

DISSERTATION

THREE ESSAYS ON THE ECONOMICS OF ENERGY, WATER AND POLLUTION

Submitted by

Zarif I. Rasul

Department of Agricultural and Resource Economics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2024

Doctoral Committee:

Advisor: Jesse Burkhardt

Jordan Suter
Dale Manning
Ellison Carter

Copyright by Zarif I. Rasul 2024

All Rights Reserved

ABSTRACT

THREE ESSAYS ON THE ECONOMICS OF ENERGY, WATER AND POLLUTION

This dissertation covers three policy-oriented topics in environmental economics. The first chapter explores the short-term impacts of US liquefied natural gas (LNG) exports on the domestic natural gas and electricity markets. Using a structural vector autoregression framework, we find that unexpected LNG exports corresponded to a 34% price increase in domestic natural gas spot prices in 2022. We illustrate the impact of this price increase on the electricity market by constructing counterfactual dispatch curves for Texas. Had gas prices not increased by 34% in 2022, our results indicate that NO_x, SO₂, and CO₂ emissions from electricity generation in Texas would have been lower by 11.2%, 47%, and 10.7% respectively. The second chapter explores whether utilities can use pro-social appeals to influence residential water consumption by evaluating a program implemented by Denver Water. The program aimed to influence both the timing and volume of residential consumption to eliminate peaking by asking different customer groups to water their lawns on alternate days. Using difference-in-difference models, we find that while the timing of outdoor watering remains unchanged, there is weak evidence that one group of households water on assigned days while another does not, but overall we find no evidence of a significant average treatment effect. These results support previous findings that non-price interventions that include conservation tips or pro-social appeals have little or no impact, as opposed to interventions that contain some element of social comparisons or other incentives. The third chapter investigates the impact of concentrated animal feeding operations (CAFOs) on ammonia and methane concentrations across the United States, leveraging ground-level monitoring data for ammonia and satellite-based observations for methane. For ammonia, we find that wind direction and proximity to CAFOs within 10 km significantly influence concentrations. Additionally, the impact varies by CAFO type. Methane concentrations are similarly elevated in grids containing CAFOs, with

the presence of shale plays and seasonal variations also playing significant roles. Our findings underscore the local environmental impacts of CAFOs, while also highlighting the challenges of isolating these effects from other sources.

ACKNOWLEDGEMENTS

First of all, I would like to extend my gratitude to The Almighty for giving me the ability to undertake this study. I thank my advisor, Dr. Jesse Burkhardt, for his continuous guidance, encouragement, and patience throughout this journey. I am also deeply indebted to the rest of my committee, Dr. Jordan Suter, Dr. Dale Manning, and Dr. Ellison Carter who have all offered their profound knowledge and invested time and effort to elevate the quality of this dissertation. I would like to thank the Department of Agricultural and Resource Economics for the assistantships which not only funded my studies but also provided invaluable learning opportunities. My sincere thanks to our Program Coordinator, Kathy Bruce, for her continuous support in helping us navigate this program. I also thank my colleagues in the program for their counsel and camaraderie. Last but not least, I would like to thank my family, and especially my parents, for their countless sacrifices which have paved the way for all of my achievements.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vi
LIST OF FIGURES	vii
Chapter 1 The Short Term Impacts of US LNG Exports	1
1.1 Introduction	1
1.2 Econometric Framework	5
1.2.1 Structural VAR Model	5
1.2.2 Impulse Response Functions and Forecast Error Variance Decomposition	9
1.3 Data	10
1.3.1 Prices and exports	10
1.3.2 Electricity generation and emissions	11
1.4 SVAR Results	12
1.4.1 Main Findings	12
1.4.2 Robustness checks	15
1.5 Emissions impact	17
1.6 Discussion and Conclusion	22
Chapter 2 Can utilities influence residential water consumption through prosocial ap- peals? Evidence from a pilot study	25
2.1 Introduction	25
2.2 Program Design and Hypotheses	28
2.2.1 Program Design and Context	28
2.2.2 Statistical Identification	31
2.2.3 Data	32
2.3 Methods and Results	35
2.3.1 Estimating Average Treatment Effect	35
2.3.2 Average Treatment Effect among Compliers	47
2.3.3 Empirical Models with Logarithmic Transformation of Consumption . .	50
2.4 Conclusion	55
Chapter 3 The Impact of CAFOs on Ammonia and Methane Concentrations	57
3.1 Introduction	57
3.2 Data and Identification Strategy	62
3.2.1 Data	62
3.2.2 Identification Strategy and Empirical Specifications	68
3.3 Results	70
3.4 Discussion, Conclusions, and Future Research	77
Bibliography	80

LIST OF TABLES

1.1	Summary Statistics	11
1.2	Unit root tests	11
1.3	Estimated SVAR contemporaneous coefficients	13
1.4	Forecast Error Variance Decomposition for Henry Hub (one to ten steps ahead)	15
2.1	Summary Statistics of Daily Residential Water Use (in gallons)	33
2.2	Results for Base Models	37
2.3	Regressions with weather controls	40
2.4	Regression results	42
2.5	Regression results for models with both day and hour	44
2.6	Regression Results	45
2.7	Regression Results for Hourly Effects	46
2.8	Regression results for compliers	50
2.9	Regression results	52
2.10	Regression results	53
2.11	Regression results	54
3.1	Summary Statistics	64
3.2	Summary Statistics (at the grid cell level)	66
3.3	Results for the Base Regression Model	70
3.4	Results for Regressions for CAFO types	72
3.5	Results for regressions with wind interaction	73
3.6	Regression Results for Recent Years (2021-2023)	75
3.7	Regression Results	76

LIST OF FIGURES

1.1	Natural gas and oil prices	3
1.2	Impulse Response Function of Henry Hub to a one-standard deviation shock in LNG exports	14
1.3	Cumulative Impulse Response Function	14
1.4	IRFs under different models	16
1.5	IRFs in different periods	16
1.6	Illustrative Monthly Dispatch Curves	19
1.7	Illustrative Hourly Dispatch Curves	19
1.8	Average hourly emissions	21
1.9	Equilibrium prices	21
2.1	Average Water Use	30
2.2	Sample Mailer	30
2.3	Average hourly use in June and October	34
2.4	Average hourly use on dry treatment days	34
2.5	Average daily use on dry days	43
3.1	CAFO locations and types	63
3.2	Ammonia monitors and nearby CAFOs	65
3.3	Mean Ammonia Concentrations	65
3.4	Mean Methane Mixing Ratio	66
3.5	Average ammonia levels at monitor sites near CAFOs	67
3.6	Average methane levels	67

Chapter 1

The Short Term Impacts of US LNG Exports

1.1 Introduction

Natural gas plays a pivotal role in the US energy landscape, constituting 43% of electricity generation in 2023 (EIA, 2024). Declines in natural gas production around the mid-2000s led to concerns about an impending shortfall in the US, prompting the construction of liquefied natural gas (LNG) import terminals to enable the expansion of natural gas imports beyond pipelines.¹ However, technological advancements in fracking led to a surge in shale gas supply around 2007, leading to the conversion of many import terminals into export terminals and the start of LNG exports in 2016 (EIA, 2023)². The US had the highest LNG export growth in 2021 and became the top exporter in 2023, capturing nearly 21% of the global demand.³ The rapid rise of LNG exports has drawn increased public scrutiny and highlighted the need for more analysis on potential market and environmental impacts.⁴

Both trade theory and previous simulated work such as Baron et al. (2012) support the notion that natural gas prices have increased following LNG exports. How this price increase has affected the electricity sector, the largest consumer of natural gas, is an issue that has not received much attention in the literature. A higher price would make natural gas a less attractive fuel in electricity generation, prompting electricity producers to switch to cheaper and dirtier alternatives such as coal in the short run. This study addresses this gap in the literature by conducting an ex-post analysis of how LNG exports have influenced domestic gas prices, and what this entails for air pollution near power plants. We utilize a Structural Vector Autoregression (SVAR) framework to estimate

¹Forbes, "How the US became the World's Top LNG Exporter, January 5, 2024

²There have been small amounts of LNG exports from Alaska prior to 2016. Unless otherwise stated, LNG exports mentioned in this paper will refer to exports from the lower 48 states.

³Reuters, "US was top LNG exporter in 2023 as hit record levels", January 3, 2024

⁴Reuters, "Biden pauses LNG export approvals after pressure from climate activists", January 26, 2024

the impact of exogenous shocks in LNG exports on spot natural gas prices. We then illustrate how a change in gas prices affect NO_x, SO₂, and CO₂ emissions from electricity generation in Texas.⁵

Estimation results show that a 1% positive shock to LNG exports corresponds with a 7.9% increase in Henry Hub prices. To put this in perspective, consider that actual exports in 2022 surpassed projected exports by 4.3%. If the difference is treated as a shock, then this translates to a 34% increase in natural gas prices. Our estimate is very close to that of 35% reported by Hausman et al. (2015). Had gas prices not increased by 34% in 2022, our results indicate that NO_x, SO₂, and CO₂ emissions from electricity generation in Texas would have been lower by 11.2%, 47%, and 10.7% respectively.

Our paper makes two contributions to the literature. First, we employ econometric methods using observed data, as opposed to utilizing general equilibrium models and conducting ex-ante analysis of future impacts that dominate existing literature. Second, we extend the literature on natural gas prices and power sector emissions by incorporating LNG exports, compared to previous studies that largely focus on fracking. Findings of this study can improve understanding of potential impacts and can inform future export projects as the US establishes itself as the top LNG exporter.

Early studies like Baron et al. (2012) and Hausman and Kellog (2015) projected overall macroeconomic gains from LNG exports, despite a rise in domestic natural gas prices. Hausman and Kellog (2015) estimate a 35% increase in domestic natural gas prices following the completion of all LNG export projects that were under consideration in 2015. These along with other market studies predate LNG exports. Recent work explores the relationship between spot natural gas prices with both domestic and international oil and gas prices. Although international gas markets are not as well connected as oil markets, LNG exports have facilitated greater integration of Asian, European, and North American markets (Halser et al., 2023; Aguiar-Conraria et al., 2022). Asian and European LNG contracts are typically linked to international oil prices, although transactions based on spot prices are gaining traction (Han et al., 2021). Figure 1.1 illustrates how Henry Hub,

⁵Throughout this paper, we use the word 'gas' to refer to natural gas as opposed to gasoline.

the benchmark for natural gas spot prices in the US, move together with US LNG export prices and Brent since the start of LNG exports in 2016. Stock and Zaragoza-Watkins (2024) explore the long-term relationship between domestic natural gas and oil spot price under various natural gas regimes. They estimate that gas prices would be approximately \$2.30 higher during 2010-2015 had LNG exports occurred in this period.

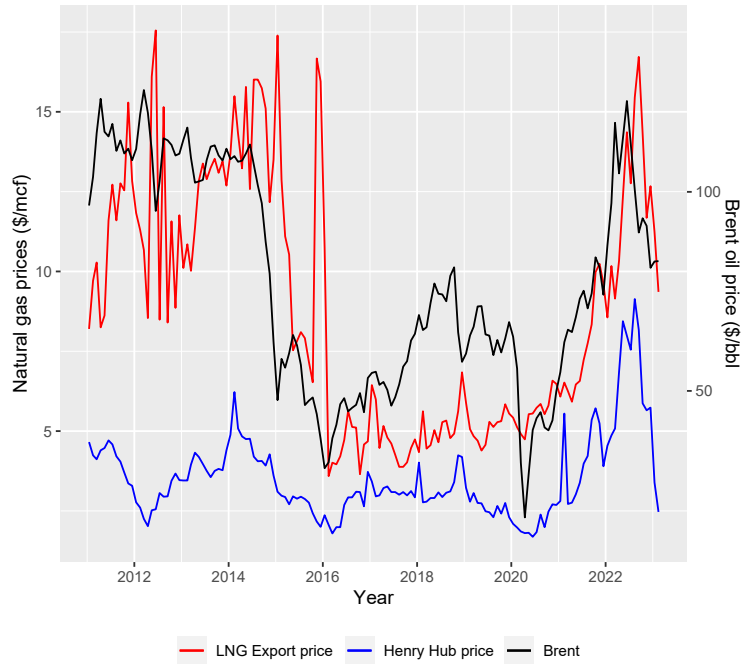


Figure 1.1: Natural gas and oil prices

While there is a vast literature on how coal and gas prices influence fuel choices in electricity generation, these studies primarily focus on the impact of *lower* gas prices resulting from fracking. Johnsen et al. (2019) observed a 28% reduction in coal-fired generation as power plants switched to gas due to lower gas prices. Similar findings were reported by Brehm (2015), Cullen and Mansur (2014), Holladay and LaRiviere (2017), Fell and Kaffeine (2018), and Knittel et al. (2014). The transition from coal to gas-fired generators is associated with lower point-source emissions, leading to improved air quality. Studies examining this transition induced by fracking, such as Johnsen et

al. (2019) and Lu et al. (2012), reported decreases in PM 2.5 and carbon dioxide emissions due to reduced coal generation and increased gas usage. Brehm (2015) additionally found that lower gas prices spurred investments in gas-fired capacity, further decreasing emissions. While Cullen and Mansur (2014) linked emissions responses to carbon pricing, the literature largely overlooks the electricity sector's response to higher gas prices resulting from LNG exports. To date, only Stock and Zaragoza-Watkins (2024) have quantified the impact of LNG exports on future CO₂ emissions from the power sector, focusing on a quantitative equilibrium model without assessing local air quality impacts.

In contrast to Stock and Zaragoza-Watkins (2024) who explore the long-term relationship between energy prices, we incorporate the volume of US LNG exports in an SVAR framework. This is a standard framework in not only macroeconomic studies but also in energy policy studies that do not utilize large scale general equilibrium models, such as Kilian (2009) and Hou and Nguyen (2018).⁶ As explained further in Section 2, we estimate a system of equations involving US natural gas production, storage, Henry Hub prices, LNG exports, and an index for global economic activity. Our aim is to isolate how an exogenous shock, or unexpected innovation, in LNG exports affects Henry Hub prices. After finding the change in Henry Hub prices, we turn to the electricity sector and construct counterfactual supply curves for power plants in Texas had this price change not taken place. This then allows us to compare NO_x, SO₂, and CO₂ emissions between the observed scenario with a shock to LNG exports and the counterfactual scenario with no LNG exports shocks.

The remainder of this paper is structured as follows: Section 2 outlines the econometric framework, followed by data presentation in Section 3. Section 4 details the SVAR analysis results and robustness checks, while Section 5 explores the impacts on the electricity sector. Finally, Section 6 concludes the paper.

⁶The EIA uses the National Energy Modeling System (NEMS) to quantify projected impacts of LNG exports and other energy-related developments.

