



ECONOMETRIC ANALYSIS ON THE IMPACT OF WATER RATES ON RESIDENT USAGE AND CITY REVENUE: A STUDY IN SALT LAKE COUNTY, UTAH

Nate Nellis

Faculty Mentor: Scott Schaefer - Department of Finance, QAMO

Abstract

This project aims to examine what effect a change in municipal water rates has on water consumption. Data was provided by two cities in Salt Lake County, Utah and consists of monthly water usage for 26,510 households from July 2016 to July 2020, yielding a total of 1,192,463 observations. Beginning the month of July 2018, one city increased the rates charged per thousand gallons used, creating a natural experiment, while the other did not, serving as a control effect. City officials estimated a 5% increase in water rates would yield a 5% increase in city revenue. The year after the rate change however, water revenue decreased by 1%. The question of city officials was whether this was due to the rate change or other factors. The goal of this study is to determine if the rate increase had a measurable effect on water consumption and if so, whether or not it was enough to cause revenue city-wide to decrease. Using linear regression techniques and controlling for several variables that impact water consumption, I find average usage should have decreased between 119-238 gallons per month for the average household after the rate change. Even with the decrease in water usage, I estimate the average monthly bill should have risen by approximately \$3.43. This information is important, as it shows that without other factors that affect usage, the city should have expected to see a small impact in conservation and an increase in city revenue after the rate change. In normal years, this small 5% rate increase would grow city revenue by an estimation of \$889,718 per year, increasing the ability of the city to provide for essential public goods and services to its residents.

Table of Contents

1. Introduction - *page 5*

2. Literature Review - *page 5*

Figure 1: Treatment City Water Rate Table - *page 6*

3. Data - *page 7*

3.1 Data Overview - *page 7*

3.2 Summary Statistics of Treatment City Accounts - *page 7*

Figure 2: Account Level Summary Statistics - *page 7*

Figure 3: Lot Size Density Chart - *page 8*

Figure 4: Home Value Density Chart - *page 8*

Figure 5: Usage Density Chart - *page 9*

3.3 Summary Statistics by Month - *page 9*

Figure 6: Summary Statistics by Month - *page 9*

Figure 7: Density Plot of Usage by Month - Treatment City - *page 10*

Figure 8: Density Plot of Usage - Summer vs Winter - *page 11*

Figure 9: Density Plot of Usage - Treatment vs Control - *page 11*

Figure 10: Median Usage by Month by Fiscal Year - *page 12*

4. Demographics - *page 12*

Figure 11: Census Data - *page 12*

Figure 12: Map of Home Value - *page 14*

Figure 13: Map of Usage - *page 14*

Figure 14: Map of Average Cost (Bill Total Divided by Usage) - *page 15*

Figure 15: Average Cost as Usage Increases - 0.75-inch meter - *page 16*

5. Estimation and Results - *page 16*

5.1 Usage - *page 16*

Figure 16: Regression on Usage - *page 17*

Figure 17: Months on Usage - *page 18*

Figure 18: Other Variables on Residuals of Figure 17 - *page 19*

Figure 19: Regression on Usage with Fixed Effects - *page 19*

Figure 20: Difference in Differences Usage Regression with Fixed Effects - *page 20*

5.2 Bill Total - *page 21*

Figure 21: Regression on Bill Total - *page 21*

Figure 22: Regression on Bill Total in July 2018 Dollars - *page 22*

Figure 23: Difference in Differences Bill Total Regression with Fixed Effects - *page 23*

Figure 24: Difference in Differences Bill Total Regression with Fixed Effects (in July 2018 Dollars) - *page 24*

5.3 Elasticity - *page 24*

Figure 25: Average Price Elasticity Regression - *page 25*

Figure 26: Wealth Elasticity Regression - *page 26*

6. Discussion - *page 26*

Figure 27: Elasticity and Rate Change Comparison - *page 28*

Figure 28: Elasticity and Rate Change Comparison - July 2018 Dollars - *page 29*

7. Conclusion - *page 29*

8. Appendix - *page 30*

8.1 Acknowledgements - *page 30*

8.2 Data Dictionary - *page 30*

8.3 Figures, Charts, and Graphs - *page 32*

Figure 29: Temperature Data for each Month/Year - *page 32*

Figure 30: Inflation Rate Table - *page 33*

Figure 31: Year on Assessment - *page 34*

Figure 32: Bill Total Density Chart - *page 34*

Figure 33: Density Plot of Usage by Month - Treatment City - *page 35*

Figure 34: Map of Lot Size - *page 35*

Figure 35: Map of Bill Total - *page 36*

Figure 36: Home Value on Average Cost - *page 36*

Figure 37: Average Cost as Usage Increases - 1-Inch Meter - *page 37*

Figure 38: Average Cost as Usage Increases - 1.5-Inch Meter - *page 37*

Figure 39: Months on Rain Regression - *page 38*

Figure 40: Temperature on Rain Regression - *page 38*

Figure 41: Months on Snow Regression - *page 38*

Figure 42: Temperature on Snow Regression - *page 39*

Figure 43: Rain and Month Interaction Regression - *page 39*

Figure 44: Rain on Usage - *page 40*

Figure 45: Temperature and Month Interaction Regression - *page 40*

Figure 46: Usage Regressions Compared - *page 41*

Figure 47: Regression on Usage - Home Value Buckets - *page 42*

Figure 48: Temp on Usage - *page 43*

Figure 49: Other Variables on Residuals of Figure 48 - *page 43*

Figure 50: Account Variables on Usage - *page 44*

Figure 51: Other Variables on Residuals of Figure 50 - *page 44*

Figure 52: Rate Changes Table - *page 44*

Figure 53: Regressions on Usage by Month - *page 45*

Figure 54: Usage Fixed Effects - Block 1 - *page 45*

Figure 55: Usage Fixed Effects - Block 2 - *page 46*

Figure 56: Usage Fixed Effects - Block 3 - *page 46*

Figure 57: Usage Fixed Effects - Block 4 - *page 47*

Figure 58: Median Usage Difference in Differences - *page 47*

Figure 59: Bill Total Regressions Compared - *page 48*

Figure 60: Bill Total Regressions Compared (in July 2018 Dollars) - *page 48*

Figure 63: Regressions on Bill Total by Month - *page 49*

Figure 64: Regressions on Bill Total by Month (in July 2018 Dollars) - *page 50*

Figure 61: Regression on Bill Total - Home Value Buckets - *page 51*

Figure 62: Regression on Bill Total - Home Value Buckets (in July 2018 Dollars) - *page 52*

Figure 65: Regression on Bill Total with Fixed Effects - *page 53*

Figure 66: Regression on Bill Total with Fixed Effects (in July 2018 Dollars) - *page 53*

Figure 67: Marginal Price Elasticity Regression - *page 54*

9. Works Cited - *page 55*

1. Introduction

Utah is the third driest state in the United States (Nag 2019) with average annual precipitation of 15.78 inches (Figure 29) compared to the national average of 32 inches (NOAA 2018). Due to the climate and large need of water, Utah is the state with the third highest consumption per capita with domestic water usage at 186 gallons a day, behind Maine at 190 gallons and Idaho at 187 (Kearney 2014). In many instances, water is provided by municipal governments, entities which must balance the imperative to provide residents with water against the need to generate revenue in order to provide for other services, and in some cases to encourage conservation. As policy makers strive to balance these objectives, it is important for them to know what effects they can expect from changing the rates at which they bill water. This paper aims to answer that question with a focus on residential water rates. As far as I am aware, this is the first study of this kind local to Utah.

In this paper, I will be performing econometric analysis on data provided by Sandy City, Utah. The data consists of monthly residential water usage from 21,604 observations from July of 2016 to June of 2020, resulting in 1,022,922 total observations. The city raised their water rates in July of 2018, creating a natural experiment I draw on to form conclusions about the impact of water rates on usage. In this paper I will refer to Sandy as the treatment city. Draper City borders the treatment city and did not change their rates during the same period of time I study. I will use an additional 169,541 observations provided from Draper to control for influences outside the data that would have impacted both cities equally. I will refer to Draper as the control city in this paper. The end result of my analysis is a difference in differences regression with fixed effects that explains the impact the change in rates had on usage. My analysis finds that the rate change caused consumer usage to decrease by 119-238 gallons per household in the average month (Figure 20 and Figure 19), but increased the average monthly bill total by \$3.26 in terms of July 2018 dollars (Figure 24), which results in an increase in yearly city revenue of \$845,042.

2. Literature Review

At the pinnacle of the residential water debate is whether pricing techniques should be used over command and control methods. Command and control is when the government uses non-price directives, permits, and legislation to encourage water conservation, especially in times of drought (Porto 2004). The most common command and control techniques include outdoor watering restrictions and low-flow fixture requirements (Olmstead 2009). The main criticism of command and control methods is that they are inflexible and require all parties to undertake similar shares of burden, regardless of individual cost (Olmstead 2009) and that these instruments do not take into account the differences in control costs for those who use the resource (Porto 2004). Furthermore, it has been repeatedly found that these methods are less cost effective than market-based regulations in controlling water usage (Collinge 1994) (Krause 2003) (Brennan 2007). Interestingly, nonprice programs are more expensive to society than water price increases (Olmstead 2009). Since this cost is often less visible, within tax bills for example, households often object more strongly to water rate increases than they would to increases in city tax revenue (Timmins 2002).

