushes up consumption uniformly at all hours, this metric would become smaller or closer to the ideal flat loadshape. So, despite the fact that cold-climate heat pumps tend to increase consumption and increase system coincident peaks, depending on what is most important to grid operators, this efficiency measure could be a help or a hindrance to solving the Duck Curve problem.

Lighting Measures

LED lighting installations were the closest thing to a sure-fire help to the Duck Curve problem of the measures investigated here. Here are the loadshape residuals at cold temperatures on weekdays:

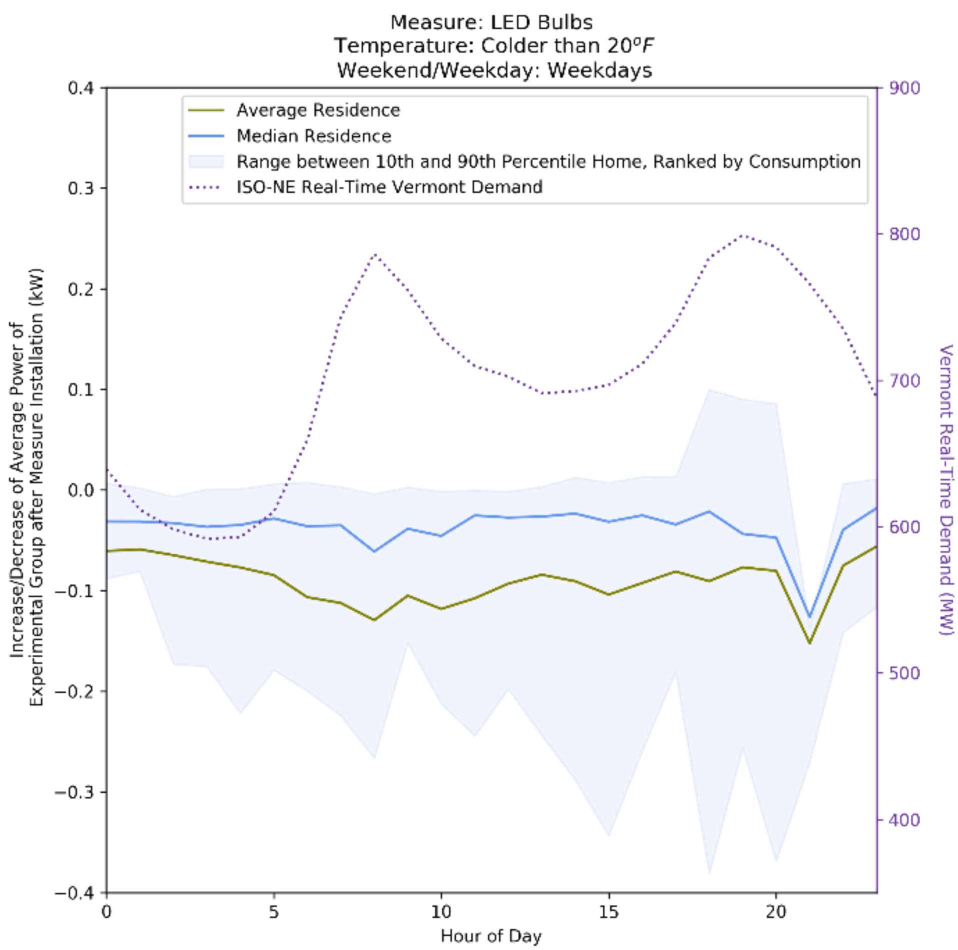

In cold weather, there is consistent savings, especially at the time of the morning Duck Curve peak. Savings associated with this measure are much less pronounced during warmer weather, probably because temperature is an effective proxy for day length. When the days are warmer (longer), there is less need for lighting. When the days are colder (shorter), there is greater need for lighting.

Unsurprisingly, LED bulb replacements offered consistent alleviation of the system coincident peak.