1.2 Econometric Framework

1.2.1 Structural VAR Model

In time series applications, Vector Auto Regressions (VAR) are commonly used to model the dynamic relationships among multiple variables over time. VAR models are essentially a system of endogenous equations. Given their endogenous nature, VAR models cannot directly disentangle the effect of any one of the variables on another, and are thus less useful for policy analysis. Moreover, if there is a persistent change in a variable, considered as a treatment, we expect economic agents to respond to the *treatment path* as opposed to a sequence of independent treatments (Ghanem and Smith, 2021). A structural VAR (SVAR) model allows one to extract treatment paths and dynamic responses by imposing structure on the system of equations. The structure entails imposing restrictions on the contemporaneous relationship between variables while allowing intertemporal dependence. Using these restrictions, one can extract sources of exogenous variation from the vector of endogenous variables.

The use of the structural time series approach in this analysis is inspired by Kilian (2009), particularly because we are dealing with a system of endogenous variables and only have time series data, without any cross-sectional variation. While alternative approaches such as event studies or regression discontinuity in time (RDiT) could be considered, they are more suitable when there is a sharp change in the time series due to an external event. For example, event studies leverage quasi-experimental methods to isolate the impact, as seen in Cozad and LaRiviere (2013), who used an event study design to estimate the effect of the 1979 oil shock on vehicle emissions. However, LNG exports have increased gradually, and natural gas prices did not exhibit sharp changes after the start of exports.

Our objective is to estimate the impact of LNG exports on natural gas prices. The primary challenge in this estimation is simultaneity—LNG exports affect natural gas prices, and natural gas prices influence exports. An instrumental variable (IV) approach could address this issue, but finding a valid instrument is challenging. Adverse weather events, such as hurricanes along the US Gulf Coast that disrupt export activity, could serve as potential instruments. However, these events

are often anticipated, allowing exporters to temporarily pause activities, and the natural gas market can adjust prices based on expected events.

This brings us to the Structural Vector Autoregression (SVAR) approach. Sims (1980) criticized the assumptions used in macroeconomic models to identify policy impacts and proposed the VAR as an alternative. This approach evolved into SVAR, which relies on the idea that unexpected movements in variables (forecast errors) are caused by structural shocks (Stock and Watson, 2018). In the context of mutually dependent natural gas prices and LNG exports, SVAR allows us to extract sources of exogenous variation, or shocks. These shocks act like randomly assigned treatments (Ramey, 2016), and their effects on the variables of interest can be quantified through impulse response functions (Stock and Watson, 2018). This method enables us to disentangle the impact of LNG exports on natural gas prices.

Accordingly, we analyze the impact of LNG exports on Henry Hub prices by setting up a 5 variable SVAR. The SVAR with 5 endogenous AR(1) variables can be expressed as

$$\mathbf{B}\mathbf{y}_t = \mathbf{\Gamma}_0 + \mathbf{\Gamma}_1\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t \quad (1.1)$$

where

$$\mathbf{y}_t \equiv \begin{bmatrix} \text{Production}_t \\ \text{Real Economic Activity}_t \\ \text{LNG exports}_t \\ \text{Storage}_t \\ \text{Henry Hub}_t \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 1 & b_{12} & b_{13} & b_{14} & b_{15} \\ b_{21} & 1 & b_{23} & b_{24} & b_{25} \\ b_{31} & b_{32} & 1 & b_{34} & b_{35} \\ b_{41} & b_{42} & b_{43} & 1 & b_{45} \\ b_{51} & b_{52} & b_{53} & b_{54} & 1 \end{bmatrix}$$

$$\mathbf{\Gamma}_0 = \begin{bmatrix} b_{10} \\ b_{20} \\ b_{30} \\ b_{40} \\ b_{50} \end{bmatrix}, \quad \mathbf{\Gamma}_1 = \begin{bmatrix} \gamma_{11} & \gamma_{12} & \gamma_{13} & \gamma_{14} & \gamma_{15} \\ \gamma_{21} & \gamma_{22} & \gamma_{23} & \gamma_{24} & \gamma_{25} \\ \gamma_{31} & \gamma_{32} & \gamma_{33} & \gamma_{34} & \gamma_{35} \\ \gamma_{41} & \gamma_{42} & \gamma_{43} & \gamma_{44} & \gamma_{45} \\ \gamma_{51} & \gamma_{52} & \gamma_{53} & \gamma_{54} & \gamma_{55} \end{bmatrix}, \quad \boldsymbol{\varepsilon}_t = \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \\ \varepsilon_{4,t} \\ \varepsilon_{5,t} \end{bmatrix}$$

y_t is our vector of variables of interest at the month level. We start with natural gas production in the US, as this is the starting point in Kilian (2009) for the oil market that has been replicated for natural gas modeling as well (Hou and Nguyen, 2018). Real Economic Activity (REA) is an index of global economic activity proposed in Kilian (2009). This business-cycle index is derived from a panel of dollar-denominated global bulk dry cargo shipping rates. It serves as a proxy for the volume of shipping in global industrial commodity markets. In addition to LNG exports and Henry Hub prices, we also include natural gas storage. Several studies (Kilian and Murphy, 2014; Janzen et al., 2017) have emphasized the importance of including storage or inventory in SVAR models of storable commodities. Modeling storage is crucial because it allows one to capture the effect of price expectations on the spot price of a commodity (Kilian and Zhou, 2020). Expectations about prices affect how much inventory an economic agent would hold. Given that the aim is to capture the role of price expectations, the literature acknowledges that this can be done through either futures prices or physical inventories. Janzen et al. (2017) use futures price spread to model the cotton market due to limitations on physical inventories data. In contrast, Kilian and Murphy (2014) rely on inventories to model the oil market, emphasizing that including inventory not only renders futures prices redundant but also aligns better with their focus on modeling the oil price in the physical market, rather than the oil futures market. We therefore follow Kilian and Murphy (2014) in our main specification but also estimate the same model where future prices replace storage, as a robustness check.

\mathbf{B} is the matrix of contemporaneous coefficients, such that b_{12} describes the contemporaneous effect of a change in REA on US natural gas production, b_{21} describes the contemporaneous effect of a change in production on REA, and so on. ε_t follows a white noise process with constant variance terms, and $E(\varepsilon_s \varepsilon_t) = 0, s \neq t$. In line with SVAR terminology, we refer to elements of ε_t as *shocks*. The uncorrelatedness of the shocks allows us to identify the effect of each independent shock. Premultiplying (1) by \mathbf{B}^{-1} yields the reduced form VAR:

$$\mathbf{y}_t = \mathbf{A}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \mathbf{U}_t \quad (1.2)$$

where $\mathbf{A}_0 = \mathbf{B}^{-1}\mathbf{\Gamma}_0$, $\mathbf{A}_1 = \mathbf{B}^{-1}\mathbf{\Gamma}_1$, and $\mathbf{U}_t = \mathbf{B}^{-1}\boldsymbol{\varepsilon}_t$. Since $\boldsymbol{\varepsilon}_t$ is white noise, \mathbf{U}_t has zero means and constant variance. Although the SVAR in (1) cannot be estimated directly due to feedback effects from contemporaneous variables, the reduced form in (2) can be consistently estimated by the least squares method. We can thus generate values for the residuals \mathbf{U}_t and obtain the coefficients in \mathbf{A}_0 and \mathbf{A}_1 . These can then be used to recover the SVAR parameters using the above-mentioned relationships.

However, the SVAR model with 2 lags (as will be explained in Section 4) contains 80 parameters⁷, while the reduced form model has 70 parameters⁸. We therefore need to calibrate 10 parameters in the SVAR model so that the number of unknown parameters is the same for both the SVAR and reduced VAR models, resulting in an exactly-identified system. This is where the restrictions on the contemporaneous relationships come in.

In this model, natural gas production or supply has a contemporaneous effect on all variables except REA. US natural gas production is unlikely to instantaneously impact dry cargo shipping rates. Moreover, natural gas production not being instantaneously responsive to other variables such as price is plausible because changing production is costly - once a well is drilled, extraction is not dependent on variables such as REA and price. This is also in line with Kilian (2009).

We let REA contemporaneously affect US LNG exports, natural gas storage, and Henry Hub prices. Janzen et al. (2017) allow REA to affect cotton prices and price expectations/storage, while Kilian and Murphy (2014) allow REA to affect oil storage and prices. It is also unlikely that US LNG export activity and other aspects of the US natural gas market would affect dry cargo shipping rates, also in line with the aforementioned studies for their respective commodities.

⁷Each equation in the 5 equation system has 1 intercept, 4 contemporaneous parameters, 10 lagged parameters, and own variance of the error term, resulting in 16 parameters each. With 5 equations, total number of parameters to be estimated is 80.

⁸Each equation has 1 intercept, 10 lagged parameters, and own variance of the error term, resulting in 60 parameters from all 5 equations. Adding the 10 covariance parameters, we end up with a total of 70 parameters.

US LNG exports are assumed to contemporaneously affect only storage and Henry Hub prices. This is plausible because higher export volumes can reduce the amount of natural gas available in the domestic market, leading to increased utilization of inventory. Additionally, LNG exports are assumed to be instantaneously influenced by production and REA.

It is essential to account for the endogeneity of storage flows concerning changes in gas prices, following Nick and Thoenes (2014) and Janzen et al. (2017). Gas storage facilities are likely to respond instantaneously to price fluctuations due to the economic rationale of inter-temporal price arbitrage for any commercial storage operator. Moreover, storage flows are expected to balance temporary supply and demand divergences caused by unforeseen market shifts. Therefore, we allow for the direct effects of production, REA, LNG exports, and Henry Hub prices on storage flows.

Imposing the above restrictions imply calibrating elements in the upper triangle of \mathbf{B} to 0, except for b_{45} which allows for the endogeneity of storage with Henry Hub prices. This results in

$$\mathbf{B} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ b_{31} & b_{32} & 1 & 0 & 0 \\ b_{41} & b_{42} & b_{43} & 1 & b_{45} \\ b_{51} & b_{52} & b_{53} & b_{54} & 1 \end{bmatrix}$$

Once the restrictions are imposed and the reduced form is estimated, the structural shocks can be recovered from the reduced form shocks using the relationship $\mathbf{U}_t = \mathbf{B}^{-1}\boldsymbol{\varepsilon}_t$.

1.2.2 Impulse Response Functions and Forecast Error Variance Decomposition

The coefficients obtained from the SVAR or reduced form models are not straightforward to interpret on their own. Instead, time series analysis utilize Impulse Response Functions (IRF) and Forecast Error Variance Decomposition (FEVD). IRFs characterize how each of the structural

shocks affect the variables in the system, tracing out their dynamic path, while FEVD describes the relative importance of structural shocks to the forecast error or expected variance in a variable. Following Hausman et al. (2012), we can utilize the VAR's moving average representation to write the structural IRF as: $y_t = \mu + \sum_{i=0}^{\infty} \psi_i \varepsilon_{t-i}$. Here, μ represents the mean, while each ψ_i quantifies the effect of a one-standard deviation shock to the error term ε_{t-i} at lag i on y_t . Inference is based on bootstrapped confidence intervals.

The forecast error variance decomposition (FEVD) is derived from impulse response coefficient matrices Ψ_h , which allows one to assess how variable j contributes to the forecast error variance of variable k over h steps. The forecast error variance $\sigma_k^2(h)$ for variable k at horizon h is calculated by summing the squares of the impulse response coefficients across all variables:

$$\sigma_k^2(h) = \sum_{n=0}^{h-1} (\psi_{k1,n}^2 + \dots + \psi_{kK,n}^2)$$

This represents the total variance of the forecast error of variable k at horizon h . To express the forecast error variance decomposition as a percentage, each term $(\psi_{kj,0}^2 + \dots + \psi_{kj,h-1}^2)$ is divided by $\sigma_k^2(h)$. This normalization yields the percentage contribution of each lagged impulse response coefficient to the total forecast error variance of variable k at horizon h .

1.3 Data

1.3.1 Prices and exports

We obtain monthly data from January 2011 to December 2023 from the Energy Information Administration (EIA) for the following variables: US natural gas production, storage, Henry Hub spot prices, and US LNG exports. The REA index is obtained from the Federal Reserve Bank of Dallas. Table 1.1 provide summary statistics for these variables. These variables are non-stationary at levels which may complicate inference and findings. We log transform and first difference all variables to achieve stationarity. Table 1.2 shows the results of unit root tests. We conduct

Augmented Dickey-Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests without trends for each variable. Both tests confirm the stationarity of the transformed variables.

Table 1.1: Summary Statistics

Variable	Unit of Measure	Mean	Median	Standard Deviation
Natural Gas Production	million cubic feet	2998043	2809094	491454
Real Economic Activity	% deviations from trend	-29	-30	47
LNG exports	million cubic feet	116043	53027	145205
Storage	million cubic feet	2731760	2773000	752336
Henry Hub	\$/mmBtu	3.35	2.975	1.281

Table 1.2: Unit root tests

Logged first differenced variables	ADF test statistic	KPSS test statistic
Production	-6.6	0.033
Real Economic Activity	-6.9	0.037
LNG exports	6.3	0.085
Storage	-9.3	0.016
Henry Hub	-5.1	0.043

The null hypothesis of the ADF test is that the variable contains a unit root. The null hypothesis of the KPSS test is that the variable is level stationary. The 5% critical values for the ADF and KPSS tests are 1.95 and 0.463 respectively. Thus we reject the ADF null of a unit root and fail to reject the KPSS null of stationarity for all variables.

1.3.2 Electricity generation and emissions

For the second component of our study, we illustrate how a change in natural gas prices following a shock to LNG exports affects electricity generation patterns and subsequent NO_x, SO₂, and CO₂ emissions. We limit our focus to power plants in Texas as much of the export activity takes place along the US Gulf Coast. We obtain monthly plant-level generation, fuel use, and fuel cost data from the Energy Information Administration (EIA) Form 923, “Power Plant Operations Report”. Although this dataset is available up to 2023, we limit our focus to 2022 only for two reasons. First, our objective here is to demonstrate how a change in natural gas prices affects power

sector emissions in the short run, as opposed to calculating historical emissions attributable to price changes. Second, we choose 2022 due to data limitations related to emissions. The Environmental Protection Agency (EPA) records power plant emissions through the Emissions and Generation Resource Integrated Database (eGRID), and 2022 is the latest year for which data is available. We obtain emissions factors (amount of emissions per unit of generation) for each power plant from this EPA dataset.