The main pushback against the market-based approach is that lower-income families show larger demand responses to price increases than higher-income households (Agthe 1987) (Renwick 1998) (Mansur 2012). This is due to the higher willingness to pay for water by higher-income families and may lead to some households not being able to meet their water needs. Since water is considered “vital to human life” (152 Cong. Rec. E1083, 2006), this can be an issue, especially where governments themselves often undertake the role of providing water to their residents. The United States Environmental Protection Agency has determined water bills are affordable at 2.0-2.5% of median household income (U.S. EPA 2002). This figure is normally used when determining affordability of pricing structures (Mack 2017), although it has been questioned by some as not measuring affordability on the household level, but rather the community level (Teodoro 2018). In some cases, a water rebate is proposed for low income families or cities may offer special “lifeline” rates to ensure water bills are not unduly burdensome for low-income households (Olmstead 2009).

One market-based approach that is often used to decrease the impact on lower-income consumers is the increasing block tariff approach, which has become the prevailing method among market-based approaches (Mack 2017). Increasing block tariffs place differing ranges of usage into several blocks and a higher marginal price is charged for each successive block. See Figure 1 below for an example of this, which is the treatment city’s increasing block structure. Not only is each successive gallon increasing the household’s bill total, but it is doing so at a greater rate than the previous gallon (subject to the tier structure). Pricing water in this way allows for “necessary” usage to

be priced at a lower rate than “excess” usage, meaning that low-income families are not as heavily impacted by the prices since their usage is generally within the lowest blocks.

Figure 1: Treatment City Water Rate Table

	2017	2018	2019	2020
Base Rate (Monthly Charge)				
0.75-inch meter	\$17.00	\$19.00	\$19.95	\$19.95
1-inch meter	\$23.20	\$25.20	\$26.46	\$26.46
1.5-inch meter	\$29.40	\$31.40	\$32.97	\$32.97
Cost Per Thousand Gallons*				
Block 1	\$1.71	\$1.71	\$1.80	\$1.80
Block 2	\$2.62	\$2.62	\$2.75	\$2.75
Block 3	\$3.09	\$3.09	\$3.24	\$3.24
Block 4	\$3.54	\$3.54	\$3.72	\$3.72

	Block 1	Block 2	Block 3	Block 4
*Block tiers based on meter size and consumption				
0.75-inch	1-6	7-40	41-80	81+
1-inch	1-7	8-65	66-130	131+
1.5-inch	1-35	36-200	201-400	401+

Source: Sandy City Approved Annual Budget Fiscal Year 2019-2020

The main difficulty with the increasing block structure is customers often have difficulty understanding how their bill is calculated (Nieswiadomy 1989), which means they don’t always follow how utility maximization theory would indicate (Nataraj 2011). Utility maximization theorizes that rational consumers will balance the marginal utility gained from a good with the marginal cost that must be spent on that good. This means they will stop spending money on a product when the cost of doing so does not provide them with enough marginal benefit. What causes confusion in this model is that with the increasing block structure, the marginal cost of water is different depending on the block, and a consumer’s usage is often spread across several blocks. In many cases, pricing information is not available to consumers, only the bill itself, and only about 10% of customers seek out how their bill is calculated (Carter 2005). Even when this information is available, many consumers do not understand the structure and even when the structure is understood, an additional confounding effect is customers often do not know their exact water usage until they receive their bill in the next month (Bell 2008). During the month, consumers usually do not notice when they have passed from one tier to the other, and so are less able to respond to marginal price changes (Bell 2008) (Nataraj 2011). This can cause consumers to be more influenced by non-price factors or by using the previous month’s bill as a proxy for the current month’s usage decisions (Wichman 2014).

The assumption that water users are rational consumers is an important aspect when creating an elasticity. Elasticity measures the percentage change in consumption given the percentage change in the price (Schotter 2008). If price increases by 1% and usage decreases by more than 1%, then this is considered elastic, but if the usage decreases by less than 1%, then it is considered inelastic. Knowing the elasticity for a product is important, because when considering rate increases, elasticities can help to explain how consumers will react at different levels of the price. Water is generally considered to be an inelastic good, and a meta-analysis of 124 estimates of water price elasticity found about 90% of them were between 0 and -0.75 (Espey 1997). This means that in the average case, people will not decrease their usage by much when water rates are increased. Normally, elasticities are calculated by measuring marginal price, but as marginal price increases in an increasing block structure, some studies have used average price (Foster 1979) (Ito 2012). Unfortunately, average price tends to produce excessively large estimates (Billings 1980), resulting in many studies proposing other measures (Taylor 1975) (Billings 1980) (Nieswiadomy 1991) (Barkatullah 1996) (Bell 2008).

This study contributes to previous research in several ways. This is the only study of note on residential water rates in Utah. Utah is one of the fastest growing areas in the country and this information will be valuable as water providers adjust for the incoming residents it is required to serve (Kearney 2014). This study also contributes in drawing from its large volume of data, which consists of 21,604 accounts over 48 months yielding 1,022,922 total observations. This data is then compared against a control set of 4,906 accounts over the same 48 months yielding

169,541 total observations. In total this results in 1,192,463 observations for 26,510 accounts. Many other studies only use aggregate data (Hortová 2013) (Bell 2005), or do not have as many observations (Nataraj 2011 [103,044 observations]) (Wichman 2014 [96,702 observations]). Account-level data is also included, such as lot size, home value, and meter size, which gives greater insight at the granular level.

3. Data

3.1 Data Overview

Data was mainly provided by the Public Utilities department of the treatment city. Data consists of the water usage of 21,604 accounts during the 48-month span between July 2016 and June 2020 with a total of 1,022,922 individual observations. The bill total for each month is determined using a rate schedule provided by the city (Figure 1). Additional data provided by the city includes the meter size of the account, the size of the lot where the account resides, and location data. In addition to latitude and longitude, the city provided GIS Parcel IDs for the accounts and these were used to retrieve the 2020 home value from the Salt Lake County Assessor’s website for the respective accounts. Data was retrieved June 6-8 by running a web automation script in Python. The 2020 home value is used as a variable that correlates with household income. Only the 2020 home assessment is used, as a yearly assessment was found to be strongly correlated with time and would affect the result of the study (Figure 31). Inflation was accounted for using the U.S. Bureau of Labor Statistics’ CPI Inflation Calculator and many statistics are denoted in terms of July 2018 dollars throughout the paper (Figure 30). Weather data for the 48 months were retrieved from the National Weather Service Forecast Office in Salt Lake City, consisting of average temperature, and monthly precipitation in the form of rain and snow. For a full table of this data, see Figure 29 in Appendix 8.3.

Additional data was provided by the Recorder of the control city. Data is similar in format to the treatment city for usage, monthly bill total, and meter size, but is lacking in that no location data, home value, or lot size was provided. As this data is meant to only serve as a control against the treatment city, this was deemed acceptable, and account-level fixed effects were used to address this lack of information. A complete data dictionary is available for reference in Appendix 8.2.

3.2 Summary Statistics of Treatment City Accounts

This section will be devoted to looking at the summary statistics of account-level variables, such as lot size, home value, and usage. Figure 2 below gives these summary statistics and afterwards a further look is taken into individual variables.

Figure 2: Account Level Summary Statistics


	Lot Size	Home Value	Usage	Bill Total
Min	0.02	\$ 104,000	0.00	\$ 17.53
Mean	0.27	\$ 423,805	18.24	\$ 65.89
Max	6.05	\$ 4,205,300	662.70	\$ 2,334.42
Q1	0.19	\$ 319,600	4.00	\$ 26.41
Median	0.22	\$ 374,600	7.86	\$ 35.32
Q3	0.27	\$ 474,000	25.82	\$ 83.61
S.D.	0.19	\$ 186,631	23.16	\$ 69.04

Bill total is in July 2018 dollars

The median of lot size is approximately a quarter of an acre. The inner quartile range, which is the difference between quartile 3 and quartile 1, is close around the median, which indicates that lot size has a small variance. This is further illustrated by the dense cluster around 0.22 as seen in Figure 3 below with very few observations towards the tail of the distribution.

Figure 3: Lot Size Density Chart


Figure 3: Lot Size Density Chart



The standard deviation of home value shows much more variation than lot size. The median is also much less than the mean, due to a large skew with several high value outliers. Parts of the treatment city are very wealthy, with 344 accounts with a home value over \$1,000,000. This increased variation in the distribution can be seen clearly in Figure 4 below.


Figure 4: Home Value Density Chart

Figure 4: Home Value Density Chart



In comparing the median and the mean of Usage, there are some very large outliers, as the mean is 2.3 times greater than the median. The standard deviation is also quite large, which indicates a large amount of variance in the data. In looking at Figure 5 below, the mode is at 3,000 gallons and usage in general appears to be mostly under 10,000 gallons. Due to the high correlation with usage, the density plot for bill total is not given here, but is included in Appendix 8.3 (Figure 32). The mean for bill total is \$65.34 and the median is \$35.18, which shows that there is a skew from large outliers.

Figure 5: Usage Density Chart



3.3 Summary Statistics by Month


As there is a very marked difference in how usage changes depending on the month (Figure 17), this section will examine how the summary statistics of the dataset change with each month on average.