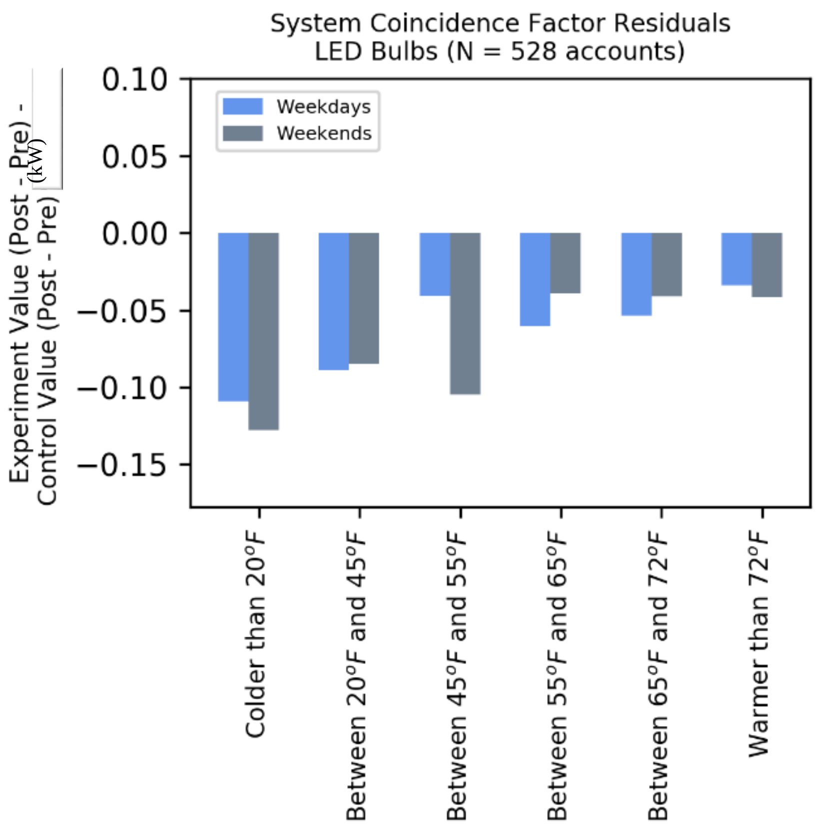

Other Results

Relative to cold-climate heat pumps and LED lighting installations, the other metrics had little effect at the grid peak times. Summarized in the chart below is the behavior of all metrics in cold weather/weekday conditions at Vermont's peak hour of consumption.

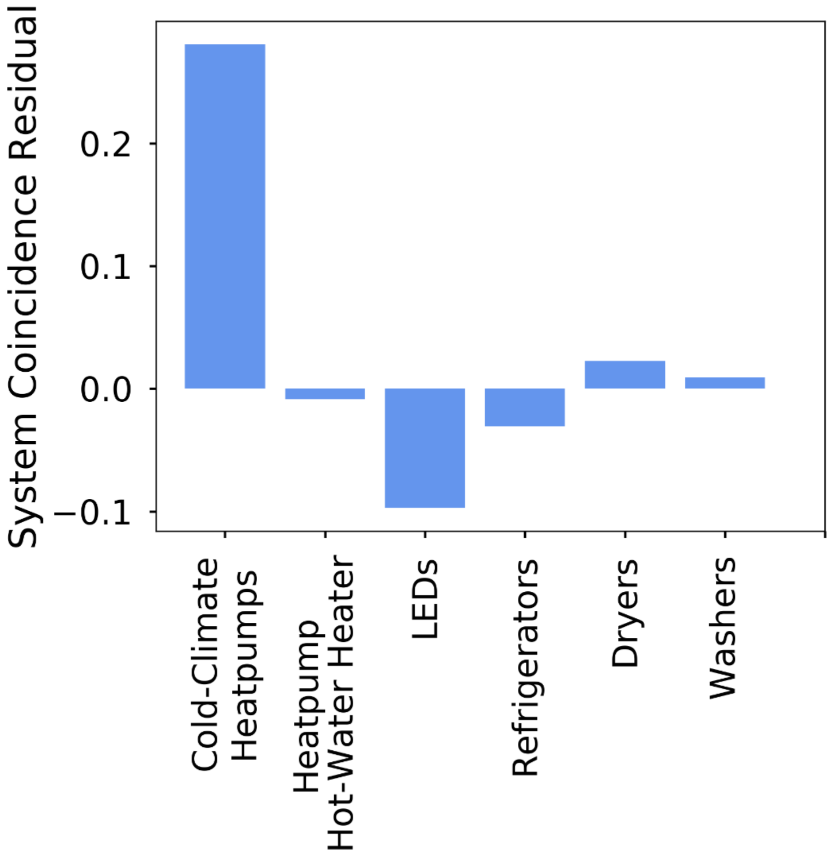

Aside from the two previously mentioned metrics, effects on the system coincident peak are minimal or not statistically significant.

Future Work

While these findings are suggestive, it should be reiterated that this is not a proper experiment with randomly selected participants in the experimental groups. If not cost-prohibitive, such a study might be the next logical step.

Of course, time-targeted energy efficiency measures is only the first of ten measures that Mr. Lazar suggests for confronting the Duck Curve problem. While some of the other nine suggestions will obviously alleviate the Duck Curve, there are some that merit testing before implementation.