The EIA dataset contains information on all power plants in the US but has some missing values. We retain power plants in Texas with no missing values and which use one of these fuels – nuclear, coal, natural gas, and diesel fuel oil (DFO). These four categories make up the bulk of all capacity, not including solar and wind. The final dataset for 2022 results in 189 unique power plants. However, a power plant may have multiple generating units with different fuels. We match these power plants by fuel type with those in the EPA dataset to obtain emissions factors.

Given our focus on the Texas market, we retrieve hourly load, solar generation and wind generation data for 2022 from the Electric Reliability Council of Texas (ERCOT), the state grid operator. As will be discussed in Section 5, we use these variables to calculate residual demand.

1.4 SVAR Results

1.4.1 Main Findings

We estimate the 5-equation SVAR model as outlined in Section 2 with 2 lags as it is suggested by the Schwarz Information Criterion. Table 1.3 shows the contemporaneous coefficients. For ease of exposition, we focus on the IRFs and FEVD. Figure 1.2 shows the IRF for the effect of a one-standard deviation shock to logged LNG exports to Henry Hub prices. The impulse response coefficient is 0.15, implying that a one-standard deviation shock in logged LNG exports corresponds to a 0.16% move in Henry Hub instantaneously.⁹ However, the effect quickly falls to zero. The impulse response coefficient of 0.15 for a one-standard deviation shock to logged LNG exports can be scaled to percentage changes by incorporating the standard deviation of logged

⁹ $e^{0.215} - 1 \approx 16\%$

LNG exports - a 1% shock to LNG exports corresponds to approximately 7.9% change in Henry Hub prices, although the effect does not persist for long.

Table 1.3: Estimated SVAR contemporaneous coefficients

	Production	REA	LNG exports	Storage	Henry Hub
Production	1.000	0.0000	0.000	0.000	0.000
REA	0.000	1.0000	0.000	0.000	0.000
LNG exports	-0.540	0.1265	1.000	0.000	0.000
Storage	-0.190	-0.1341	0.048	1.000	-0.930
Henry Hub	0.130	0.0026	-0.128	-0.170	1.000

Note: In our notation, contemporaneous effects make up the B matrix, on the right-hand side of the set of equations. Moving these effects to the left-hand side reverses the signs. Thus, LNG exports have a positive effect on contemporaneous Henry Hub prices (with a coefficient of 0.128). These coefficients should be interpreted with caution because of the recursive nature of the system. For instance, Production contemporaneously affects LNG exports, which in turn contemporaneously affects Henry Hub prices. Thus the coefficient on Production in the LNG exports row is only part of the total effect, which also includes effects that feed through REA.

The above is a response to a one-time shock. A more likely scenario is that there would be more frequent shocks as LNG export capacity builds up. We thus compute the cumulative IRF (CIRF) for repeated shocks to LNG exports. As shown in Figure 1.3, the CIRF shows a persistent effect. However, a potential drawback of CIRFs is that they do not capture changes in expectations. Economic agents can incorporate repeated shocks into their expectations of the market, and shocks will no longer be considered as a true shock.

The FEVD in Table 1.4 shows the contribution of a shock to each variable to the expected variance in Henry Hub over 10 steps. Shocks to LNG exports explain 11.4% of the one-step ahead forecast error, and decrease to around 9.4% in the remaining steps. Storage explains around 3.6% and the bulk of the variation comes from unobserved variables (85%). Production and storage gradually become more prominent in explaining the variation in Henry Hub prices. Overall, these results suggest that shocks to LNG exports contribute a sizeable portion to the expected variance in Henry Hub.

Having identified the increase in natural gas prices following a shock to LNG exports, we now calculate the increase in natural gas prices in 2022 for the electricity component of our study. Ac-




Figure 1.2: Impulse Response Function of Henry Hub to a one-standard deviation shock in LNG exports
 Note: The shaded area shows the 95% confidence interval.




Figure 1.3: Cumulative Impulse Response Function

Table 1.4: Forecast Error Variance Decomposition for Henry Hub (one to ten steps ahead)

Production	REA	LNG Exports	Storage	Henry Hub
0.00093	0.00022	0.114	0.036	0.85
0.12354	0.00027	0.099	0.035	0.74
0.12848	0.00036	0.098	0.041	0.73
0.13649	0.00098	0.096	0.042	0.72
0.14794	0.00104	0.095	0.043	0.71
0.15483	0.00104	0.094	0.042	0.71
0.15593	0.00106	0.094	0.043	0.71
0.15576	0.00112	0.094	0.043	0.71
0.15569	0.00114	0.094	0.044	0.71
0.15577	0.00114	0.094	0.044	0.71

According to EIA’s Short Term Energy Outlook in 2021, projected 2022 LNG exports were 3704.75 bcf. Actual exports were 3865.64 bcf, and the difference of 4.3% can be regarded as a shock as the market did not anticipate these exports. The shock of 4.3% in exports thus translates to an approximately 34% increase in natural gas prices. This suggests prices would have been roughly 74% of the observed prices had the unexpected LNG exports not taken place.

1.4.2 Robustness checks

We consider some robustness checks. As mentioned in Section 2.1, an alternative to storage in modeling price expectations is to use futures prices. We therefore estimate the SVAR model where storage is replaced with one-month futures prices. Figure 1.4 compares the IRF from a shock to LNG exports on Henry Hub prices for the two models. The instantaneous IRF coefficient in the futures model is only slightly smaller, and exhibit a similar pattern with that of the storage model. This supports the idea that storage and futures can be treated as alternatives, as highlighted by Kilian and Murphy (2014) and Janzen et al. (2017).

We also estimate the SVAR under various periods (Figure 1.5). The IRFs for periods after 2016 are very similar, notwithstanding the differences in the error bands, with the impulse response coefficients varying between 0.12 and 0.16. In contrast, the coefficient for 2011-2015 is 0.02, suggesting a much smaller impact before the expansion of LNG exports from the lower 48 states. This reinforces the notion that LNG exports post 2016 have a different impact on Henry Hub.




Figure 1.4: IRFs under different models

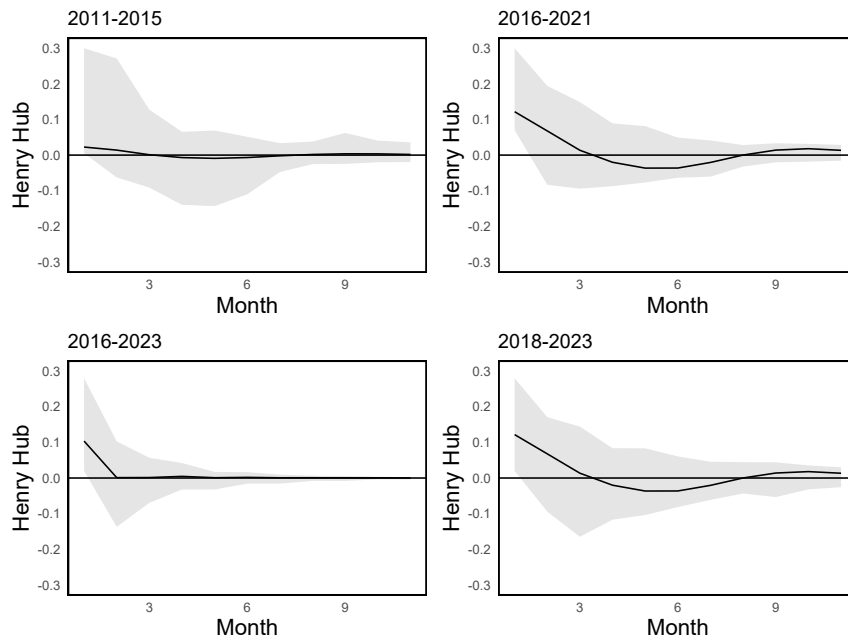


Figure 1.5: IRFs in different periods

1.5 Emissions impact

In the previous section, we calculated a 34% increase in natural gas prices in 2022 corresponding to a shock in LNG exports. We now turn to calculating the emissions impact of this increase. We construct monthly and hourly electricity dispatch curves to assess how electricity generation in Texas responds to changes in fuel prices, specifically a change in gas price. First, we obtain hourly load or demand data, along with hourly solar and wind generation. All of these variables are expressed in megawatts (MW). We subtract hourly solar and wind generation from hourly load to calculate the residual demand. The idea is that solar and wind generation, if available, are used first to meet demand, and the remaining or residual demand is met through other fuels. We then aggregate the hourly residual demand to the month level. The hourly and monthly residual demand (which are inelastic) serve as the respective demand curves.

Next, we compile fuel use or input data for each power plant (except solar and wind) in each month, expressed in million British thermal units (mmBtu). We then compile generation or output data expressed in megawatt hours (mwh), resulting in a dataset with fuel use and generation data for each power plant in each month. We use this data to calculate the heat rate of each power plant. The heat rate is a measure of efficiency and is simply the ratio of fuel use to electricity generated, expressed in mmBtu/mwh. Power plants either have positive or zero values for generation. Calculating the heat rate is potentially complicated because the zero generation values mean that we cannot calculate the heat rate of those plants. One potential workaround is to look for power plants that had zero generation values in some months but positive values in other months. This approach is appropriate as the heat rate or efficiency of a power plant is unlikely to vary across months. However, this approach drops all power plants that had no positive generation in 2022.

Once the heat rates are calculated, we import fuel cost data (expressed in cents/mmBtu) for each power plant and multiply it with the heat rate to calculate the variable cost of generating electricity, expressed in cents/mwh. Some power plants have excessively high variable costs which are outliers. We drop power plants with variable costs above the 95th percentile. We then sort power plants in terms of ascending variable costs. For each power plant, we obtain its nameplate

capacity in MW. The nameplate capacity refers to the maximum electricity a power plant can produce in an hour. For the monthly analysis, we multiply the capacity with 24 and 30 to obtain monthly capacity. For the hourly analysis, we retain the original nameplate capacity. With the capacity figures now available, power plants ordered in terms of ascending variable costs depict the monthly and hourly supply curves. Although our main interest lies in the hourly analysis, we constructed the monthly supply curve first because the variable costs are calculated based on monthly data. Data on power plant input, output, and fuel costs are not available at a more granular level. The hourly supply curve is simply a scaled down version of the monthly supply curve, and is based on the same variable costs.

With the hourly and monthly residual demand and supply curves now available, we intersect the demand and supply curves to identify which power plants are dispatched. Figures 1.6 and 1.7 illustrate sample monthly and hourly dispatch curves. Power plants on the left of the residual demand line are dispatched, while the ones on the right are not. There are four fuel types - natural gas, coal, diesel fuel oil (DFO), and nuclear. Nuclear is always dispatched first due to its low costs. Although natural gas power plants are usually dispatched first to meet the remaining demand, there are many natural gas plants that have higher variable costs than coal, which move them higher up the order.

In order to evaluate how power plants respond to changes in natural gas prices, we construct counterfactual dispatch curves. These curves illustrate the dispatch order had gas prices not increased by 34%, holding coal and oil prices constant. Given changes in natural gas prices, it is possible that some power plants which were not previously dispatched due to higher costs would now be dispatched.

Figures 1.6 and 1.7 also show the counterfactual dispatch curves for February 2022. There is a concentration of natural gas power plants earlier in the dispatch curve as natural gas is now cheaper. There are also fewer coal and oil fired plants on the left as some of these plants are displaced by natural gas generation.

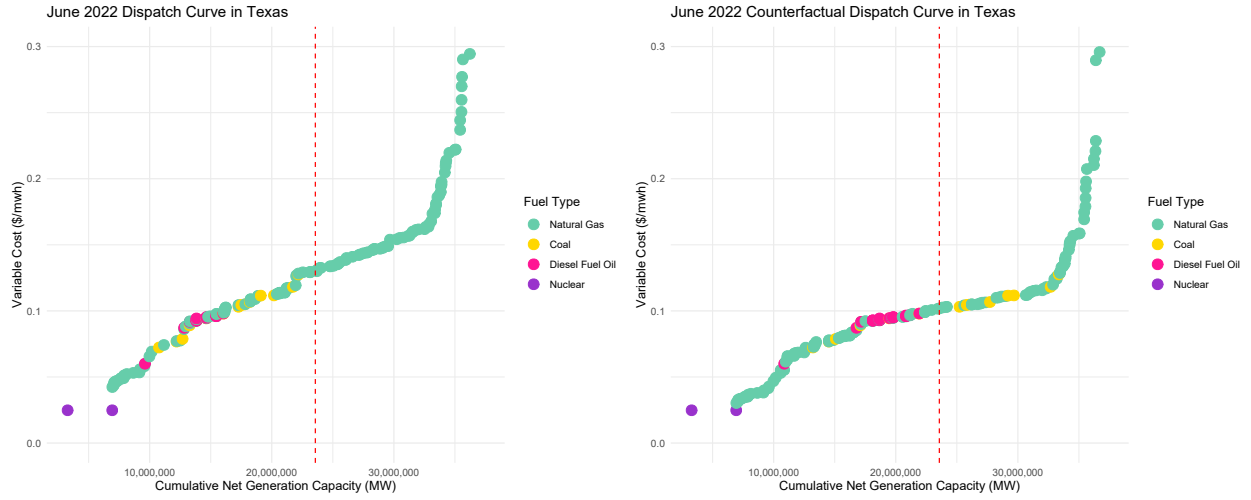


Figure 1.6: Illustrative Monthly Dispatch Curves

Note: The red line shows the residual demand.

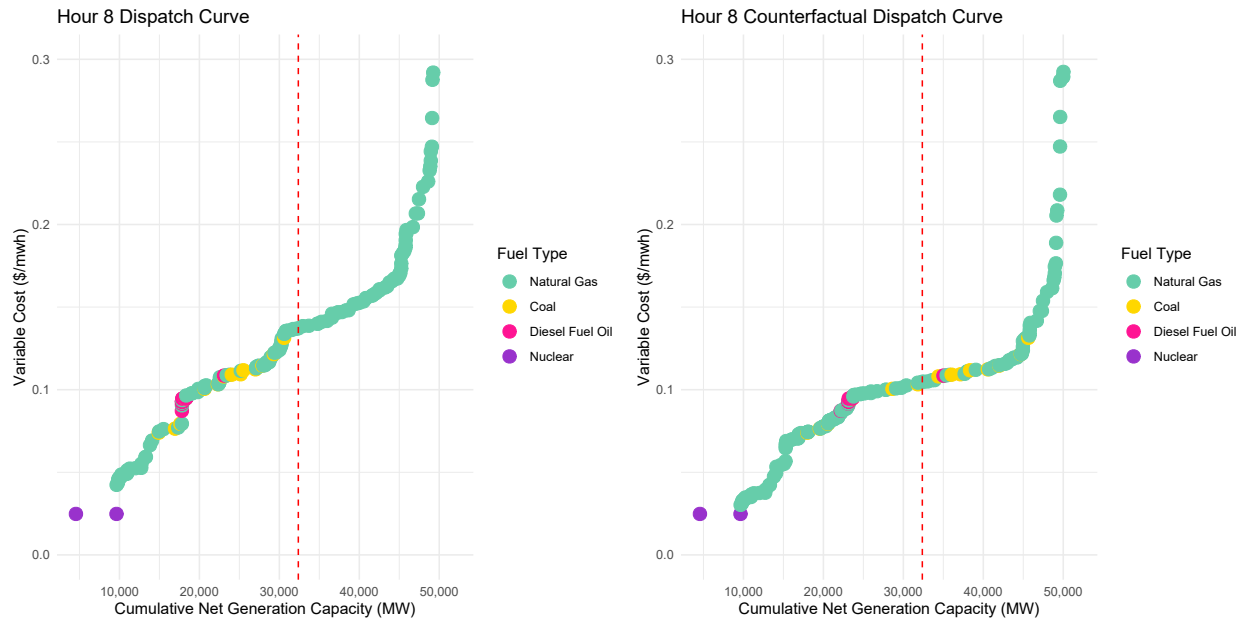


Figure 1.7: Illustrative Hourly Dispatch Curves

Note: The red line shows the residual demand.