Figure 6: Summary Statistics by Month

Month	Median Usage	Median Bill Total	Average Temperature	Average Rainfall	Average Snowfall	Observations
Jan	4.5	27.19	33.3	1.6	12.9	92,468
Feb	4.3	26.84	37.0	1.3	11.5	83,205
Mar	3.9	26.29	46.4	2.9	8.1	82,205
Apr	4.6	27.53	52.7	2.2	0.7	87,082
May	8.7	37.02	62.4	1.5	0.0	84,599
Jun	22.6	73.98	72.8	0.6	0.0	85,864
Jul	37.7	114.58	83.3	0.3	0.0	84,567
Aug	38.8	117.97	79.8	0.3	0.0	89,352
Sep	32.3	100.00	68.1	1.4	0.0	83,547
Oct	14.5	52.50	52.6	1.0	0.0	84,774
Nov	5.4	29.00	43.9	1.3	4.3	82,953
Dec	4.3	26.99	32.1	1.4	8.6	82,306


As expected, usage, bill total, and temperature all rise in the summer months and rainfall and snowfall drop. Figure 7 below shows the median usage by month along with inner-quartile bands. During the winter months, about November through April, usage is tightly clustered around 4,000 gallons and has very little variation. During the peak months of summer, between June and September, the usage at the median is normally above 30,000 gallons and the variation expands considerably.

Figure 7: Density Plot of Usage by Month - Treatment City



Given this large peak, it is useful to summarize summer and winter months together and see the difference between them. In Figure 8 below, usage is grouped by season, where November through April are winter months and May through October are summer months. The dotted lines show median usage for both seasons. As seen previously in Figure 7, there is a large cluster during winter months around 4,000 gallons. This changes in summer months, where there are two modes, one again around 4,000 and another near the summer median at 25,150 gallons. This shows that during summer months there is a portion of the population that does not increase their usage from winter months, while another section of the population increases their usage by a large amount.

Figure 8: Density Plot of Usage - Summer vs Winter



A similar graph to Figure 7 for the control city can be found in Appendix 8.3 (Figure 33). Below is a comparison between the median usage of the two cities, showing how correlated they are with each other. While the treatment city tends to have a higher median usage in each month, they run together very closely, which gives validity to the effectiveness of the control city in this study.

Figure 9: Density Plot of Usage - Treatment vs Control





Figure 10: Median Usage by Month by Fiscal Year



While the focus of this paper is not about how the COVID-19 pandemic impacted water rates, it is something that needs to be addressed with additional control effects. Figure 10 shows median usage by month separated into the four fiscal years being studied. The red line, which represents the 2020 fiscal year, shows a large increase in usage for April through June compared to previous years. This is likely due to people being at home more, as they will use dishwashers, bathrooms, and sinks more frequently. In Figure 16, a regression of various factors on usage, the coefficient on the COVID-19 dummy variable indicates an average increase of 8,431 additional gallons per month during those three months. This is not something focused on extensively in the paper, but is factored into my analysis.

4. Demographics

Figure 11: Census Data

Measure	Sandy	Draper	County	Utah
Population estimates July 1, 2019	96,380	48,587	1,160,437	3,205,958
Population, Census, April 1, 2010	87,461	42,274	1,029,655	2,763,885
Population per square mile 2010	3,824.8	1,405.2	1,387.1	33.6
Persons per household	3.09	3.32	3.00	3.13
Households 2014-2018	30,722	13,542	369,429	957,619
Owner-occupied housing unit rate	78.0%	80.0%	66.7%	69.9%
Median value of owner-occupied housing units	\$334,200	\$435,200	\$281,800	\$256,700
Median selected monthly owner costs (with mortgage)	\$1,732	\$2,218	\$1,587	\$1,497
Median selected monthly owner costs (without a mortgage)	\$502	\$609	\$465	\$418
Median household income (in 2018 dollars)	\$91,836	\$115,618	\$71,230	\$68,374
Per capita income in past 12 months (in 2018 dollars), 2014-2018	\$36,862	\$44,067	\$31,648	\$28,239
Total retail sales per capita, 2012	\$24,615	\$29,235	\$16,148	\$13,317
Bachelors degree or higher, percent of persons age 25+	40.8%	50.1%	34.5%	33.3%
Persons in poverty, percent	5.2%	5.0%	9.0%	8.9%

All above facts were found in the Census.gov Quick Facts pages for each city, Salt Lake County, and the state. Data was retrieved on 9/18/2020.

Figure 11 (above) shows demographic data provided by the Census Bureau's Quick Facts pages for the treatment city, control city, county, and state. The treatment city has almost twice the population of the control city and has just under one-tenth the population of Salt Lake County. The treatment city has a much higher population density as well, double the county and control city levels. According to the Census estimate there are 30,722 households in the treatment city, but only 21,604 accounts in my dataset, a little over 75% of the household estimate. This difference is due to various factors. According to city officials, over 2,000 homes have water provided to them by Salt Lake County, and are not included in the treatment city system. Another potential reason is due to how the Census Bureau counts a household, as apartment complexes are not included in my data. My data only includes accounts with an individual residential meter between 0.75 and 1.5 inches. The data was further filtered by removing any observation that did not have a lot size, home value, or location coordinates. This filtering removed less than a thousand accounts from the final dataset. As this study is focused on how water rates impacted usage for residential customers of the treatment city, the data in this study is sufficient to meet that goal.

The poverty level of the treatment city is almost half of the county and state levels, indicating the area is more affluent. This is also supported by the median household income, which is over \$20,000 more than the county and state levels, although the per capita income is only \$5,000 more than the county level. The median value of the owner-occupied housing unit is over \$50,000 higher than the county level, and over \$75,000 higher than the state level, although it is \$100,000 lower than the control city's levels. The median home value in my dataset is \$374,600, which is about \$40,000 higher than the census estimate. Other than the excluded data mentioned previously, this can be attributed to inflation, as I use the 2020 home assessment, as well as differences in how the Census Bureau and Salt Lake County calculate home value.


The Environmental Protection Agency advocates that the average cost of water should be between 2.0 and 2.5% of a community's median household income (U.S. EPA 2002). In the treatment city, annual median household income is \$91,836, thus monthly income is \$7,653. For a bill to be 2% of monthly income, as the EPA has suggested, that would entail an average bill of \$153.06, which is 2.3 times higher than the current average bill total, \$65.89. Currently the average bill total is 0.86% of median household income. Even in summer months, when usage and water bills are at their highest, the average bill is \$100.77, which is 1.3% of median household income, still less than the EPA standard. This doesn't necessarily mean water rates need to be drastically increased, as some question the true affordability of the EPA standard (Teodoro 2018) (Olmstead 2009), but it does allow some leeway into the increases policymakers can lobby for before they hit the EPA buffer. According to officials of the treatment city, the government does not subsidize the water. The full cost to the city of providing water is entirely met through the revenue generated from the water fund. Due to the increasing block rate structure, an individual may pay more or less than their exact cost, but generally low users are those who will pay below their cost and high users will pay more than their actual cost. When compared with the median monthly expenses of homeowners, \$1,732 for those with a mortgage and \$502 for those without, the median bill total would be 2.03% and 7.01% of monthly bills respectively. During the summer months the median bill total is \$81.16, resulting in 4.69% and 16.17% of summer monthly bills respectively. During the winter months the median bill total is \$27.21, resulting in 1.57% and 5.42% of winter monthly bills respectively.

Aggregating all accounts in my dataset gives a total combined usage of 388,652,537 gallons in the average month resulting in 4,663,830,443 gallons in an average fiscal year. During the summer season (May through October), there is an aggregate of 657,285,540 gallons used each month, an increase of 169% from the average usage. Average revenue per month is \$1,392,531 with average summer revenue being \$2,133,311, an increase of 153.2% from the average monthly revenue. In a fiscal year, this would result in average revenue of \$16,849,449 in July 2018 dollars.

The control city borders the treatment city to the south and is similar in demographics. The median home value and median household income indicates that the control city is more affluent and it is smaller in size given the number of households and population estimates. The household estimate for the control city according to the Census.gov QuickFacts is 13,542. Only 4,906 accounts were provided by the control city, for a few factors, one being the filtering methods mentioned earlier, and the other being that less than half of the water in the city is provided by the city itself, and the rest is provided by a private company, WaterPro. As the only purpose of this data is to control against the treatment city, this is sufficient, as no price change was made in the 2017-2020 fiscal years for the control city. WaterPro did increase their rates, and thus data was not solicited from them.

Thanks to latitude and longitude coordinates provided by the city I am able to create maps that show each account I have data for. Some areas are owned by other cities or have county-based water, which is why there are blank spots. Figure 12 below shows that the east side of the city has higher home values and would be considered wealthier. Figure 34 in Appendix 8.3 shows that this area also tends to have the homes with the largest lot sizes.


Figure 12: Map of Home Value



All values over \$1,000,000 were given the value of \$1,000,000 so that outliers wouldn't skew the color gradient.


Figure 13 has a very similar color gradient to Figure 12 and Figure 34, indicating that those with a higher home value and lot size are generally those who use more water. There are several accounts in locations other than the southeast corner of the city, but the large concentration is in the southeast. A map of bill total, Figure 35, is in Appendix 8.3 and follows the same trend as Figure 13, which makes sense since bill total is derived from usage.

Figure 13: Map of Usage



All values over 50 were given the value of 50 so that outliers wouldn't skew the color gradient.


Figure 14: Map of Average Cost (Bill Total Divided by Usage)



All values over \$5 were given the value of \$5 so that outliers wouldn't skew the color gradient.