Once the actual and counterfactual dispatch curves are constructed, we turn to calculating emissions. We bring in NO_x, SO₂, and CO₂ emissions factors for each power plant from the EPA dataset and multiply it with electricity generation to obtain hourly and monthly emissions from each power plant. This approach of calculating emissions can yield more accurate estimates of emissions compared to measuring emissions using monitors which may have confounding factors (Linn et al., 2014). Aggregating the monthly emissions, we find that actual emissions in 2022 were 175,367,697 lbs, 217,885,378 lbs, and 307,515,005,118 lbs for NO_x, SO₂, and CO₂ respectively. Counterfactual emissions were 155,602,363 lbs, 115,362,301 lbs, and 274,630,847,644 lbs respectively. Thus the 34% price increase associated with the shock to LNG exports in 2022 corresponds to a 12.7% higher NO_x emissions, 88% higher SO₂ emissions, and 12% higher CO₂ emissions from the electricity sector in Texas.¹⁰

Figure 1.8 presents average hourly emissions for each pollutant. Hourly emissions follow electricity generation patterns, remaining low in the early morning hours before picking up later in the morning. Emissions reach their peak around 8 pm when electricity generation is usually the highest. Although there is a substantial gap between observed and counterfactual emissions, this gap is lower when around 8 pm when electricity generation and emissions reach their peak. During hours of high demand, more of the available capacity is dispatched, and there is less scope for substitution. In line with our earlier finding, the difference between observed and counterfactual emissions is largest for SO₂, mainly due to the large differences in SO₂ emissions rate between coal and natural gas power plants.

We also calculate equilibrium prices, based on the intersection of residual demand and supply curves, at the hour and month level (Figure 1.9). Both monthly and hourly actual prices are more variable compared to counterfactual prices. Monthly prices peak in the summer when demand for cooling raises electricity demand. Hourly prices similarly reach their peak around 8 pm when demand is usually the highest.

¹⁰The large increase in SO₂ compared to other pollutants is likely due to the substantial difference in SO₂ emission rates between coal and natural gas plants, relative to the emission rates of other pollutants.

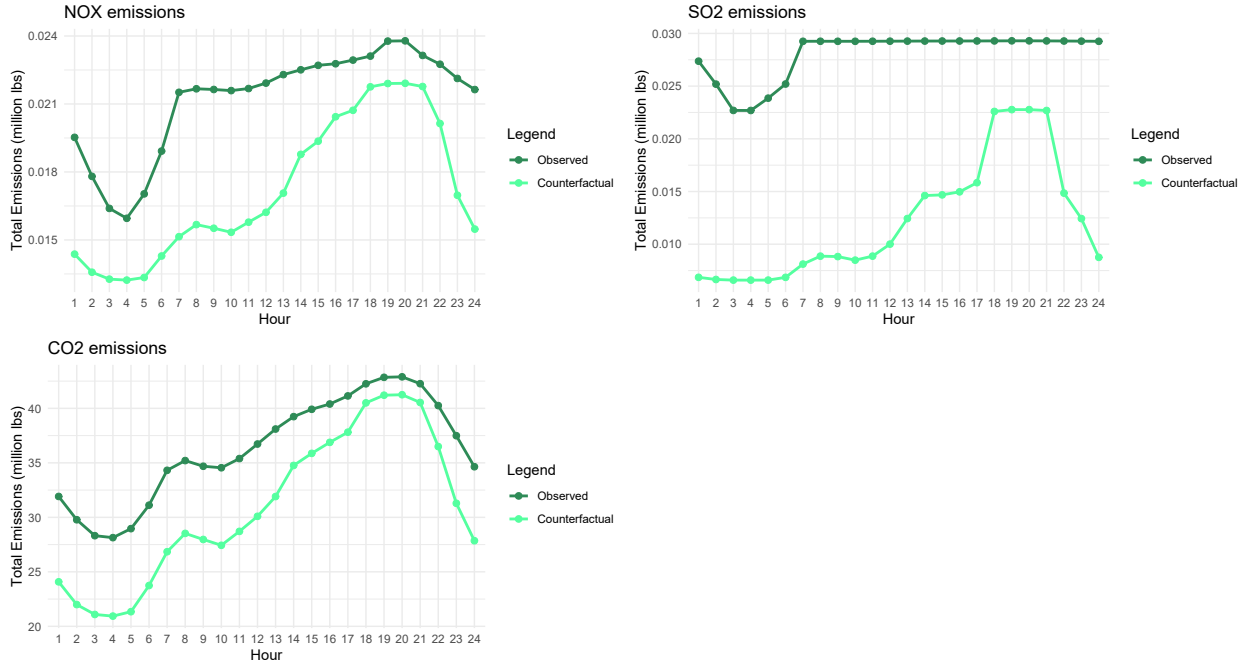


Figure 1.8: Average hourly emissions

Note: SO2 emissions between hours 7 and 24 vary only slightly, are are not constant.




Figure 1.9: Equilibrium prices

1.6 Discussion and Conclusion

In this study, we examined the impact of LNG exports on Henry Hub natural gas prices using an SVAR model and assessed the consequent emissions impact from changes in electricity generation in Texas. Our analysis provides insights into the short run relationship between LNG exports and natural gas prices, and the broader environmental implications stemming from shifts in the energy market. Our SVAR model indicated that a one-standard deviation shock to logged LNG exports corresponds to an immediate 0.16% increase in Henry Hub prices. This translates to a 34% increase in natural gas prices in 2022.

We then explored how this increase in natural gas prices affected the electricity sector in Texas by constructing monthly and hourly dispatch curves under the observed scenario with higher natural gas prices and a counterfactual scenario without the price increase. The higher gas prices make natural gas generators a less attractive option in electricity generation, and many of these generators were replaced by dirtier coal power plants. This shift resulted in a notable increase in emissions. Specifically, we observed a 12.7% increase in NO_x emissions, an 88% increase in SO₂ emissions, and a 12% increase in CO₂ emissions from the electricity sector in Texas, compared to the counterfactual scenario without the gas price increase. The difference between observed and counterfactual emissions is more pronounced for SO₂ given the larger difference in SO₂ emission rates between natural gas and coal plants, compared to the difference in emissions rates for the other pollutants. Hourly emissions show a substantial gap between observed and counterfactual scenarios, although this gap narrows during hours of peak demand later in the evening. We also explore how equilibrium electricity prices are affected, finding that observed prices fluctuate more than counterfactual prices, both at the month and hour level.

It is important to note some of the limitations of this study. Our analysis explores only the short term impacts of US LNG exports on natural gas prices, and the subsequent impact on electricity sector emissions. The electricity component of the study assumes constant coal and oil prices. This is a reasonable assumption, especially for coal, as coal contracts are typically long-term due to the extended time required for mining and transportation (Stock and Zaragoza-Watkins, 2024).

However, if higher natural gas prices persist, it could deter investments in natural gas capacity and delay the retirement of coal plants. This, in turn, could result in further changes in the relative prices of coal and gas in electricity generation. Stock and Zaragoza-Watkins (2024) investigate the long run equilibrium relationships between natural gas, oil, and coal prices under various regimes that include periods with and without LNG exports. They find that prior to LNG exports, the expansion of fracking depressed domestic natural gas prices and disconnected gas prices from oil and international gas prices. But since LNG exports began, domestic gas prices started to reconnect to international fuel prices, as well as domestic coal prices. They highlight two main mechanisms. First, Asian and European LNG contracts are typically linked to oil prices. US LNG exports thus facilitated greater integration of the US natural gas market to international markets. Second, gas and coal are competing fuels in domestic electricity generation. So both gas and coal prices have moved together with oil prices since 2016. Thus, if we are to consider long run impacts, it would be important to model how coal prices would respond to higher gas prices from LNG exports. It could be the case that if both gas and coal prices are substantially higher with LNG exports, then there will be increased investments in non-fossil generation capacity, thereby leading to *lower* emissions in the long run. This is what Stock and Zaragoza-Watkins (2024) find, reporting that power sector CO₂ emissions would be 145 million metric tons, or one-third, less in 2030 with LNG exports compared to a scenario without. Therefore, whether we focus on short-term or long-term impacts can substantially change the study's conclusions.

This would then affect the preference of specific fuels in electricity generation leading to further changes in emissions.

Another major limitation is that we focus only on the electricity sector in Texas to illustrate the emissions impact due to higher gas prices following LNG exports. Although a thorough welfare analysis is beyond the scope of this study, we do not consider the emissions implications in importing countries. If the importing country uses the natural gas to generate electricity, displacing dirtier coal and oil power plants, there would be positive environmental outcomes. If these changes are taken into consideration, it could offset the environmental damages within the US, especially

when considering global pollutants like CO₂. This would be an important factor in the discussion surrounding the economic and environmental consequences of allowing LNG exports from the US.

Our findings have important policy implications. The significant impact of LNG exports on natural gas prices and subsequent emissions highlights the need for comprehensive energy policies that consider both market dynamics and environmental consequences. The US is currently the largest LNG exporter in the world and is expected to substantially increase export volumes in the near future. At present, more than a dozen export terminals are either under construction or proposed. The findings of this study can inform the policy dialog on further expansion of export capacity. If the downstream air quality impacts are taken into consideration while conducting the social and environmental impact assessments of LNG export projects, it could lead policymakers to revise their conclusions on the benefits of such projects. The impact on hourly emissions is also worth considering. Higher gas prices result in a smaller reduction in electricity emissions, particularly SO₂, during off-peak hours compared to what would have occurred without the price increase. Although we hold demand constant, this could have implications for critical peak pricing and whether the emissions impact can be factored into such pricing schemes.

Although not accounted for in the methods, the findings can also have severe distributional impacts. Pollutants such as NO_x and SO₂ have adverse effects on public health, and given that low-income communities are more likely to be in closer proximity to power plants, they are more at risk. This can aggravate existing inequalities and lead to long term disadvantages in socioeconomic outcomes. The nexus between energy policy, public health, and socioeconomic realities illustrates the need for well-informed policymaking, and this study aims to support that by providing empirical evidence on the negative impact stemming from selected air pollutants.

Chapter 2

Can utilities influence residential water consumption through prosocial appeals? Evidence from a pilot study

2.1 Introduction

Prolonged drought conditions that persisted for much of the past decade have exacerbated water scarcity in the arid Southwest (Mankin et al., 2021). This region also experienced higher average temperatures between 2000 and 2020 compared to other periods, and future climate projections show a sharper increase in temperatures as well as higher variability in precipitation and dry days (USGCRP, 2023). These realities have led to a greater concern for water conservation and efficiency. Given that raising rates to encourage conservation is politically sensitive, water utilities have largely resorted to non-price interventions that include social comparisons and pro-social appeals among other means. Many of these non-price interventions aimed to lower *monthly* consumption. However, water utilities also experience peak use, occurring at particular hours of a day and particular days of a week, which impose substantial costs to their systems. With the increased proliferation of smart water meters, utilities can now observe real-time consumption and take steps to shift the timing of water use, both at the hour and day levels.

In this paper, we utilize hourly-level data to evaluate the effectiveness of a non-price voluntary program on both the timing and volume of water use. First, we explore whether households shift the timing of outdoor watering in response to appeals by the utility, which also include technical tips. Second, we test whether households respond to appeals for watering on assigned days. We find that while the timing of outdoor watering remains unchanged, there is weak evidence that one group of households water on assigned days while another does not, but overall we find no evidence of a significant average treatment effect. These results support previous findings that non-

price interventions that include conservation tips or pro-social appeals have little or no impact, as opposed to interventions that contain some element of social comparisons or other incentives.

This paper makes two contributions. First, it adds to the growing literature on evaluating non-price interventions, particularly those that are voluntary and seek to influence both timing and volume of consumption. Utilities are increasingly employing such interventions to encourage conservation and better management of limited resources. An improved understanding of the type of non-price interventions that are not effective in modifying behavior can help utilities in designing programs that can help achieve their objectives. This is particularly important for public water utilities as they are non-profits and have limited budgets for implementing these programs. The second contribution of this paper is that it demonstrates some of the pitfalls associated with nonrandomized programs and how analyzing their causal impact is complicated by the lack of a suitable counterfactual.

One type of non-price intervention that has gained traction throughout the US is the mailing of Home Water Reports (HWR) that compare a household's use with that of their neighbors. These reports are often accompanied by injunctive norms, such as a "smiley" face or a "frowny" face, that can nudge users to socially desirable behavior. Similar efforts are also widespread among energy utilities. Previous studies in the context of both water and energy find that social comparisons reduce consumption. Brent et al. (2015) report a 5% decrease in water consumption, while similar findings are reported in Allcott and Kessler (2019) and Ferraro and Price (2013). Studies also report heterogeneity in treatment effects, with high-use households responding more to social comparisons either because they have lower opportunity costs or because they are more likely to receive strong signals (Lurbe et al., 2023, Brent et al., 2020). Although social comparisons are a major component of non-price interventions, utilities also leverage other mechanisms to encourage conservation and modify behavior. Common practices include conveying technical tips, such as ways to identify leaks, and appealing to pro-social behavior, such as invoking better stewardship of limited resources (Stitzel and Rogers, 2023, Ferraro and Price, 2013).

Non-price interventions such as HWRs often contain multiple components that seek to influence user behavior through multiple channels. The incremental impact of these components and how they interact have been widely explored in the literature, with mixed findings. Technical tips alone have either a negligible impact on consumption (Ferraro and Price, 2013) or no impact at all (Brent et al., 2017, Seyranian et al., 2015, Schulze et al., 2016). However, augmenting technical tips with pro-social appeals can not only lead to significant reductions (Brent et al., 2017) but also have a lasting impact (Ferraro and Price, 2013). Results from Nemati et al. (2016) corroborate the notion that users do not respond to generic tips, but information on consumption such as monthly report emails and leak alerts can influence consumption. In the context of electricity consumption, however, Burkhardt et al. (2019) find minimal response to information provision or pro-social appeals.

The evaluation of these non-price interventions depends on how they are designed and implemented. In many cases, the interventions are field experiments in which treatment is randomly assigned. The vast majority of the literature cited above analyzes programs with an experimental design with one or multiple treatment groups and a control group that does not receive the treatment. However, utilities often implement programs in which customers are nonrandomly assigned to a particular treatment and participation is voluntary. This is particularly prevalent in the arid Southwest where utilities are increasingly implementing voluntary programs to reduce outdoor watering, which can account for up to 50% of residential water use.¹¹ These are mostly rebates for xeriscaping or adopting efficient irrigation technologies.¹² Voluntary programs require an active decision to participate and suffer from some degree of selection bias because there will be inherent differences between those who participate and those who do not. Sherwin et al. (2022) question whether it is even possible to formulate a suitable counterfactual for voluntary programs using data available to utilities. Although the non-randomized nature of such programs compli-

¹¹Outdoor watering is also often subjected to mandatory restrictions, especially during dry periods. Renwick and Green (2000) find that mandatory restrictions, violations of which can result in fines, have a stronger effect on reducing consumption compared to voluntary measures.

¹²It could be argued whether rebate programs can be categorized under non-price interventions, since they entail financial incentives. However, the term ‘price’ in non-price interventions refer to water rates paid by users, which are not directly impacted by rebate programs.

cates the finding of a treatment effect, the literature is replete with studies that evaluate voluntary programs but caution that results are not necessarily causal (Boomhower and Davis, 2022, Fielding et al., 2013, Kurtz et al., 2013, Nemati et al., 2016, Seyranian et al., 2015, Schulze et al., 2016).

In addition to total consumption, utilities are often also concerned with the timing of consumption. Large peaks in consumption can impose added costs to utilities and stress their systems. There is a vast literature on how electric utilities can influence the timing of consumption (Boomhower and Davis, 2022, Burkhardt et al., 2019). Although most of the studies in the water context analyze the effect of non-price interventions on monthly consumption, an increased proliferation of smart meters that provide higher-frequency consumption data allows a more granular analysis. Nonetheless, the literature is largely limited to evaluating the effect on daily consumption (Brent et al., 2015, Brent et al., 2020) as opposed to comparing consumption between days of week. Moreover, even when hourly data is available, the analysis is aggregated at the daily level (Nemati et al., 2016).

We aim to address this gap in the literature by evaluating an intervention in Denver, Colorado that aimed to reduce peaking by using a pro-social appeal to shift the time and days of outdoor watering. In line with much of the previous literature, we find no significant effect of the pro-social appeal, highlighting the need to augment such appeals with economic incentives or social comparisons to make them more effective.

The rest of the paper proceeds as follows. In Section 2, we present the program context and our data. Section 3 describes the methodology and main results. Section 4 concludes.

2.2 Program Design and Hypotheses

2.2.1 Program Design and Context

Most residential sprinklers in Denver are set to turn on at 5 am on Mondays, Wednesdays, and Fridays¹³. The timing of outdoor watering coincides with heavy indoor usage as people begin their

¹³Figure 2.1 illustrates how water use peaks on Mondays, Wednesdays, and Fridays in a given week, while Figure 2.3 shows the peak at 5 am on a given day. The timing of outdoor watering can be inferred by comparing the difference in average hourly consumption between June and October, for instance. Outdoor watering usually takes place in the

days. This results in residential water use peaking on these three days within a week, and early in the morning within those days. Between 2017 and 2022, the day with the maximum consumption had on average 40,000 gallons more than the day before and after.¹⁴ This peaking, both at the hour and day level, results in infrastructural, operational, and energy costs. Beginning in May 2023, Denver Water implemented the Peak Shaving Pilot program which aimed to shift the timing of outdoor watering and eliminate these peaks. The program divided all single-family residences in the service area that have smart meters into odd and even groups based on whether their address is odd or even. Only households with smart meters were selected because the utility can observe real-time water consumption every 15 minutes.¹⁵ Starting on May 2023, households in both groups received a mailer detailing the problem of peaking and requesting them to shift both watering time and day (Figure 2.2). The suggested watering time was between 1 and 3 am for both groups. The suggested watering days were Wednesdays and Saturdays for odd-numbered households, and Thursdays and Sundays for even-numbered households. These messages were not tailored in that all households received the same mailer that contained all this information. If the program were 100% effective, we would expect to see a reduction in overall consumption because houses are told to reduce watering from 3 to 2 days and we expect to see a shift in the timing of consumption. We evaluate the effectiveness of this pilot program in shifting peak time and day of water use.

summer months. The difference between average consumption at 5 am in June and that in October is the amount attributed to outdoor watering.

¹⁴Correspondence with Denver Water

¹⁵In contrast, water consumption for households without smart meters is observed once a month when their meters are read and the monthly bill is issued.

Average water use in 2022




Figure 2.1: Average Water Use

Outdoor Watering Rules are in effect May 1 to Oct. 1.

These times are based on a sprinkler system running at an efficient level. Watering is not allowed between 10 a.m. and 6 p.m.

This chart is a starting point for the number of minutes to water each sprinkler zone on the days you determine watering is needed. These times are averages. Adjust your watering minutes based on rainfall, type of grass or plants, sunny or shady locations and other characteristics.

Watering two days a week should be sufficient during most of the summer. If needed, water a third day only during extreme heat or dry periods. Water trees and shrubs as needed, but not between 10 a.m. and 6 p.m. Shrubs and perennials need half as much water as lawns.

There are no assigned watering days, but never water a zone more than three days per week.

Lawn Watering Times

Minutes to Water Per Zone
(For lawns, based on watering three days per week.)

Watering Months	Water trees/shrubs as needed.			
	Flood Spray Heads	Rotor Heads	Rotary Nozzles	Manual Sprinklers
January–April	Water trees/shrubs as needed.			
May	12	24	30	18
June	17	35	43	26
July	18	36	45	27
August	14	27	34	20
September	11	23	28	17
October–December	Water trees/shrubs as needed.			

These times are based on a sprinkler system running at an efficient level. Watering is not allowed between 10 a.m. and 6 p.m.

Visit denverwater.org/Lawn for more information and tips.

Shave The Peak
Some days just aren't your day. Know when to water.

Did you know most sprinkler clocks in Denver are set to run at 5 a.m. on Monday, Wednesday and Friday during the summer? It's also about the time many people wake up, take showers and start their morning routines.

We know this, because our system operators see water use peak at this time throughout the summer.

The thing is, building our system to meet these peak water use times carries a cost. We must size and build our system and incur higher energy costs to meet peak demand times. (And you might notice the water pressure is slightly less during these times, too.)

Help us shave the peak by taking a few small steps:

Shift your watering time:
Instead of 5 a.m., set your sprinkler clock to run between 1-3 a.m.

Then shift your watering days:

- For EVEN-numbered addresses (ending in 0, 2, 4, 6 or 8), set your sprinkler clock to run **Thursday and Sunday**. (Add Tuesday only during extreme heat or dry periods, if needed.)
- For ODD-numbered addresses (ending in 1, 3, 5, 7 or 9), set your sprinkler clock to run **Wednesday and Saturday**. (Add Monday only during extreme heat or dry periods, if needed.)

Figure 2.2: Sample Mailer

2.2.2 Statistical Identification

In contrast to other interventions, such as households receiving or not receiving home water reports, there is no specified control group in this setting. The entire population of households with a smart meter received the mailer and is therefore treated. We need an appropriate comparison group and period to estimate the treatment effect. One option is to compare the treated group with those residential users who do not have a smart meter and are therefore not treated. However, we only observe monthly consumption for users without a smart meter, so comparing time and day of use is not possible. While a monthly analysis could provide some insights on overall consumption, the aim of the program is to shift the time and day of outdoor watering. We focus only on the users with smart meters in our analysis and explore options such as comparing each group with their behavior on days they are not supposed to water as well as days prior to treatment. We further explore comparing the odd houses on days when they are treated to even houses on days when the odd houses are treated. Although this method violates the stable unit treatment values assumption (SUTVA), it could provide a lower bound on an estimate of the Average Treatment Effect.

The program is voluntary and there are no punitive measures for not following the recommended schedule. Our identifying assumption is based on households not knowing about the program prior to receiving the mailer, particularly which days would be assigned to which group. Although it could be argued that households *can* influence whether they live in an odd or even numbered address by moving accordingly, the last digit of the address is unlikely to be a major criteria for moving decisions. Given that the population is single-family residences, households that moved within the study period are unlikely to constitute a large portion of the data, as opposed to households living in multi-family buildings who are more likely to be renters and move more frequently.

Moreover, we assume we are estimating a lower bound on the treatment effect for the following reason. If all houses perfectly responded to treatment, we would expect to see no outdoor consumption on days houses were asked not to water. This means, for example, odd houses would only water on Wednesdays and Saturdays and no other day. Therefore odd houses, since their wa-

tering schedule is randomly assigned, is a sufficient control group for even houses on even house watering days (Thursdays and Sundays), and vice versa. The estimated size and statistical strength of the "treatment effect" depends on the magnitude of the response from each group of houses, just like a randomized control trial.

Our hypothesis is that odd households will consume more water on Wednesdays and Saturdays between 1 am and 3 am relative to other days, and even households will consume more water on Thursdays and Sundays between 1 am and 3 am following the intervention because this is what they were asked to do. If true, the results would suggest that the program was effective in shifting the day and time of use as desired, and that households in our study respond to recommendations about water use.

2.2.3 Data

We obtain hourly water use data from January 2021 to December 2023 for all residential users equipped with a smart meter. The number of users with smart meters increased each year of the study period, ending with 2607 unique users with an almost equal split of odd and even numbered addresses. Table 2.1 shows the summary statistics of daily water use. Water consumption is higher and more variable during the summer due to outdoor watering, which does not typically occur during the winter months. Odd households consume more water on average, and the differences between the two groups are statistically significant. This is a somewhat puzzling finding because we would expect water consumption to be uncorrelated with home addresses being odd or even. Odd houses tend to face east in our sample, while even houses tend to face west. This could drive the difference in consumption. Given that outdoor watering is sensitive to weather conditions, we combine the hourly water use data with temperature and precipitation data recorded at the Denver Water Department weather station which is retrieved from the Colorado Climate Center.

When aggregated across all households, there is substantial variation in hourly consumption and across years. However, hourly consumption follows a similar pattern across years; consumption peaks early in the morning around 5 am, remaining stable for the rest of the day before in-

Table 2.1: Summary Statistics of Daily Residential Water Use (in gallons)

	N	Mean	Median	SD
All months				
Overall	2607	294	115	2386
Odd	1307	305	119	2307
Even	1300	282	111	2467
Difference		23		
Summer months (June, July, August)				
Overall	2607	533	190	1249
Odd	1307	557	200	1208
Even	1300	508	180	1290
Winter months (January, February, March)				
Overall	2607	123	94	309
Odd	1307	125	96	392
Even	1300	121	93	180

creasing again around 8 pm. Figure 2.3 illustrates the average hourly consumption for June and October in all three years. Outdoor watering is prevalent in June, which explains the peak early in the morning. In contrast, this peak is absent in October when water use is mostly limited to indoors. Average hourly consumption is lower in 2023, but this decrease is not necessarily attributable to the intervention given the yearly variation in weather. 2022 was a dry year with total precipitation of 1.74 inches in June, whereas 2023 was a relatively wet year with total precipitation of 5.61 inches in June thereby reducing the demand for outdoor watering.

Given that outdoor watering is mostly prevalent in summer, we narrow our focus to the months of June, July, and August.¹⁶ Figure 2.4 shows the average hourly consumption on dry summer days before and after the intervention for odd and even households. We focus on dry days, defined as days with zero precipitation on that day and the day before, to reduce the impact of precipitation on water use. Recall that odd households are asked to water on Wednesdays and Saturdays, while even households are asked to water on Thursdays and Sundays. There is an average decrease

¹⁶Although June is sometimes considered to be part of Spring, we define summer as these three months throughout this paper.