Figure 14 above shows the average cost per thousand gallons. This is calculated by taking the bill total and dividing it by the usage. While the marginal cost is increasing as consumers move across different blocks, this doesn't mean their average cost is always increasing, which is why it appears that the areas with low home values have a high average cost (supported by Figure 36 in Appendix 8.3). This is due to the fixed cost in the treatment city's water rate structure (Figure 1). For example, in 2020 the fixed cost of a 0.75-inch meter each month is \$19.95. This fixed cost is spread out across the entire usage, so initially, the average cost goes down, but then spikes up each time usage passes into a new rate tier. Due to this effect, the lowest average cost for a 0.75-inch meter is at 40,000 gallons when the average cost is \$3.25, which is only at the end of the second block. Before that point, the average cost is generally decreasing, except for a small bump at 7,000 gallons when the usage passes from block 1 to block 2. After this point average cost decreases, but since the marginal price is \$3.24 until 80,000 gallons and \$3.72 thereafter, the average cost will never get below where it is at 40,000 gallons. This isn't information that a normal consumer would consider, as their total bill is still rising, but it is interesting to examine, as those with very low usage have a large average cost, especially if they use only a few thousand gallons. This is due to the base fee only being spread across a few of units of usage. The highest average cost after 40,000 gallons is at 81,000 gallons, when the consumer moves into Block 4, and the average cost is \$3.97. Figure 15 below shows average cost for the 0.75-inch meter from 0 to 100 gallons. Figure 37 and Figure 38 in Appendix 8.3 show the same for the 1-inch and 1.5-inch meters. As an example of how the fixed cost is spread out, if a consumer were to only use the median winter usage (November through April) of 4,450 gallons, that \$19.95 would be spread to be an additional \$4.48 per thousand gallons. If a person were to instead use the median summer usage (May through October) of 25,150 gallons that same \$19.95 would be spread to be an additional \$0.79 per thousand gallons. This example illustrates why the average cost continually gets lower, subject to the tiered rate structure.

Figure 15: Average Cost as Usage Increases - 0.75-inch meter



5. Estimation and Results

5.1 Usage

To understand the different effects variables have on usage, I first run a regression with level independent variables on level usage. The corresponding table is in Figure 16 below. Notice that all coefficients are statistically significant with the exception of `meter_size1.5`. The R Squared term is 0.485, which suggests there are other factors not in the dataset that also influence usage. A home value squared term was included in the regression to better explain how wealth affects usage, and dummy variables were created for the meter sizes and the months. `Post` is a dummy variable with a value of one to indicate that an observation is after the rate increase (July 2018) and zero if it was before. Here the coefficient shows a decrease on average of 335 gallons when the rate was changed. This is not the best estimate, as controlling for fixed effects shows a decrease of 238 gallons (Figure 19) and the difference in differences regression gives an estimate of 119 gallons (Figure 20).

Figure 16: Regression on Usage

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-23.36	0.35	-66.71	0.00
post	-0.33	0.04	-8.52	0.00
home_value	3.05	0.02	140.46	0.00
home_value2	-0.07	0.00	-64.55	0.00
lot_size	8.38	0.11	78.33	0.00
rain	0.58	0.02	24.40	0.00
snow	0.16	0.01	22.71	0.00
temp	0.36	0.01	45.25	0.00
meter_size1	4.49	0.08	56.90	0.00
meter_size1.5	0.26	0.96	0.27	0.78
month_nameFeb	-1.11	0.08	-13.29	0.00
month_nameMar	-5.30	0.13	-39.69	0.00
month_nameApr	-7.11	0.16	-43.96	0.00
month_nameMay	-2.27	0.22	-10.37	0.00
month_nameJun	7.63	0.29	26.63	0.00
month_nameJul	22.64	0.37	62.02	0.00
month_nameAug	24.76	0.34	73.07	0.00
month_nameSep	21.44	0.26	83.03	0.00
month_nameOct	9.01	0.15	58.33	0.00
month_nameNov	-0.27	0.11	-2.43	0.01
month_nameDec	1.17	0.09	13.46	0.00
covid	8.43	0.08	102.95	0.00

For interpretability, home value is in units of \$100,000.

R Squared	Observations
0.485	1,022,922

This regression can help to explain several effects on the account level. The 1-inch meter size has the highest usage of the three meters, with an average of 4,492 gallons more than the 0.75-inch meter, which serves as the base in the regression. Interestingly, even though the 1.5-inch meter is the biggest and has the largest space between tiers in the rate structure, (Figure 1) it does not have as high an average usage as the 1-inch meter size. Home value has a positive coefficient that has the highest t-value in the table. The home value squared term has a negative coefficient, but the effect is not extremely large and even at the maximum home value of 42.1 (home value is in units of \$100,000), the total effect of home value is 1,152, gallons, which is still a positive number. Evaluating the coefficients of home value at the median value of \$374,600 gives a value of 10,421 gallons. Increasing that value by one standard deviation (\$186,631) results in 14,860 gallons, an increase of 4,439 gallons per month. Lot size evaluated at its median of 0.22 acres gives a value of 1,856 gallons. An increase of one standard deviation of lot size, which is 0.19 acres, results in an increase in usage of 1,592 gallons at a value of 3,448 gallons.

This regression also gives some insight into other factors that influence usage. The coefficients on the month dummy variables follow the same pattern that is seen in Figure 7, where summer months have a higher usage than winter months. The months are contrasted against January so the coefficient shows the difference between a given month and what usage in January would look like, controlling for other effects. As expected, the coefficient on temperature is positive, meaning on average, a hotter month sees more usage than that same month in a cooler year. The standard deviation of temperature is 17.34, an increase of 2,705 gallons in the average month. Interestingly, rain and snow both have positive coefficients, which would indicate more rain or snow leads to additional usage in a month. This deviation from expectation is likely the result of multicollinearity, as rainfall and snowfall are correlated with the months and temperature (Figure 39, Figure 40, Figure 41, and Figure 42). Figure 44 in Appendix 8.3, where only rain is regressed on usage, shows a negative coefficient, but with a very low R squared. This demonstrates that increased rainfall does decrease usage, but the effect is very small in explaining usage. Figure 43 in Appendix 8.3 is an interaction regression between rain and months with other effects controlled for. Rain within a month does not appear to impact usage, as half of the interaction coefficients are positive, and the other half are negative. While

this could be due to the lack of observations, as breaking down the interaction like this takes the number of months from 48 to four, it could also be due to Utah receiving so little rain that it doesn't make much of a difference. Figure 45 is a very similar regression, but with temperature instead of rain. Three of the twelve coefficients are negative, but the other nine coefficients are positive. Some coefficients, such as May through July, show an impact of over 1,000 gallons per extra degree of temperature. The coefficient of June shows a decrease of 2,002 for a degree increase in temperature, which is a huge amount. While this regression does suffer from the same lack of observations as Figure 43, it does make a strong case that temperature has more of an influence on usage than rain does.

Figure 46 in Appendix 8.3 compares different variations of level and logged independent and dependent variables, as well as interactions between post and home value. R2 of Figure 46 is a regression of level usage and level dependent variables with interaction terms between home value and post. In R2, the interaction coefficient for post and home value is -0.26 and the coefficient for post and home value squared is 0.01. This suggests that those with higher home values decreased their usage more than those with lower home values. Figure 47 in Appendix 8.3 has usage regressions broken out by subsets of the data where the buckets are broken out by home value in jumps of \$100,000, and it shows that users across all buckets decreased their usage after the rate change. This is a key insight, as one of the concerns of treatment city officials was the rate increases would disproportionately hurt the poor, but this shows that the rich were also affected when the rate increases occurred, and in many cases decreased their usage more than those with lower home values (Figure 57). This doesn't mean the poor didn't experience a negative impact, as their usage is already lower in general, but it is helpful knowledge for city officials to consider.

Figure 16 explains how several different factors impact usage, but there is a question as to how much of the variation in the data can be attributed to each variable, and which variables have the greatest effects in determining usage. One technique to see the influence of a variable compared to others is to take one or more variables and regress them on the independent variable and then run another regression with the remaining variables on the residuals from the first regression. Doing this gives the same coefficients as running the regression with all of the variables together, but it does show how the different variables' importance in explaining the independent variable. I used this technique to examine how much usage is explained by the month variables over the account-specific data. The first regression is given below in Figure 17, which has twelve month dummy variables on usage. In Figure 18 the other variables, other than post, are then regressed on the residuals from Figure 17.

Figure 17: Months on Usage

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	5.44	0.06	94.02	0.00
month_nameFeb	-0.24	0.08	-2.84	0.00
month_nameMar	-0.71	0.08	-8.38	0.00
month_nameApr	0.39	0.08	4.68	0.00
month_nameMay	8.24	0.08	98.45	0.00
month_nameJun	21.33	0.08	255.89	0.00
month_nameJul	37.80	0.08	451.57	0.00
month_nameAug	38.73	0.08	469.38	0.00
month_nameSep	31.75	0.08	378.13	0.00
month_nameOct	13.54	0.08	161.85	0.00
month_nameNov	1.97	0.08	23.42	0.00
month_nameDec	-0.17	0.08	-2.02	0.04

R Squared	Observations
0.423	1,022,922

The R Squared of Figure 17 shows that that 42.3% of the variation in usage during each billing period can be explained alone by the month, and the other factors only explain another 9.6% of the variation (Figure 18). Interestingly, running only temperature on usage Figure 48 shows that 34.7% of the usage can be explained by temperature alone. This shows how there is multicollinearity between temperature and the months, which makes sense since they run together (Figure 29). Using this method and running only account-specific variables on usage, such as home value, lot size, and meter size, demonstrates that only 5.5% of usage is explained by the individual characteristics of a given account (Figure 50 and Figure 51). This makes sense, as an account does not vary much

from its typical usage, and more of the variation for the account is due to the season and the temperature. For more information in how the rate change impacted the months see Figure 53 in Appendix 8.3, which breaks out usage regressions by month. The R Squared terms are lower than in Figure 16, which makes sense since in the basic regression there are 48 months to draw upon and for Figure 53 there are only four observations of each month. It shows that during the peak months of June through August, there was a reduction in usage after the rate change. Other months display an increase in usage when the rate was increased, thus the usability of those estimates is questionable.