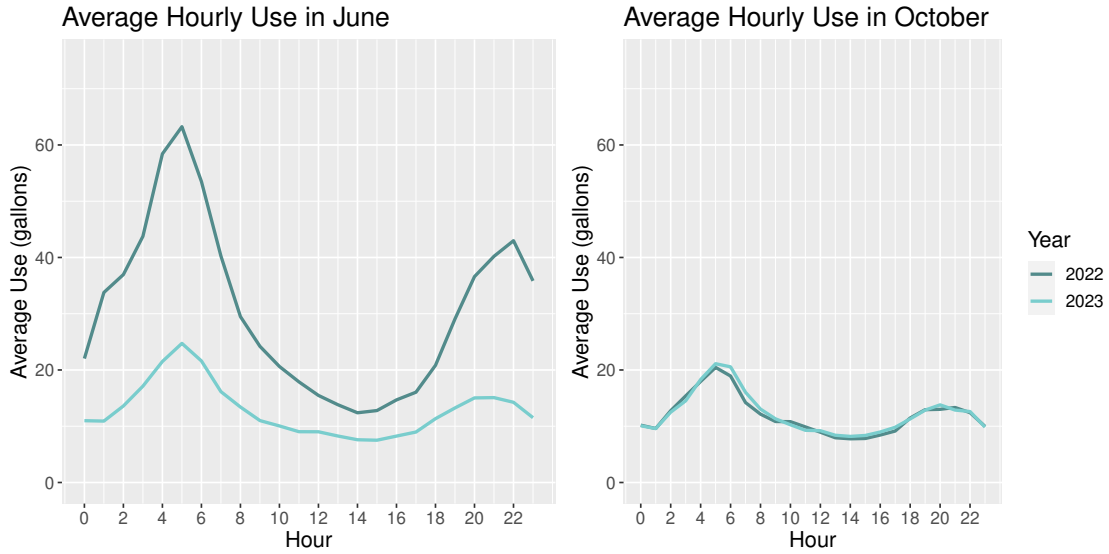
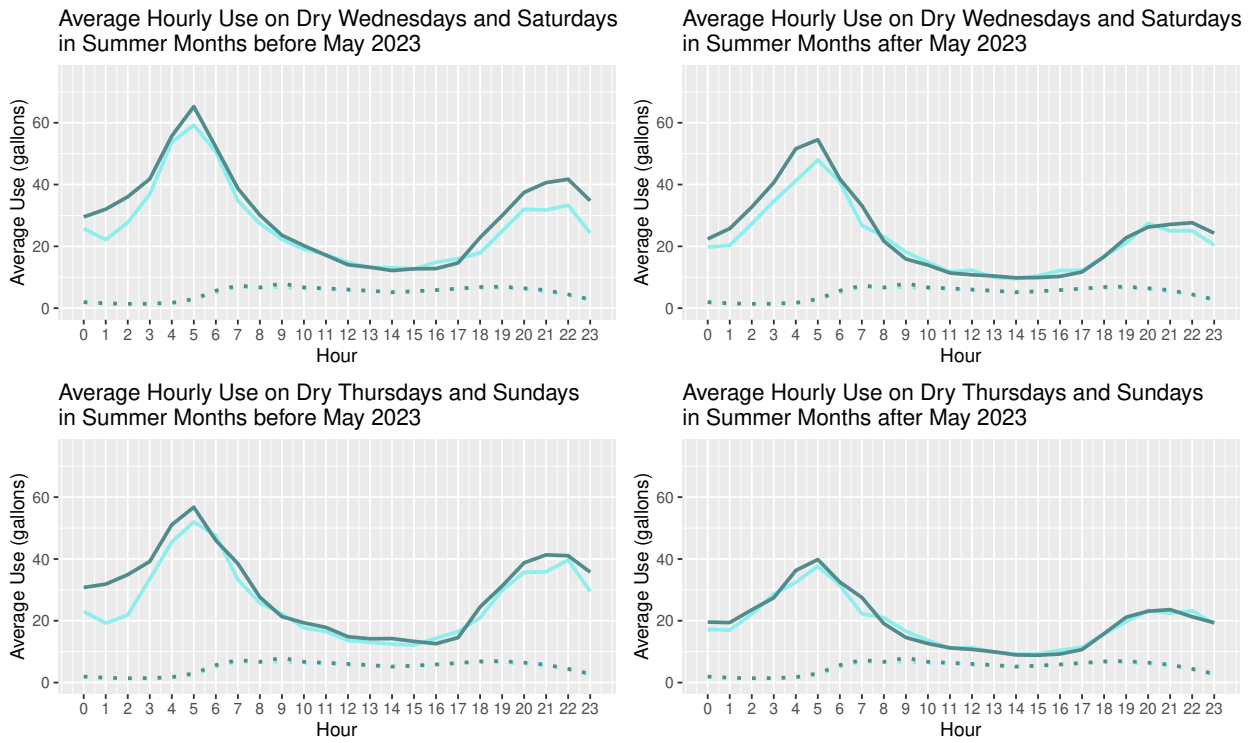


Figure 2.3: Average hourly use in June and October



Note: The dark and light lines refer to odd and even houses respectively. The dotted lines are the average winter consumptions (AWC), a proxy for indoor water use.

Figure 2.4: Average hourly use on dry treatment days

of around 10 gallons for both odd and even households on dry Wednesdays and Saturdays after the intervention (comparing the Pre-May 2023 to the Post-May 2023), although the peak time of 5 am remains unchanged. This illustrates that both odd and even households are turning on their sprinklers (as demonstrated by the peak) on dry days, but they are turning them off sooner which results in lower use. One plausible explanation is that the effect of precipitation on outdoor watering can persist for many days; if it rains heavily on just one day of a week, households can still alter their watering behavior and use less water for the rest of the week. There is a larger decrease of around 20 gallons for odd households on dry Thursdays and Sundays after the intervention, while even households decrease their consumption by around 10 gallons. Overall, the total consumption appears to decrease among both even and odd houses from 2022 to 2023, but the timing of consumption remains unchanged.

Another important takeaway from Figure 2.4 is that both odd and even houses appear to have the same distribution of consumption, using roughly the same amount of water during peak hours, despite the day of the week. In other words, there is no obvious gap between odd and even house consumption on Wednesdays and Saturdays or on Thursdays and Sundays. There are houses in both groups that disregard the program and water on days they were asked not to water on. This provides preliminary visual evidence that the appeals for conservation had no effect.

Given that there is a change in overall consumption, we now turn towards isolating the impact of the intervention on consumption aggregated at the daily level. We also use empirical methods to analyze changes in the timing of water use.

2.3 Methods and Results

2.3.1 Estimating Average Treatment Effect

Analyzing changes in watering days

The methodology and results are presented in several sections. First, we estimate difference-in-difference models for the two sets of dry ‘treatment days’, that are Wednesdays/Saturdays and Thursdays/Sundays. The models are as follows:

$$\text{Consumption}_{it} = \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Odd}_i + \beta_4 (\text{Post}_t \times \text{Odd}_i) + \gamma_i + \phi_t + \epsilon_{it} \quad (2.1)$$

if day = Wed, Sat

$$\text{Consumption}_{it} = \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Even}_i + \beta_4 (\text{Post}_t \times \text{Even}_i) + \gamma_i + \phi_t + \epsilon_{it} \quad (2.2)$$

if day = Thu, Sun

where Consumption_{it} is the total water consumption of household i on date t , Post_t is a post treatment indicator equal to 1 for period after May 2023 and 0 for before, Odd_i and Even_i indicate whether household i is odd or even, γ_i is a household fixed effect, ϕ_t is a month fixed effect, and ϵ_{it} is a normally distributed error term. The "control group" in equation 2.1 is even houses on Wednesdays and Saturdays and the "control group" in equation 2.2 is odd houses on Thursdays and Sundays.

Table 2.2 presents the results of estimating equations 2.1 and 2.2 in the first two columns. The coefficients on Odd and Even are omitted in equations 2.1 and 2.2 respectively due to multicollinearity with the fixed effects. The coefficients on both interaction terms are positive which aligns with our apriori expectation that consumption would increase on the 'treated' days, but they are insignificant. However, note that total water consumption is significantly lower in the post period. This can include the effects of increased precipitation, because even though we are focusing on dry days, the impact of heavy precipitation can persist for multiple days.

Table 2.2: Results for Base Models

Dependent Variable: Model:	Water Consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Post	-76.58*** (23.07)	-186.4*** (31.32)	-85.90*** (25.14)	-123.6* (42.02)	-156.0*** (36.50)	-199.8* (55.81)
Post × Odd	3.488 (41.55)		-56.07 (44.97)	-70.71** (12.22)		
Post × Even		61.21 (38.25)			65.07 (43.88)	70.70** (12.20)
Wed_Sat			3.066 (17.15)	8.345 (11.49)		
Post × Wed_Sat			12.26 (19.60)	6.468 (13.63)		
Odd × Wed_Sat			-8.730 (22.72)	-8.891 (8.092)		
Post × Odd × Wed_Sat			36.50 (25.14)	36.57** (7.115)		
Odd				86.98** (9.413)		
Thu_Sun					-15.40 (22.28)	-10.97 (11.82)
Post × Thu_Sun					-35.06 (21.84)	-40.36** (5.120)
Even × Thu_Sun					2.860 (33.24)	2.300 (22.57)
Post × Even × Thu_Sun					6.600 (33.29)	7.351 (28.00)
Even						-86.93** (9.398)
<i>Fixed-effects</i>						
Meter_ID	Yes	Yes	Yes	No	Yes	No
Month	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	51,342	49,223	128,848	128,848	126,729	126,729
R ²	0.65391	0.63704	0.42701	0.00302	0.39247	0.00448

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Next, we test whether odd households consume more on Wednesdays/Saturdays after the intervention relative to days other than Thursdays/Sundays. The idea is that by excluding Thursdays/Sundays, days in which even households are "treated", we are again setting even households as the "control" group. The model is as follows:

$$\begin{aligned}
\text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Odd}_i + \beta_3 (\text{Wed, Sat}) + \beta_4 (\text{Post}_t \times \text{Odd}_i) \\
& + \beta_5 (\text{Post}_t \times \text{Wed, Sat}) + \beta_6 (\text{Odd}_i \times \text{Wed, Sat}) \\
& + \beta_7 (\text{Post}_t \times \text{Odd}_i \times \text{Wed, Sat}) + \gamma_i + \phi_t + \epsilon_{it}
\end{aligned} \tag{2.3}$$

if day \neq Thu, Sun

where *Wed, Sat* is an indicator for whether the day is Wednesday or Saturday. β_7 is the coefficient of interest which shows the difference between odd and even houses before and after the intervention and on Wednesdays/Saturdays compared to other days except Thursdays/Sundays. Equation 2.3 is estimated with and without household fixed effects with results shown in columns (3) and (4) of Table 2.2. β_7 is positive as expected but insignificant with household fixed effects. However, it becomes positive and significant when leaving out household fixed effects, suggesting odd households consume 36 more gallons on dry Wednesdays/Saturdays after the intervention. The insignificant results with fixed effects can indicate substantial heterogeneity among households that mask the treatment effect.

We conduct a similar analysis for even houses, testing whether they consume more on Thursdays/Sundays after the intervention relative to days other than Wednesdays/Saturdays. The model is as follows:

$$\begin{aligned}
\text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Even}_i + \beta_3 (\text{Thu, Sun}) + \beta_4 (\text{Post}_t \times \text{Even}_i) \\
& + \beta_5 (\text{Post}_t \times \text{Thu, Sun}) + \beta_6 (\text{Even}_i \times \text{Thu, Sun}) \\
& + \beta_7 (\text{Post}_t \times \text{Even}_i \times \text{Thu, Sun}) + \gamma_i + \phi_t + \epsilon_{it}
\end{aligned} \tag{2.4}$$

if day \neq Wed, Sat

where Thu, Sun is an indicator is an indicator for whether the day is Thursday or Sunday, and with results shown in the last two columns of Table 2.2. Unlike the case with odd houses, the coefficient on the triple interaction term is not significant in either specification, and much smaller in magnitude, suggesting that the treatment appears to have little or no effect on even houses. We can also compare how odd and even houses behave on Mondays, Tuesdays, and Fridays (untreated days) after the intervention by looking at the sum of β_1 and β_4 in Equations (2.3) and (2.4) estimated without household fixed effects (columns 4 and 6 in Table 2.2). As expected, odd houses reduced consumption by 194 gallons on these days after the intervention compared to before, while even houses reduced consumption by 129 gallons. This supports our later conclusion that some odd houses are more responsive to the treatment relative to even houses.

Recognizing that focusing on dry days may not adequately control for the persistent effects of precipitation, we estimate equations 3 and 4 for all days and add average temperature and precipitation as additional controls.¹⁷ Table 2.3 shows the results in columns 2 and 4 with the previous estimates reproduced in columns 1 and 3 for ease of comparison. Temperature and precipitation are positively and negatively correlated with water consumption respectively, as expected. But even with additional weather controls, the triple interaction terms are insignificant as in the previous estimates. If the intervention were effective, we would expect odd and even houses to have significantly higher consumption on their treated days after the intervention. We would also expect both groups to significantly reduce their consumption on untreated days. These results imply that there is no statistically significant evidence of either phenomenon, and leads us to conclude that there is no treatment effect.

¹⁷For each day, these variables are the averages of the last 7 days.

Table 2.3: Regressions with weather controls

Dependent Variable: Model:	Water Consumption			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
Post	-85.90*** (25.14)	-113.1*** (20.65)	-156.0*** (36.50)	-155.0*** (34.62)
Wed_Sat	3.066 (17.15)	10.97 (12.18)		
Post × Odd	-56.07 (44.97)	-37.75 (40.74)		
Post × Wed_Sat	12.26 (19.60)	-15.58 (13.74)		
Odd × Wed_Sat	-8.730 (22.72)	-5.538 (17.48)		
Post × Odd × Wed_Sat	36.50 (25.14)	28.04 (17.86)		
Avg_Temp		10.15*** (0.5009)		9.207*** (0.4638)
Avg_Precip		-293.9*** (32.74)		-287.2*** (29.39)
Thu_Sun			-15.40 (22.28)	-30.26 (20.21)
Post × Even			65.07 (43.88)	43.03 (40.30)
Post × Thu_Sun			-35.06 (21.84)	-9.846 (18.42)
Even × Thu_Sun			2.860 (33.24)	9.164 (29.21)
Post × Even × Thu_Sun			6.600 (33.29)	2.433 (26.64)
<i>Fixed-effects</i>				
Meter_ID	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	128,848	266,938	126,729	267,825
R ²	0.42701	0.44317	0.39247	0.42274

Clustered (Meter_ID) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The next set of results is for models where odd and even households serve as their own controls.

We estimate separate regressions for odd and even households as follows:

$$\text{Consumption}_{it} = \beta_0 + \beta_1(\text{Wed, Sat}) + \beta_2\text{Post}_t + \beta_3(\text{Post}_t \times (\text{Wed, Sat})) + \gamma_i + \phi_t + \epsilon_{it} \quad (2.5)$$

if address = Odd

$$\text{Consumption}_{it} = \beta_0 + \beta_1(\text{Thu, Sun}) + \beta_2\text{Post}_t + \beta_3(\text{Post}_t \times (\text{Thu, Sun})) + \gamma_i + \phi_t + \epsilon_{it} \quad (2.6)$$

if address = Even

Table 2.4 shows the results under different specifications. Odd households consume more water on Wednesdays/Saturdays after the intervention, and the effect is stronger in the summer. In contrast, there is no corresponding effect for even households. It could be the case that even households, which are supposed to water on Thursdays/Sundays, are not responding to the intervention by sticking to their usual watering days. Higher usage on Wednesdays/Saturdays among odd households is also not necessarily fully attributable to the intervention. It is possible that odd households, which previously watered on Wednesdays among other days, are continuing to water on Wednesdays as usual and not because they are following the recommendations. Finally, another explanation could be that some households respond while some do not, yielding positive but statistically insignificant treatment effects.

The overall notion of no treatment effect is further corroborated by Figure 2.5 which shows average consumption on dry summer days. In 2023, after the intervention, Odd houses consume the highest on Wednesdays, but only 7 gallons more than their consumption on Mondays. If the program were effective in shifting watering days to Wednesdays and Saturdays, we would expect the highest consumption on these days. But Odd houses consume less on Saturdays compared to Mondays. For even houses, the highest consumption after the intervention occurs on Mondays, an untreated day. They also consume the least on Sundays, one of their treated days. This visual evidence, combined with the previous regression results, supports the overall conclusion that the program is not effective in shifting watering days.