Figure 18: Other Variables on Residuals of Figure 17

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-14.38	0.13	-110.32	0.0
home_value	3.05	0.02	139.39	0.0
home_value2	-0.07	0.00	-63.96	0.0
lot_size	8.38	0.11	77.84	0.0
rain	-0.16	0.02	-9.03	0.0
snow	0.03	0.00	6.65	0.0
temp	0.01	0.00	6.31	0.0
meter_size1	4.51	0.08	56.80	0.0
meter_size1.5	0.37	0.96	0.38	0.7

For interpretability, home value is in units of \$100,000.

R Squared	Observations
0.096	1,022,922

Figure 19: Regression on Usage with Fixed Effects

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-27.76	0.30	-92.56	0
post	-0.24	0.03	-6.92	0
rain	0.56	0.02	27.00	0
snow	0.16	0.01	25.92	0
temp	0.36	0.01	52.45	0
month_nameFeb	-1.17	0.07	-16.10	0
month_nameMar	-5.39	0.12	-46.23	0
month_nameApr	-7.23	0.14	-51.16	0
month_nameMay	-2.49	0.19	-13.01	0
month_nameJun	7.35	0.25	29.37	0
month_nameJul	22.28	0.32	69.88	0
month_nameAug	24.46	0.30	82.64	0
month_nameSep	21.17	0.23	93.86	0
month_nameOct	8.82	0.13	65.39	0
month_nameNov	-0.38	0.10	-4.02	0
month_nameDec	1.11	0.08	14.57	0
covid	8.58	0.07	119.90	0

R Squared	Observations
0.522	1,022,922

The purpose of Figure 16 is to give some explanation to the different variables that affect usage, however, as the main question of this study is how the rate increase affected usage, I ran a regression using account-level fixed effects, shown below in Figure 19. Fixed effects control for anything specific to the account, including things outside the

data provided. For example, while two different households may have the same lot size and same home value, they could have different incomes, landscaping, or sprinkling systems, which is not accounted for in the dataset. Fixed effects address this by creating a dummy variable for each account within the regression. As adding an additional 21,604 regressors would be a large burden on computational power, I use another method where the average usage of the account is taken and subtracted from the actual usage of that month. This serves the same purpose as dummy variables in creating fixed effects. Due to adding fixed effects to the regression, no coefficients are created for home value, meter size, or lot size, but those effects were addressed in Figure 16 and this regression only focuses on the post variable, which here shows a decrease of 238 gallons. This estimate is likely more precise than the 335-gallon estimate found in Figure 16, supported by the higher R Squared value of 0.522. Note that fixed effects were not added for each month of the four-year span. This is because the goal of this paper is not to determine how the rate change only impacted usage in July 2018, but to generalize how rate changes made at any time affect usage.

Appendix 8.3 contains several instances where I use this method to look at how the rate change impacted usage within each block. All of the blocks seem to be impacted, with usage in block one decreasing by 202 gallons (Figure 54), usage in block two decreasing by 224 gallons (Figure 55), usage in block three decreasing by 237 gallons (Figure 56), and usage in block four decreasing by 1,440 gallons (Figure 57). These regressions only examine the usage within each block, so if a 0.75-inch meter were to use 8,500 gallons before the rate change, only the 2,500 gallons that are in block two would be examined. These estimates show the rate changes affected all tiers and did so about equally for each block. This is interesting, as it shows that the rate changes affected all users, not just those in the lowest blocks.

Figure 20: Difference in Differences Usage Regression with Fixed Effects

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-26.61	0.28	-94.33	0.00
treatment_post	-0.12	0.08	-1.56	0.12
post	-0.11	0.07	-1.46	0.14
treatment	0.03	0.05	0.52	0.60
rain	0.49	0.02	25.42	0.00
snow	0.15	0.01	26.82	0.00
temp	0.35	0.01	53.57	0.00
month_nameFeb	-1.13	0.07	-16.65	0.00
month_nameMar	-4.99	0.11	-46.04	0.00
month_nameApr	-6.73	0.13	-51.26	0.00
month_nameMay	-1.71	0.18	-9.64	0.00
month_nameJun	8.14	0.23	35.02	0.00
month_nameJul	22.24	0.30	75.13	0.00
month_nameAug	23.97	0.27	87.22	0.00
month_nameSep	20.48	0.21	97.71	0.00
month_nameOct	7.91	0.13	63.11	0.00
month_nameNov	-0.45	0.09	-5.03	0.00
month_nameDec	1.06	0.07	15.05	0.00
covid	8.12	0.07	122.33	0.00

R Squared	Observations
0.507	1,192,463

To control for factors that would have impacted the treatment city beyond account-specific fixed effects, I further added to my analysis by including accounts from the control city. As stated previously, the control city did not increase their rates during the four years being studied, and thus provides for a control for the treatment effect. Anything that impacted the usage of both cities equally is thus accounted for. Results are given in Figure 20 above. Here the coefficient of post measures how the usage in both cities changed after July 2018, and it is apparent that usage in both cities dropped by 106 gallons after the rate change. The interaction of post and treatment (treatment_post) shows usage dropped by a further 119 gallons for the treatment city alone. One thing to note is that the treatment_post coefficient in this regression are not significant, with a t-value of only -1.56. This lack of

significance is surprising, since the t-value of Figure 16 was -8.52 and the t-value of Figure 19 was -6.92. Figure 58 in Appendix 8.3 uses the median usage from the control city in the regression, rather than difference in differences, and this shows the total post effect being the same as the combined effect of post and treatment_post in the regression below. Figure 58 has a t-value of -6.42, which is much more than the combined t-values of post and treatment_post, which alludes that interacting the terms causes a large loss in power. In the bill total difference in differences regressions with fixed effects (Figure 23 and Figure 24) the t-values are both significant at 13.98 and 13.13 respectfully. Since bill total is calculated based on usage, it is interesting the bill total coefficients would be significant at such a high level when usage using the same method would not be.

5.2 Bill Total

This second result section aims to answer the question of what happened to city revenue after the rate change. I have already established that usage would decrease with the rate change, but here I answer whether the bill totals also dropped, increased, or stayed the same.

Figure 21: Regression on Bill Total

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-58.08	1.07	-54.41	0.00
post	3.56	0.12	29.68	0.00
home_value	8.73	0.07	131.82	0.00
home_value2	-0.20	0.00	-57.96	0.00
lot_size	27.39	0.33	84.02	0.00
rain	1.56	0.07	21.58	0.00
snow	0.45	0.02	21.71	0.00
temp	1.10	0.02	45.09	0.00
meter_size1	17.42	0.24	72.36	0.00
meter_size1.5	6.04	2.92	2.07	0.04
month_nameFeb	-3.32	0.25	-13.04	0.00
month_nameMar	-15.43	0.41	-37.92	0.00
month_nameApr	-21.23	0.49	-43.06	0.00
month_nameMay	-9.74	0.67	-14.61	0.00
month_nameJun	15.86	0.87	18.15	0.00
month_nameJul	59.08	1.11	53.09	0.00
month_nameAug	65.64	1.03	63.53	0.00
month_nameSep	56.22	0.79	71.40	0.00
month_nameOct	21.80	0.47	46.31	0.00
month_nameNov	-2.08	0.33	-6.25	0.00
month_nameDec	3.51	0.27	13.19	0.00
covid	22.83	0.25	91.42	0.00

For interpretability, home value is in units of \$100,000.

R Squared	Observations
0.449	1,022,922

The post variable in Figure 21 above shows an increase of \$3.56 for the average bill in the average month. This regression does not account for inflation, which is correlated with time, so a second regression was run where the bill total of each month was changed to the value of the dollar in July 2018 (Figure 22 below). Controlling for inflation shows an increase in real dollars of \$0.67 for the average account in the average month. One thing to note is that the inflation was not controlled for just the post coefficient, but for all 1,022,922 observations of bill total. In both cases, there was an increase in city revenue, although not by much in terms of real purchasing power. The regression on nominal dollars suggests an increase in city revenue of \$76,863 each month, resulting in an increase in yearly revenue of \$922,361. From the regression using July 2018 dollars, there appears to only be an increase in revenue of \$14,558 in the average month. In a year, this would increase city revenue by \$174,700 in July 2018

dollars. The reason that these two figures can be so far off when inflation over the four year span was small (Figure 30) is due to the way that the regression works, where every individual observation is adjusted for inflation rather than only the post variable.

Figure 22: Regression on Bill Total in July 2018 Dollars

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-44.14	1.07	-41.33	0.00
post	0.67	0.12	5.55	0.00
home_value	5.48	0.03	159.58	0.00
lot_size	28.54	0.33	86.52	0.00
rain	1.57	0.07	21.44	0.00
snow	0.42	0.02	19.86	0.00
temp	1.02	0.02	41.41	0.00
meter_size1	18.44	0.24	75.77	0.00
meter_size1.5	-4.39	2.95	-1.49	0.14
month_nameFeb	-3.35	0.26	-12.98	0.00
month_nameMar	-14.89	0.41	-36.14	0.00
month_nameApr	-20.50	0.50	-41.05	0.00
month_nameMay	-8.59	0.68	-12.71	0.00
month_nameJun	18.10	0.89	20.45	0.00
month_nameJul	64.41	1.13	57.14	0.00
month_nameAug	70.12	1.05	67.01	0.00
month_nameSep	59.28	0.80	74.34	0.00
month_nameOct	23.05	0.48	48.33	0.00
month_nameNov	-1.54	0.34	-4.56	0.00
month_nameDec	3.23	0.27	11.99	0.00
covid	22.90	0.25	90.56	0.00

For interpretability, home value is in units of \$100,000.

R Squared	Observations
0.448	1,022,922

Figure 59 (nominal dollars) and Figure 60 (July 2018 dollars) in Appendix 8.3 compare different variations of level and logged independent and dependent variables, as well as interactions between post and home value. The R2 regression has an interaction term between home value and post. In Figure 59, the interaction coefficient for post and home value is -0.13 and the coefficient for post and home value squared is 0.02. In Figure 60, the interaction coefficient for post and home value is -0.66 and the coefficient for post and home value squared is 0.03. This suggests those with higher home values decreased their bills along with their usage (Figure 46) more than those with lower home values. This is a key insight, as one of the concerns of treatment city officials was that the rate increases disproportionately hurt the poor, but here it is seen that the rich were more affected than the poor when the rate increases occurred. This doesn't necessarily mean the poor were not impacted, as their bill totals are already lower in general, but it is useful knowledge for city officials to consider. Figure 61 in Appendix 8.3 supports this, as it shows that those with higher home values on average had higher bills. Figure 62 is a similar table, but with bill total adjusted for inflation. Figure 63 and Figure 64 in Appendix 8.3 show bill total regressions broken out by month. Nothing very notable was found, but they are included in this paper for reference.

Figure 65 and Figure 66 in Appendix 8.3 are bill total regressions with fixed effects. The coefficient of nominal bill total is 3.81 and the coefficient of bill total in July 2018 dollars is 0.94. These estimates imply higher revenue gains than the ones projected in Figure 21 and Figure 22, which make sense, since the fixed effects regression on usage (Figure 19) found a smaller decrease in usage than the level regression (Figure 16).

Figure 23: Difference in Differences Bill Total Regression with Fixed Effects

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-79.70	0.91	-87.77	0.00
treatment_post	3.43	0.25	13.98	0.00
post	0.36	0.23	1.55	0.12
treatment	-1.82	0.17	-10.54	0.00
rain	1.34	0.06	21.72	0.00
snow	0.45	0.02	24.95	0.00
temp	1.08	0.02	51.90	0.00
month_nameFeb	-3.63	0.22	-16.65	0.00
month_nameMar	-15.12	0.35	-43.35	0.00
month_nameApr	-20.80	0.42	-49.20	0.00
month_nameMay	-7.79	0.57	-13.63	0.00
month_nameJun	19.58	0.75	26.15	0.00
month_nameJul	61.77	0.95	64.81	0.00
month_nameAug	66.81	0.88	75.51	0.00
month_nameSep	56.23	0.67	83.33	0.00
month_nameOct	19.23	0.40	47.63	0.00
month_nameNov	-2.78	0.29	-9.74	0.00
month_nameDec	3.10	0.23	13.59	0.00
covid	22.87	0.21	107.03	0.00

R Squared	Observations
0.457	1,192,463

Figure 23 (above) and Figure 24 (below) are bill total difference in differences regressions with fixed effects. In both regressions, the most meaningful coefficient is treatment_post, which is the effect that the rate change had on the treatment city due to the rate increase. Post in this case measures how much the bill total changed for both cities after July 2018, and treatment_post shows the effect that only the treatment city experienced. Notice that unlike in Figure 20, the difference in differences regression on usage, the coefficient of treatment_post for both bill total regressions is very significant. Interestingly the nominal estimate and the real estimate of the rate change are very similar at \$3.43 and \$3.26 respectfully. The coefficient for the inflation-modified regression may seem surprising, as this estimate is much higher than those seen in Figure 22 and Figure 66. An explanation for this can be given by the coefficient of post, which shows a decrease for both cities after the rate change of \$2.34. This decrease was an effect seen in both cities and thus cannot be attributed to the rate change. In the two previous regressions this systematic decrease was hidden and is only shown through the use of the control city data. This is evidence of the value of the difference in differences approach, as it shows the true impact was somewhat hidden in the previous regressions. From these estimates the treatment city can expect to increase their nominal yearly revenue by \$889,718 and have a real increase in yearly revenue of \$845,042.

Figure 24: Difference in Differences Bill Total Regression with Fixed Effects (in July 2018 Dollars)

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-75.54	0.92	-82.24	0
treatment_post	3.26	0.25	13.13	0
post	-2.34	0.24	-9.89	0
treatment	-1.72	0.17	-9.84	0
rain	1.33	0.06	21.32	0
snow	0.41	0.02	22.76	0
temp	1.00	0.02	47.57	0
month_nameFeb	-3.58	0.22	-16.21	0
month_nameMar	-14.44	0.35	-40.94	0
month_nameApr	-19.96	0.43	-46.68	0
month_nameMay	-6.51	0.58	-11.26	0
month_nameJun	21.94	0.76	28.98	0
month_nameJul	67.28	0.96	69.78	0
month_nameAug	71.37	0.90	79.74	0
month_nameSep	59.43	0.68	87.07	0
month_nameOct	20.52	0.41	50.25	0
month_nameNov	-2.17	0.29	-7.51	0
month_nameDec	2.89	0.23	12.54	0
covid	22.96	0.22	106.22	0

R Squared	Observations
0.458	1,192,463

5.3 Elasticity

As mentioned in the literature review of this paper, computing an elasticity is difficult when working with an increasing block rate structure due to the different marginal prices at different levels of usage. The equation for elasticity is the percentage change in quantity over the percentage change in price. If price changes by a percentage and quantity demanded of that good changes by a larger percentage, then the elasticity measure is greater in magnitude than one and the good is considered elastic. If the price changes by a percentage and the quantity changes by a smaller percentage, then the elasticity measure is smaller in magnitude than one and the good is considered inelastic. A demand elasticity, in theory, is always negative, because when faced with a price increase, the rational consumer uses less of the product. This is where the elasticity breaks down in the increasing block structure. Normally in econometrics, the method for deriving an elasticity from data is to put price in log terms and quantity in log terms, and then the coefficient of the price variable is the elasticity. This fails because at higher prices, usage is higher in the data, and thus when running a regression using logged price, the result is that usage increases, which is why in Figure 67 in Appendix 8.3, the coefficient is 3.24, which is a positive number.

Due to this inconsistency, some researchers have used average price when attempting to determine elasticities (Foster 1979) (Ito 2012). Unfortunately, average price tends to produce excessively large estimates that do not accurately reflect the true elasticity of water demand (Billings 1980). Figure 25 below is a regression using average price instead of marginal price. Here the coefficient of logged average price is -1.47. As expected, it is negative, but excessively large with a magnitude greater than 1, suggesting water is an elastic good, when in reality, the expectation would be the opposite, as most studies have found water to be an inelastic good (Espey 1997).

Figure 25: Average Price Elasticity Regression

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	3.69	0.01	506.63	0.00
log(av_price)	-1.47	0.00	-1904.90	0.00
home_value	0.07	0.00	149.84	0.00
home_value2	0.00	0.00	-85.73	0.00
lot_size	0.12	0.00	49.45	0.00
rain	0.01	0.00	20.77	0.00
snow	0.00	0.00	10.19	0.00
temp	0.01	0.00	40.71	0.00
meter_size1	0.24	0.00	134.56	0.00
meter_size1.5	0.36	0.02	16.24	0.00
month_nameFeb	-0.03	0.00	-14.53	0.00
month_nameMar	-0.10	0.00	-35.59	0.00
month_nameApr	-0.18	0.00	-51.39	0.00
month_nameMay	-0.01	0.00	-1.10	0.27
month_nameJun	0.33	0.01	54.83	0.00
month_nameJul	0.75	0.01	101.53	0.00
month_nameAug	0.80	0.01	115.41	0.00
month_nameSep	0.71	0.01	132.78	0.00
month_nameOct	0.31	0.00	93.50	0.00
month_nameNov	0.00	0.00	2.00	0.05
month_nameDec	0.02	0.00	10.74	0.00
covid	0.28	0.00	155.75	0.00

For interpretability, home value is in units of \$100,000.

Values with zero usage were filtered away for the log transformation.

R Squared	Observations
0.902	1,018,037

Since using regression techniques does not give useful elasticity estimates, I computed the measure through the base formula. For the expected percentage change in quantity, I took the post coefficients from Figure 16 (level-level usage regression), Figure 19 (fixed effects usage regression), and Figure 20 (difference in differences usage regression with fixed effects) and took that as the change in quantity. Dividing those measures by the median usage from before the rate change gives the percentage change in quantity. From looking at the rate changes, and consulting with the treatment city, they increased the marginal cost from each block by 5%, which makes computing the elasticity relatively easy, as I can use it as the percentage change in price (Figure 52). By computing with each of the post estimates, I get the following elasticities: -0.84 from Figure 16, -0.59 from Figure 19, and -0.30 from Figure 20. All three of the estimates are negative and all three are smaller in magnitude than one, which is expected. Comparing the three elasticities above for an account with median usage (7,860 gallons) would result in decreases of usage of 66 (-0.84 elasticity), 47 (-0.59 elasticity), and 23 gallons (-0.30 elasticity). For an account that uses the third quartile of usage, 25,820 gallons, this would result in decreases of 216 (-0.84 elasticity), 154 (-0.59 elasticity), and 77 gallons (-0.30 elasticity).