Table 2.4: Regression results

Dependent Variable: Model:	Water Consumption			
	Odd	Odd & Summer	Even	Even & Summer
<i>Variables</i>				
Wed_Sat	1.378 (3.359)	-0.8592 (13.54)		
Post	-89.62*** (26.61)	-161.7*** (33.00)	-27.78*** (10.67)	-85.50*** (19.74)
Wed_Sat × Post	51.74** (20.11)	67.38*** (14.84)		
Thu_Sun			-1.065 (5.417)	-14.30 (21.78)
Thu_Sun × Post			-64.36 (43.85)	-33.96 (22.01)
<i>Fixed-effects</i>				
Meter_ID	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	496,748	91,948	469,926	86,123
R ²	0.02314	0.57365	0.01992	0.34831

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Analyzing changes in timing of water use

We can analyze changes in the timing of water use through the following specifications:

$$\begin{aligned}
\text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Early}_t + \beta_3 (\text{Wed, Sat}) + \beta_4 (\text{Post}_t \times \text{Early}_t) \\
& + \beta_5 (\text{Post}_t \times \text{Wed, Sat}) + \beta_6 (\text{Early}_t \times \text{Wed, Sat}) \\
& + \beta_7 (\text{Wed, Sat} \times \text{Post}_t \times \text{Early}_t) + \gamma_i + \phi_t + \epsilon_{it}
\end{aligned} \tag{2.7}$$

if address = Odd

$$\begin{aligned}
\text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Early}_t + \beta_3 (\text{Thu, Sun}) + \beta_4 (\text{Post}_t \times \text{Early}_t) \\
& + \beta_5 (\text{Post}_t \times \text{Thu, Sun}) + \beta_6 (\text{Early}_t \times \text{Thu, Sun}) \\
& + \beta_7 (\text{Thu, Sun} \times \text{Post}_t \times \text{Early}_t) + \gamma_i + \phi_t + \epsilon_{it}
\end{aligned} \tag{2.8}$$

if address = Even




Figure 2.5: Average daily use on dry days

where the dependent variable is consumption of household i in hour t during dry summer days. $Early_t$ is a dummy for hours between 1 and 3 am, the hours in which people are asked to water. We can see whether odd and even houses are watering at the appropriate hour and appropriate days by looking at the coefficient on the triple interaction term. The results are shown in Table 2.5. Odd houses consume significantly more during the early hours after the program relative to other hours on Wednesdays and Saturdays, relative to other days. The effect is not significant for even houses. The results conform with the overall conclusion that odd houses are more responsive, but on average, there is no treatment effect.

We can also consider a simple model to assess whether households change the watering time, ignoring the effect of days. This can be done through a simple difference in difference model as follows:

$$\text{Consumption}_{it} = \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Early}_i + \beta_4 (\text{Post}_t \times \text{Early}_i) + \gamma_i + \phi_t + \epsilon_{it} \quad (2.9)$$

Table 2.6 shows the estimation results for all houses, as well as for odd and even houses. There appears to be no significant change in consumption in the appropriate watering hours after the program. Another way to model hourly consumption is to create hourly bins and see how consumption varies relative to the appropriate watering hours. Table 2.7 shows the results of

Table 2.5: Regression results for models with both day and hour

Dependent Variable: Model:	Water Consumption	
	Odd	Even
<i>Variables</i>		
Wed_Sat	0.0009 (0.5502)	
Post	-6.746*** (1.371)	-4.020*** (0.8137)
Early	5.576** (2.829)	-1.185 (2.417)
Wed_Sat × Post	2.544*** (0.6070)	
Wed_Sat × Early	-0.8168 (1.410)	
Post × Early	-1.868 (2.445)	3.573 (2.415)
Wed_Sat × Post × Early	3.544* (1.846)	
Thu_Sun		-0.2387 (0.9174)
Thu_Sun × Post		-1.664* (0.9353)
Thu_Sun × Early		-3.996** (1.903)
Thu_Sun × Post × Early		2.654 (2.132)
<i>Fixed-effects</i>		
Meter_ID	Yes	Yes
Month	Yes	Yes
<i>Fit statistics</i>		
Observations	2,202,949	2,063,375
R ²	0.14760	0.02734

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2.6: Regression Results

Dependent Variable: Model:	Water Consumption		
	All	Odd	Even
<i>Variables</i>			
Post	-5.807*** (0.7746)	-6.383*** (1.328)	-5.189*** (0.7208)
Early	1.789 (1.779)	5.154* (2.729)	-1.949 (2.201)
Post × Early	0.6741 (1.611)	-1.577 (2.349)	3.271 (2.161)
<i>Fixed-effects</i>			
Meter_ID	Yes	Yes	Yes
Month	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	5,760,865	2,970,686	2,790,179
R ²	0.05559	0.14242	0.03127

Clustered (Meter_ID) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

estimating Equation (2.9) with hourly bins. Each bin is a dummy for hours that fall within the specified bounds. For example, Hour_3_5 includes hours between 3 am and 5 am. The reference category is the bin for the assigned watering hours between 1 am and 3 am. If the program were effective, we would expect watering only on these hours, so consumption on the remaining hours would be lower compared to that in these hours.¹⁸ Households consume 11.8 gallons less between 5 am and 6 am, compared to their consumption during the assigned watering hours. As mentioned earlier, most sprinklers are set to turn on at 5 am, and this result suggests households are consuming significantly less in the usual watering hour compared to the assigned watering hours. However, if the program were fully effective in that no one watered between 5 and 6 am, we would expect to see a larger reduction. This finding conforms with Figure 2.4, where we still see the peak at 5 am but it is smaller compared to the previous year. Table 2.7 also shows households consuming significantly more between 6 am and 9 am compared to consumption in the assigned watering hours. This is perhaps due to heavy indoor water usage at that time. The hourly models also have

¹⁸This is assuming outdoor watering is the most water-intensive activity in a household.

poorer fit compared to the daily models. Overall, these results should be interpreted with caution as these do not incorporate the day of the week, which was a major element of the program.

Table 2.7: Regression Results for Hourly Effects

Dependent Variable: Model:	Water Consumption		
	All	Odd	Even
<i>Variables</i>			
Post	-2.446*	-2.166	-2.681
	(1.373)	(2.113)	(1.722)
Hour_3_5	19.38***	19.39***	19.37***
	(2.159)	(3.008)	(3.100)
Hour_5_6	33.12***	33.59***	32.60***
	(2.563)	(3.704)	(3.517)
Hour_13_17	-14.67***	-16.23***	-12.95***
	(1.051)	(1.424)	(1.549)
Hour_17_21	-2.624**	-2.681	-2.560
	(1.160)	(1.700)	(1.561)
Hour_21_0	5.914***	7.657***	3.978*
	(1.787)	(2.610)	(2.411)
Post × Hour_3_5	-6.716***	-6.737**	-6.694**
	(2.150)	(3.052)	(3.027)
Post × Hour_5_6	-11.85***	-12.15***	-11.50***
	(2.697)	(3.813)	(3.803)
Post × Hour_6_9	5.362***	4.864**	5.871***
	(1.487)	(2.247)	(1.943)
Post × Hour_9_13	-10.12***	-12.02***	-8.183***
	(1.239)	(1.873)	(1.615)
Post × Hour_13_17	1.451	1.394	1.373
	(1.338)	(1.836)	(1.962)
Post × Hour_17_21	-2.136	-3.509	-0.7362
	(1.468)	(2.153)	(1.985)
Post × Hour_21_1	-6.917***	-9.494***	-4.129
	(2.077)	(3.069)	(2.758)
<i>Fixed-effects</i>			
Meter_ID	Yes	Yes	Yes
Month	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	5,760,865	2,970,686	2,790,179
R ²	0.05950	0.15273	0.03340

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

2.3.2 Average Treatment Effect among Compliers

The results so far indicate a large degree of non-compliance among both odd and even households. That is, there are households in both groups that water on days other than their assigned days. Noncompliance can also manifest through households watering both on their assigned days and on days they are not supposed to. In the presence of non-compliance, the average treatment effect will be biased (Angrist et al., 1996). The issue of non-compliance has been explored in several fields, particularly medical research (Moerbeek and Schie, 2019). Consider a medical program that seeks to estimate the impact of taking a daily medicine. Those in the treated group are supposed to take the medicine daily while those in the control group are not. However, there could be many individuals in the treated group who do not follow these instructions, either taking the medicine intermittently or not taking the medicine at all. Calculating the difference between the outcomes of the treated and control groups would lead to biased estimates of the average treatment effect. The literature thus suggests estimating the average impact on those who comply with their treatment assignments (also referred to as local average treatment effect) (Bloom, 1984). Although this entails estimating the treatment effect on a subset of the treated group and may not be generalized to the broader population, it could offer some insight into what would have happened in an ideal scenario if everyone complied with their treatment status.

Given the large degree of non-compliance in this study, we now turn to estimating the average treatment effect among compliers. Since we have two sets of treated groups, odd and even households, we analyze each group separately. We begin by identifying compliers in each group, and then we estimate the impact on water consumption on assigned days.

To define compliance, we ideally want households to switch their watering days to their assigned days. This means that odd households should avoid watering on Wednesdays and Saturdays before the intervention and should water *only* on these days afterward. The same would apply for even houses on their assigned days. However, no households fit this ideal scenario. For instance, some odd households never watered on Wednesdays before the intervention but did water on Saturdays.

A potential workaround is to examine their watering frequency on their assigned days to see if they are watering more often on these days. This could help identify households that switched. Watering frequency is calculated as the ratio of the number of days they watered on their assigned days to the total number of assigned days. If this frequency is higher in the post-intervention period, the household can be considered to have switched.

To identify households that switched, we need to determine whether a household waters on any given day. This can be done by calculating the indoor water use for each household. A proxy for indoor use is the average winter consumption (AWC), which Denver Water calculates as the average consumption in January, February, and March. We calculate the daily AWC for each household in both the pre- and post-intervention periods. Then, we compare each household's daily water use in the summer months with their daily AWC. If their consumption exceeds the AWC, it indicates that the household watered that day. These steps allow us to identify the specific days of the week a household waters.

Using the above criteria, we find that 31% of odd houses and 26.5% of even houses switched. Having identified those that switched, we again utilize a triple difference model as follows:

$$\begin{aligned} \text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Switched}_i + \beta_3 (\text{Wed, Sat}) + \beta_4 (\text{Post}_t \times \text{Switched}_i) \\ & + \beta_5 (\text{Post}_t \times \text{Wed, Sat}) + \beta_6 (\text{Switched}_i \times \text{Wed, Sat}) \\ & + \beta_7 (\text{Post}_t \times \text{Switched}_i \times \text{Wed, Sat}) + \gamma_i + \phi_t + \epsilon_{it} \end{aligned} \quad (2.10)$$

if address = Odd

$$\begin{aligned} \text{Consumption}_{it} = & \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Switched}_i + \beta_3 (\text{Thu, Sun}) + \beta_4 (\text{Post}_t \times \text{Switched}_i) \\ & + \beta_5 (\text{Post}_t \times \text{Thu, Sun}) + \beta_6 (\text{Switched}_i \times \text{Thu, Sun}) \\ & + \beta_7 (\text{Post}_t \times \text{Switched}_i \times \text{Thu, Sun}) + \gamma_i + \phi_t + \epsilon_{it} \end{aligned} \quad (2.11)$$

if address = Even

In estimating equations (2.10) and (2.11), our coefficient of interest is β_7 which shows how much water those that switched consumed on their assigned days relative to those who did not switch. As shown in Table 2.8, odd houses that switched consumed 75 gallons more on Wednesdays and Saturdays after the program compared to before, as we would expect. Similarly, even houses that switched consumed 94 gallons more on Thursdays and Sundays after the program compared to before. However, one caveat is that this does not necessarily imply they are *only* watering on their assigned days. This analysis shows that those who switched consumed more on their assigned days relative to other days after the treatment compared to before. This also contrasts with our earlier finding that odd houses appear to be more responsive to the treatment as a group compared to even houses. Although odd houses appear to be more responsive to the program as a group, if we consider the subset of the population that complied or switched, then the effect is stronger for even houses. These results highlight the importance of a more nuanced approach in the presence of non-compliance.

Table 2.8 also provides insight into the behavior of those that did not switch. In each equation, β_5 shows the difference between water consumed on their assigned days and the other days after the program compared to before. Odd houses that did not switch consumed 44 gallons more on Wednesdays and Saturdays after the program compared to before. Recall that we defined switching as watering more frequently on assigned days, so even though odd houses that did not switch consume more water on their assigned days after the program, they are watering less frequently on their assigned days. In contrast, even houses that did not switch consumed 71 gallons less on Thursdays and Sundays after the program compared to before.

Table 2.8: Regression results for compliers

Dependent Variable: Model:	Water Consumption	
	Odd	Even
<i>Variables</i>		
Wed_Sat	24.00 (16.09)	
Post	-173.7*** (43.33)	-70.72*** (22.39)
Wed_Sat × Post	44.09*** (16.68)	
Wed_Sat × Switched	-88.60*** (28.83)	
Post × Switched	53.41 (49.91)	-45.90 (49.02)
Wed_Sat × Post × Switched	75.49** (35.92)	
Thu_Sun		30.83 (25.94)
Thu_Sun × Post		-70.87*** (26.21)
Thu_Sun × Switched		-183.6*** (44.21)
Thu_Sun × Post × Switched		94.46** (44.73)
<i>Fixed-effects</i>		
Meter_ID	Yes	Yes
Month	Yes	Yes
<i>Fit statistics</i>		
Observations	91,948	86,123
R ²	0.57388	0.34859
<i>Clustered (Meter_ID) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

2.3.3 Empirical Models with Logarithmic Transformation of Consumption

So far, we have explored empirical specifications where the dependent variable was total hourly or daily consumption. Although many relevant studies include the consumption variable in levels (Brent et al., 2017; Stitzel and Rogers, 2023; Bhanot, 2021; Ferraro and Price, 2013), other papers like Lurbe et al. (2023) employ the logarithmic transformation of consumption. Log-transforming water consumption is often convenient as it allows one to interpret the coefficient of interest as a percentage change. However, such a transformation helps stabilize variance and would dampen the effect of very high water users (Brent et al., 2017). Nonetheless, we estimate our base models with log-transformation of daily water consumption and compare results with earlier models.

Table 2.9 shows the results of estimating equations (2.1) and (2.2) with log-transformation of consumption. The results are unchanged in that there appears to be no significant treatment effect. The variable Post is significant in both regressions and describes the percentage changes in consumption in line with the visual evidence in Figure 2.4 that odd houses are more responsive. Odd houses reduce their consumption on Thursdays and Sundays by 23%, whereas even houses reduce their consumption on Wednesdays and Saturdays by 6.8%. However, these models have poorer fit compared to the linear (columns 1 and 2 in Table 2.2). We next estimate the triple difference models (equations 2.3 and 2.4) with and without weather controls. As shown in Table 2.10, the results are unchanged in the first two columns which shows the specifications in which Odd is the treated group. However, unlike the previous results, in the specifications in which even houses are the treated group and odd houses are the control group, there is a significant treatment effect. Even houses consume 7.9% more on Thursdays and Sundays after the program compared to before. Adding weather controls moderate the treatment effect, but it remains significant for even houses. Columns (1) and (3) also show better model fits compared to columns (1) and (3) on Table 2.3.