From the fixed effects regressions within each block in Appendix 8.3 I calculate the following for the decreases in usage for each block: 202 gallons in block one (Figure 54), 224 gallons in block two (Figure 55), 237 gallons in block three (Figure 56), and 1,440 gallons in block four (Figure 57). Using these estimates with the exact price change for each block, which varies slightly from 5% due to the number being rounded to the nearest cent, I calculate the following elasticities: -1.07 for block one, -0.30 for block two, -0.09 for block three, and -0.29 for block four. For the median usage (pre-July 2018) in block one of 3,590 gallons at the elasticity of -1.07 would result in a decrease of 38 gallons from a 1% increase in price. For the median usage (pre-July 2018) in block two of 14,880 gallons at the elasticity of -0.30 would result in a decrease of 45 gallons from a 1% increase in price. For the median usage (pre-July 2018) in block three of 52,540 gallons at the elasticity of -0.09 would result in a decrease of 49 gallons

from a 1% increase in price. For the median usage (pre-July 2018) in block four of 99,100 gallons at the elasticity of -0.29 would result in a decrease of 283 gallons from a 1% increase in price. This shows that for those in the lowest blocks, usage only dropped very slightly with the rate increase. The biggest savings of usage came from block four, where people use the greatest amount of water.

Another elasticity to consider is income elasticity, which is a measure showing how increasing income affects water consumption. The dataset provided by the city doesn't have household income, which is why I use home value as a proxy. Figure 26 below uses home value to show an elasticity of wealth, or how much an increase in wealth has on usage. It is relatively inelastic, but it does show more wealth on average results in more usage. An increase of 1% in wealth would be expected to result in a 0.61% increase in usage. To give perspective, take a home at the median value (\$374,600) that uses median usage (7,860 gallons per month). If this household were to increase its home value by 1% (\$3,746), then in expectation, its usage would increase by 0.61%, or 48 gallons. For a home at the third quartile of home value, \$474,000, that has typical usage in the third quartile, 25,820 gallons, a 1% increase in home value would be \$4,740 and this would be expected to result in a 156-gallon increase in usage.

Figure 26: Wealth Elasticity Regression

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-5.02	0.02	-252.30	0
log(home_value)	0.61	0.00	386.67	0
lot_size	0.18	0.00	63.68	0
meter_size1	0.11	0.00	53.34	0
meter_size1.5	-0.51	0.03	-20.44	0

For interpretability, home value is in units of \$100,000.

R Squared	Observations
0.226	1,022,922

6. Discussion

From my analysis I have shown that controlling for many other effects, the city should have expected to increase revenue by 1.5% to 4.2%, rather than seeing the actual decrease of 1.1% between fiscal years 2018 and 2019. According to city officials, their full revenue was \$21,228,160 for fiscal year 2018 and \$20,996,029 for fiscal year 2019, a decrease of \$232,131. Because revenue went down the year following the rate increase, city officials worried the rate change had decreased usage at such a level that revenue had also decreased. If this were the case, then the argument could be made that the rate changes were too drastic, and potentially even harmful if it caused the poor to be unable to afford the cost of necessary water. My analysis shows that in a normal year, the treatment city would have expected to see an increase in revenue rather than a decrease, and that the increases in bill total and decreases of usage would have likely been bearable to the poor.

The reason revenue went down so much for the treatment city is likely from the decrease in temperature, as fiscal year 2019 had an average temperature of 53.8 degrees compared to 56.9 degrees in 2018. In Figure 48 I have shown approximately 34.7% of the variation in usage can be described by temperature. Fiscal year 2019 also had more rainfall at 19.9 inches compared to 12.4 inches in 2018, which would additionally decrease usage (Figure 44). Looking at Figure 20, the coefficient of temperature is 0.35, meaning a decrease of 3.1 degrees would result in a decrease of usage of 1,072 gallons, everything else constant, which caused the city to see a decrease in usage much bigger than expected. This decrease in temperature would result in a loss of \$3.10 per account each month in July 2018 dollars, or \$803,727 in a full year (Figure 24). The difference of one degree in temperature can cause usage in summer months to change by over 1,000 gallons (Figure 45), which further shows how this decrease in temperature would have affected overall usage. Since the highest revenue that the city receives is during the summer months, as more gallons are used and at higher rates, a decrease in 3.1 degrees would have decreased usage substantially, and likely at such a rate that accounts for the large decrease in city revenue.

Interestingly, the revenue in my data, which only focuses on residential accounts, has total revenue in fiscal year 2018 at \$15,899,922 and \$16,444,011 in 2019, which is an increase of \$544,088, unlike the overall decrease the city

observed. As my data only includes residential accounts, I would assume the remaining discrepancy is explained by the accounts I do not have data for, mostly those with commercial-grade meters or apartment complexes. It makes sense that there would be differences in behavior between residential accounts and commercial accounts due to differences in water needs. One of these needs is that commercial accounts require more water than residential accounts and thus are given larger-sized meters. Given that increasing the marginal price will be an increase across every thousand gallons, the more water used, the higher the increase will be. When comparing the elasticities of block 2 to block 4 (see section 5.3), they are almost the same at -0.30 and -0.29 respectively, however due to the amount of water used within those blocks, this caused very different decreases at 224 and 1,440 gallons respectively (Figure 55 and Figure 55). This would indicate that the decrease in temperature would also have caused large accounts to decrease their usage by a much larger amount than residential accounts would have seen. As mentioned previously, the decrease in temperature caused an average decrease across residential accounts of 1,072 gallons (Figure 20), which was a decrease in average usage of 5.88%. While I am unable to make concrete assumptions on the impact on commercial accounts without additional data, it is plausible that decreases in commercial accounts would account for the overall decrease in revenue in the city. The main focus of this study is on residential usage and thus studying the commercial effects is outside the scope of this paper. Hopefully the city will be able to do further analysis at a later date to understand what about the other accounts resulted in the city-wide decrease in revenue.

The reason revenue among residential accounts did not decrease is due to the rate change. Water usage across the city between fiscal years 2018 and 2019 decreased from 4,499,964,320 to 4,377,535,100, a decrease of 122,429,220 gallons. Even though usage decreased, the city brought in more revenue due to the rate change, a total of \$544,088. Without the rate increase and assuming an elasticity of -0.59, this would mean the city would have probably only observed a revenue of \$15,632,285 in fiscal year 2019 among residential accounts, which is \$811,726 less than their actual 2019 revenue.

With a 5% increase in water rates to the average account usage was estimated to fall by 238 gallons each month, and perhaps as little as 119, although the difference in differences regression was just shy of statistical significance. Studying more years would be advisable to get higher significance as well as increased precision. While usage decreased for the average account, the bill total increased on average by \$3.26 in terms of real dollars. With the average bill total being \$65.23 before the rate change, this is an increase of 5.0% in real dollars, provided that the average account has the average bill. The average bill is much higher than the median bill, which is \$34.91, so if the average account instead has the median bill, this would be an increase of 9.3%. This estimate is likely inflated, as regression estimates are derived from the average of the observations, not the median, but assuming variation among accounts, this wouldn't be an impossible increase. Even for a household with this sort of increase, the treatment city's rates are still well below the EPA standard of having the average cost of water being deemed "affordable" at 2.0-2.5% of median household income (U.S. EPA 2002). Having the average bill be 2% of median household income would result in a bill of \$153.06, which is much higher than the current case. If the question of the treatment city is whether another rate increase would cause extreme harm to city residents, then the answer is that for the majority of residents, the rate increase would not be harmful. There may be a small minority of people that would be harshly affected, but the city can institute other programs, such as a water rebate for low-income residents (Olmstead 2009).

While using the control city data did not give a statistically significant estimate on usage in the difference in differences regression, it did serve to show real revenue increased very close to nominal revenue. Comparing the post estimates in the level-level regression (Figure 22) and the fixed effects regression (Figure 66) to the difference in differences regression (Figure 24) there is a marked difference between the first two estimates and the results from the difference in differences regression. Without this additional data, the real revenue increase would have only been thought to yield \$174,700 to \$243,192 instead of a real effect of \$845,042. This demonstrates that comparing across cities was able to give more precise estimates and uncover effects that would have been hidden otherwise.

The treatment city was contemplating another rate increase at the beginning of fiscal year 2021, but with the previous worry that rate increases unduly hurt residents and with concerns over the COVID-19 pandemic, this action was postponed. Under the assumption that the city would like to implement a rate increase at the beginning of fiscal year 2022, I have computed estimates of how different increases would impact residents as well as the city. In an effort to give ranges to the treatment city, I will use the three estimates of elasticity that I have derived, -0.30 (Figure 16), -0.59 (Figure 19), and -0.84 (Figure 20). Of the three, the most significant estimate, created from the largest amount of information, was -0.59, and thus would be the most accurate estimate and the best prediction my model generates. The other two estimates are included to give a range to the treatment city as they consider possible effects and other factors outside the scope of this paper. To achieve the new bill total estimates, I

multiplied the usage for each month on the account level by the elasticity, separated the usage into the respective blocks, and multiplied each new block by the new rate. This created a precise estimate of the bill total and revenue increases, as all 1,022,922 were modified according to the elasticity estimate at the new rate. The base rate did not change between fiscal years 2018 and 2021, so I will not change it in my estimate. The results are given below in Figure 27 for nominal dollars and in Figure 28 for estimates converted into July 2018 dollars.

Figure 27: Elasticity and Rate Change Comparison

Elasticity	Median Usage	Average Usage	Median Bill Total	Average Bill Total	City Usage	City Revenue	Revenue Increase	Percent Increase
Current Rates								
	7,860	18,237	\$ 35.18	\$ 65.34	4,663,830,443	\$ 16,710,375		
3% Increase								
-0.84	7,662	17,779	\$ 36.42	\$ 66.93	4,546,577,180	\$ 17,115,388	\$ 405,013	2.42 %
-0.59	7,720	17,912	\$ 36.56	\$ 67.32	4,580,590,388	\$ 17,215,732	\$ 505,358	3.02 %
-0.30	7,790	18,074	\$ 36.76	\$ 67.80	4,622,185,177	\$ 17,338,603	\$ 628,228	3.76 %
5% Increase								
-0.84	7,531	17,473	\$ 36.40	\$ 66.98	4,468,408,338	\$ 17,128,448	\$ 418,073	2.50 %
-0.59	7,626	17,695	\$ 36.66	\$ 67.64	4,525,097,019	\$ 17,298,634	\$ 588,260	3.52 %
-0.30	7,743	17,966	\$ 36.96	\$ 68.46	4,594,421,666	\$ 17,507,204	\$ 796,829	4.77 %
7% Increase								
-0.84	7,399	17,167	\$ 36.33	\$ 66.91	4,390,239,496	\$ 17,110,533	\$ 400,159	2.39 %
-0.59	7,533	17,478	\$ 36.71	\$ 67.86	4,469,603,649	\$ 17,352,701	\$ 642,326	3.84 %
-0.30	7,696	17,857	\$ 37.15	\$ 69.02	4,566,658,156	\$ 17,649,763	\$ 939,389	5.62 %
10% Increase								
-0.84	7,201	16,709	\$ 36.19	\$ 66.82	4,272,986,233	\$ 17,089,011	\$ 378,637	2.27 %
-0.59	7,392	17,152	\$ 36.73	\$ 68.21	4,386,363,595	\$ 17,443,583	\$ 733,208	4.39 %
-0.30	7,626	17,694	\$ 37.40	\$ 69.91	4,525,012,890	\$ 17,879,101	\$ 1,168,726	6.99 %

See below table for the price increases used at each percentage.

Rate Increase	Block 1	Block 2	Block 3	Block 4
Current Rate	\$1.80	\$2.75	\$3.24	\$3.72
3%	\$1.85	\$2.83	\$3.34	\$3.83
5%	\$1.89	\$2.89	\$3.40	\$3.91
7%	\$1.93	\$2.94	\$3.47	\$3.98
10%	\$1.98	\$3.03	\$3.56	\$4.09

From these tables, treatment city officials can adapt their strategy based upon the bill total they believe residents can bear, the amount of gallons the city can purchase to distribute, or the target revenue they would like to achieve. These estimates are explanatory of the data I presently have and may or may not approximate the effects that would be seen in commercial meters in these cases. Additionally, this exercise assumes a universal rate increase, whereas city officials may desire to increase different blocks at different percentages. In that case, they would want to use the elasticities that I computed for each block in section 5.3 to determine how usage would change in each block at the different rates. Since for nominal dollars some of the increase in revenue goes towards covering the rise in inflation, Figure 27 overestimates the real increase in revenue that the city would see. Figure 28 is much better at explaining the real effect of the rate increases. This is because by accounting for inflation, the percentage increases show how much more the city has the ability to accomplish with their water revenue, as the income going towards covering inflation has been controlled for and the only remaining effect is the increase in real dollars. Something interesting to note from these tables is that for the elasticities of -0.59, and -0.30, overall revenue increased with each increase of the rates, however, for the -0.84 elasticity, the increase in revenue for 7% and 10% was less than for 3% and 5%. Revenue is still expected to increase, but to do so less than at the 3% and 5% levels. Understanding and being aware of this range will aid the treatment city as they consider the possible effects of different rate changes, especially given that other factors, especially the weather, will also affect usage. Even at the highest estimates, the average bill total is still within EPA standards (U.S. EPA 2002), but as residents are more resistant to water rate

increases than other, less visible tax increases, city officials likely wouldn't want to increase rates in a large amount at one time (Timmins 2002). Smaller rate increases that are more frequent would be more likely to yield increased revenue over the long term, as a large increase would create a shock to the bill total all at one time and would be likely to have a greater impact on decreasing usage than small increases spread out over time. This also gives the city less risk, as the variation in the range of estimates grows with larger rate increases.

Figure 28: Elasticity and Rate Change Comparison - July 2018 Dollars

Elasticity	Median Usage	Average Usage	Median Bill Total	Average Bill Total	City Usage	City Revenue	Revenue Increase	Percent Increase
Current Rates								
	7,860	18,237	\$ 35.32	\$ 65.89	4,663,830,443	\$ 16,849,449		
3% Increase								
-0.84	7,662	17,779	\$ 36.22	\$ 66.37	4,546,577,180	\$ 16,971,849	\$ 122,400	0.73 %
-0.59	7,720	17,912	\$ 36.36	\$ 66.75	4,580,590,388	\$ 17,071,274	\$ 221,825	1.32 %
-0.30	7,790	18,074	\$ 36.55	\$ 67.23	4,622,185,177	\$ 17,193,017	\$ 343,568	2.04 %
5% Increase								
-0.84	7,531	17,473	\$ 36.19	\$ 66.42	4,468,408,338	\$ 16,984,809	\$ 135,360	0.80 %
-0.59	7,626	17,695	\$ 36.45	\$ 67.08	4,525,097,019	\$ 17,153,435	\$ 303,986	1.80 %
-0.30	7,743	17,966	\$ 36.76	\$ 67.88	4,594,421,666	\$ 17,360,093	\$ 510,644	3.03 %
7% Increase								
-0.84	7,399	17,167	\$ 36.14	\$ 66.35	4,390,239,496	\$ 16,967,066	\$ 117,617	0.70 %
-0.59	7,533	17,478	\$ 36.50	\$ 67.29	4,469,603,649	\$ 17,207,013	\$ 357,565	2.12 %
-0.30	7,696	17,857	\$ 37.15	\$ 69.02	4,566,658,156	\$ 17,501,352	\$ 651,903	3.87 %
10% Increase								
-0.84	7,201	16,709	\$ 36.00	\$ 66.26	4,272,986,233	\$ 16,945,743	\$ 96,294	0.57 %
-0.59	7,392	17,152	\$ 36.53	\$ 67.64	4,386,363,595	\$ 17,297,064	\$ 447,615	2.66 %
-0.30	7,626	17,694	\$ 37.40	\$ 69.91	4,525,012,890	\$ 17,728,589	\$ 879,140	5.22 %

All dollar amounts were converted to July 2018 dollars.

7. Conclusion

By performing econometric analysis on the water habits of 26,510 households in Salt Lake County, Utah I have been able to create precise measures of the elasticity of water demand. Due to the vast size of my data, as well as the ability to examine how usage contributes to the individual blocks, these estimates add to previous research by being more exact than elasticities measures derived through average price, Rate Structure Premium, and other methods that have been proposed. I have used these estimates to explain a scenario in which the rates were not increased, which should help treatment city officials better understand the observed revenue decrease in fiscal year 2019. Through manipulating individual observations of data, I simulated scenarios in which different rate increases were applied in the future to aid city officials as they plan for future increases. This will be invaluable as they balance their conflicting duties of increasing city revenue, encouraging conservation, and keeping water rates affordable to their residents. Understanding the true impact that rates have on residential water consumption habits will allow government officials to better utilize market techniques over command and control methods in driving water behavior, which will lead to a decrease in overall bills and taxes to the population (Olmstead 2009), less waste to society, increased revenue to governmental entities, and more effective conservation.

8. Appendix

8.1 Acknowledgements

Special thanks goes to my wife, Sadie, who supported me through the entire project gave me the time that I needed to complete it. Whenever we have taken long car rides she has always taken the wheel so I could work on my laptop. Thanks goes to Sandy City and Draper City for providing their water rate data. Without the volume of the data, this study would not have been possible and the measurements would not have been as precise. Special thanks goes to Brett Neumann, who helped me start on this project and supported me the entire way. I also want to thank Tom Ward, the Sandy Public Utilities Director for his support, and Zlatko Grebenar for his help with the location data. Thanks to those who proofread the paper, including Audrey Zinninger and Diana, David, Brandee, and Sadie Nellis. Thanks to Michael Blum who helped write the web automation script that allowed me to retrieve the home values from the Salt Lake County Assessors Office website. I estimated that it would have taken over 120 hours to manually retrieve it without his help. Thanks goes to the University of Utah for sponsoring this research and allowing me an forum to present it. Special thanks to my advisor, Scott Schaefer, who took a simple question that I had and encouraged me to work on it for over eight months, 300 hours, and 8,000 lines of code. Without him I would not have done this project, and without his advice and knowledge throughout it would not be to the point that it is today. I have learned that there is so much more to learn and so much deeper to dive in understanding an issue. Econometrics is a powerful tool and I have learned that it can reveal the answer to questions that could not be found in any other way.

Due to size restraints in the University Research Journal the full article including the rest of the Appendix can be found at the links below

https://s3.amazonaws.com/files.formstack.com/uploads/2733688/55776342/714793556/55776342_nate_nellis_urj_submission.pdf

<https://uofu.box.com/s/hcnyhqk2g7ha6r3uoyguu9omnwc72nfa>