Lastly, we estimate the treatment effect among compliers (equations 2.10 and 2.11) using log-transformed water consumption, with results shown in Table 2.11. Compared to the linear model which showed a stronger effect for even houses, here the effect is stronger for odd houses. Odd houses that switched consumed 24% more on Wednesdays and Saturdays after the program compared to before, whereas the effect is 18.8% for even houses that switched. The log model for even houses has a better fit compared to the linear model for even houses. Overall, the log transformation seems to narrow the gap in model fit between regressions with odd houses and regressions with even houses.

Table 2.9: Regression results

Dependent Variable:	Log Water Consumption	
Model:	(1)	(2)
Post	-0.0685*	-0.2370***
	(0.0387)	(0.0336)
Post × Odd	0.0214	
	(0.0517)	
Post × Even		0.0368
		(0.0503)
<i>Fixed-effects</i>		
Meter_ID	Yes	Yes
Month	Yes	Yes
<i>Fit statistics</i>		
Observations	48,896	46,761
R ²	0.49961	0.51950

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2.10: Regression results

Dependent Variable: Model:	Log Water Consumption			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
Post	-0.1664*** (0.0335)	-0.1782*** (0.0302)	-0.1932*** (0.0298)	-0.1811*** (0.0280)
Wed_Sat	-0.0590*** (0.0200)	-0.0307* (0.0169)		
Avg_Temp		0.0250*** (0.0008)		0.0237*** (0.0008)
Avg_Precip		-0.6957*** (0.0398)		-0.6471*** (0.0393)
Post × Odd	-0.0234 (0.0451)	-0.0094 (0.0415)		-0.0050 (0.0416)
Post × Wed_Sat	0.0917*** (0.0234)	0.0128 (0.0185)		
Odd × Wed_Sat	0.0230 (0.0273)	0.0298 (0.0237)		
Post × Odd × Wed_Sat	0.0394 (0.0326)	0.0244 (0.0260)		
Thu_Sun			-0.0550** (0.0240)	-0.0736*** (0.0237)
Post × Even			-0.0020 (0.0448)	-0.0050 (0.0416)
Post × Thu_Sun			-0.0585** (0.0256)	-0.0166 (0.0235)
Even × Thu_Sun			-0.0263 (0.0364)	-0.0168 (0.0350)
Post × Even × Thu_Sun			0.0790** (0.0385)	0.0598* (0.0347)
<i>Fixed-effects</i>				
Meter_ID	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	123,022	253,315	120,887	254,051
R ²	0.46703	0.44171	0.46174	0.43999

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2.11: Regression results

Dependent Variable: Model:	Log Water Consumption	
	(1)	(2)
<i>Variables</i>		
Wed_Sat	0.0511** (0.0201)	
Post	-0.2216*** (0.0345)	-0.1213*** (0.0393)
Wed_Sat × Post	0.0975*** (0.0245)	
Wed_Sat × Switched	-0.2403*** (0.0499)	
Post × Switched	0.0207 (0.0606)	-0.0521 (0.0710)
Wed_Sat × Post × Switched	0.2437*** (0.0659)	
Thu_Sun		0.0336 (0.0290)
Thu_Sun × Post		-0.0933*** (0.0316)
Thu_Sun × Switched		-0.3702*** (0.0671)
Thu_Sun × Post × Switched		0.1880*** (0.0703)
<i>Fixed-effects</i>		
Meter_ID	Yes	Yes
Month	Yes	Yes
<i>Fit statistics</i>		
Observations	88,186	81,597
R ²	0.46242	0.44499
Within R ²	0.00645	0.00522

Clustered (Meter_ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

2.4 Conclusion

As utilities grapple with the challenge of managing scarce resources and addressing peak demand, non-price interventions have emerged as a promising strategy to influence residential water consumption behavior. In this study, we assessed the impact of a nonrandomized intervention designed to alter outdoor watering schedules in Denver, Colorado. The intervention aimed to address peak demand by providing households with information about the challenges associated with water use during peak periods and requesting that they adjust their watering days and times according to their addresses. This approach can be characterized as a pro-social appeal intended to promote more efficient water use. Our analysis reveals that, despite the intervention's intent, the overall water consumption did not significantly change in the way the intervention had hoped. While there was a noticeable reduction in water consumption following the intervention, this decrease appears to be largely attributed to an increase in precipitation during the same period rather than the effectiveness of the intervention itself. This complicates the evaluation of the intervention's impact, as it masks the true effect of the behavioral change induced by the program.

The study encountered several methodological challenges inherent to evaluating nonrandomized interventions. One of the primary issues was the lack of a control group, which made it difficult to isolate the specific effects of the intervention from other factors influencing water consumption. To address this, we employed alternative analytical strategies, such as comparing the behavior of odd-numbered households on their assigned days with even-numbered households on the same days. This approach was intended to provide a quasi-experimental framework within the nonrandomized context. Despite these efforts, the analysis yielded weak evidence of a significant change in watering times as a result of the intervention.

The results indicate that households did not substantially alter their watering times as requested. While there was some indication that a subset of households may have responded to the intervention by changing their watering behaviors, the evidence was not robust enough to conclude a significant average treatment effect. Specifically, our analysis showed that while some households might have adjusted their behavior, the overall effect of the intervention on water consumption

patterns was not statistically significant. This suggests that, on a broad scale, the intervention did not achieve its intended impact of reducing peak demand through altered watering schedules.

These findings underscore several important lessons for the design and evaluation of similar non-price interventions. First, the challenge of isolating the treatment effect in the absence of a control group highlights the need for carefully designed experimental or quasi-experimental frameworks when evaluating the effectiveness of such interventions. Future studies should consider incorporating control groups or using other methodological approaches to better isolate and measure the impact of the intervention.

Second, the results emphasize the importance of addressing non-compliance and ensuring that the intervention is effectively communicated and implemented. In this case, the limited changes in watering behavior suggest that the pro-social appeal of the intervention may not have been sufficiently compelling or actionable for a significant portion of the population. To enhance the effectiveness of future programs, it may be beneficial to explore additional strategies, such as targeted incentives or more direct enforcement mechanisms.

Finally, the study highlights the broader challenge of evaluating nonrandomized treatments. Even when an intervention has the potential to impact consumption behavior, the precise estimation of its effects can be complicated by external factors, such as weather variations, and by the limitations inherent in the study design. As such, while nonrandomized interventions can play a significant role in resource management, careful consideration and methodological rigor are essential to accurately assess their effectiveness.

In conclusion, while the intervention aimed at altering residential water consumption in Denver did not demonstrate a clear average treatment effect, the study provides insights into the complexities of evaluating nonrandomized interventions. The findings underscore the importance of robust study designs, effective communication, and adherence strategies in achieving the desired behavioral changes. Future research should build on these insights to refine and improve the design of similar programs, ensuring that they are better equipped to achieve their objectives and contribute to more effective resource management.

Chapter 3

The Impact of CAFOs on Ammonia and Methane Concentrations

3.1 Introduction

A substantial shift in US agriculture has led to a rapid rise in the number of concentrated animal feeding operations (CAFOs), which grew more than 16% between 2011 and 2022, reaching a total of 21,539 establishments by 2022 (EPA, 2023).¹⁹ CAFOs have more adverse environmental impacts relative to smaller facilities (EPA, 2022; Copeland, 2010). Much of the impacts come from the large quantities of animal waste produced from these facilities, improper management of which can lead to environmental degradation (EPA, 2001). Examples include runoff of manure pollutants into surface and groundwater resulting in increased algae production, fish kills, reduced recreational opportunities and contaminated drinking water (EPA, 1997). CAFOs also contribute to air pollution by emitting pollutants like ammonia (NH₃), which degrade air quality, and potent greenhouse gases, like methane (CH₄), which lead to climate change (Golston et al., 2020). Although CAFO discharges to water are regulated under the Clean Water Act, air pollution regulation is less stringent.²⁰ One of the reasons behind this state of air pollution regulation is that the EPA has yet to finalize methods to reliably estimate emissions from CAFOs, which would allow the agency to determine whether CAFOs are in compliance with applicable Clean Air Act requirements.²¹ This paper provides an estimate of CAFO-related ammonia and methane contributions to

¹⁹The Environmental Protection Agency (EPA) defines CAFOs as agricultural operations where animals are raised and fed in confined spaces for at least 45 days in a year and with no sustained vegetation growth during the normal growing season within the facility. CAFOs are further divided into small, medium, and large categories based on the number of animals.

²⁰Two laws requiring facilities to report air emissions have been amended in recent years to exempt CAFOs (Missouri Independent, 2023)

²¹Draft Air Emissions Estimating Methodologies for Animal Feeding Operations, EPA

atmospheric concentrations at a broader geographic scale than previous literature and contributes the policy dialog on this topic.

CAFOs emit pollutants through two primary channels - livestock manure and the digestive process of ruminants (cows, sheep, and goats) (Clark et al., 2020). Ammonia is among the primary gases associated with manure decomposition, resulting in strong odors and chronic respiratory illness.²² The health impacts of agricultural ammonia emissions have been widely studied. Agricultural ammonia emissions in the U.S. are estimated to result in over 12,000 premature deaths annually, along with societal costs totaling approximately \$160 billion per year, primarily due to the formation of fine particulate matter (PM_{2.5})²³ (Domingo et al., 2021; Heo et al., 2016). Mirabelli et al. (2005) find that children in schools near CAFOs are more likely to have asthma. These and other studies demonstrate the substantial damages from CAFO ammonia emissions, primarily originating from manure decomposition. While livestock manure can emit methane as well, livestock digestion is the major source of CAFO methane emissions (Fisher 2022). The UN Environmental Programme estimates that livestock contribute 32% of global methane emissions (UNEP, 2021). Although methane is a short-lived climate pollutant with an atmospheric lifetime of roughly a decade, it is a more potent greenhouse gas compared to carbon dioxide; according to the Intergovernmental Panel on Climate Change, one ton of methane can be considered to be equivalent to over 80 tons of carbon dioxide over a 20-year period.²⁴ Methane contributes to the formation of ground-level ozone with wide-ranging impacts, including premature deaths, ecosystem damages, and diminished agricultural productivity (UNEP, 2021). Methane is a major component of agricultural emissions, and if these emissions persist at current levels, the Paris Climate Agreement targets would remain out of reach even if fossil fuel emissions were completely halted (Clark et al., 2020).

The salience of CAFO air pollutants to environmental and climate outcomes makes it necessary to reliably estimate emissions. This paper aims to measure changes in ammonia and methane

²²CDC, Ammonia: Exposure, Decontamination, Treatment

²³Ammonia is a key precursor to the formation of PM 2.5 through atmospheric reactions.

²⁴Methane and Climate Change, IEA

contributions at the surface and in the atmospheric column respectively near and around CAFOs in the contiguous US. Our first contribution is that we assemble a mostly complete dataset of CAFOs at the national level, as opposed to existing studies that are mostly specific to farms, regions, or industries (such as hog production). A major challenge in CAFO-related research is the absence of consistent and publicly available data on the exact locations of CAFOs and their characteristics, such as animal type and population (Burns et al., 2023). This is the gap we are attempting to fill.

The second contribution is in adding to our understanding of CAFO emissions. Not only is there no consensus on what emissions actually are, there is also no national consensus on how best to measure these emissions, as acknowledged by the EPA.²⁵ This paper adds to the scarce economic literature on CAFO emissions and introduces a data point to the existing estimates.

The economics literature on this topic is relatively scarce. Some of the papers investigate the effects of ammonia discharge on ambient water quality, as CAFO discharges to water are regulated under the Clean Water Act.²⁶ According to Raff and Meyer (2021), the introduction of an additional CAFO in Wisconsin results in a 2.7% increase in surface water ammonia levels, with the median household estimated to be willing to pay up to \$12 per year to avoid this damage. The interaction of water quality regulation with ambient air quality has also been explored. Sneeringer (2010) investigates how water regulations, which requires CAFOs to store manure in open-air ponds, can increase ammonia emissions from hog production. In terms of the socioeconomic impacts, Sneeringer (2009) finds that a doubling of production in CAFOs leads to a 7.4% increase in infant mortality at the national level, and that this effect comes from air pollution rather than water.

Quantifying CAFO emissions has been largely explored in the environmental science literature (Todd et al, 2008; Pollack et al., 2022; Burns et al., 2023). There are four general ways of measuring pollutants from CAFOs, each with benefits and drawbacks. On-the-ground monitors can offer more precise measurements if located on CAFO premises or nearby. However, a monitor

²⁵Improving Air Quality, IEA

²⁶In contrast, two laws requiring facilities to report air emissions have been amended in recent years to exempt CAFOs, and although the Clean Air Act still applies to livestock operations, the EPA is yet to finalize reliable pollutant estimation methods. (Missouri Independent, 2023)

located near multiple CAFOs would pick up aggregate emissions and make it difficult to isolate emissions from individual CAFOs, as well as being susceptible to other sources of the pollutant. Another disadvantage is that monitors are on the ground whereas plumes are dispersed throughout the atmosphere, thereby understating the extent of pollution. The second way of measurement is through satellite observations. While satellite observations can offer more frequent data points at the national level, relatively large grid sizes make the measurements imprecise. The drawbacks of on-the-ground monitors and satellite observations can be somewhat addressed by the third way of measurement which is aircraft flyovers. Although very expensive, research aircrafts provide precise measurements of multiple pollutants on site. However, they provide few data points making them appropriate for a local analysis of a few farms. The fourth way of measurement, particularly for methane, is through wearable devices equipped with sensors that are fitted over an animal's snout. These devices can provide good estimates of methane emitted from the digestive process of ruminants. However, they do not capture methane and ammonia emissions from manure decomposition and are thus difficult to scale up.

CAFOs represent a relatively undersampled source of air pollutants (Eilerman et al., 2016), and the various measurement techniques lead to emission estimates that are often highly uncertain (Hayek and Miller, 2021; Vechi et al., 2023). Hayek and Miller (2021) highlight the discrepancy between "top-down" methods, such as aerial surveillance and satellite observations, and "bottom-up" methods, which entail collecting information on animal and farms and aggregating to calculate total emissions. Their top-down estimates of methane emissions in the US are 39%-90% higher than predictions from bottom-up models. Similar discrepancies are reported by Vichi et al. (2023), who find ammonia and methane emissions from dairy CAFOs in California were 28% and 60% higher, respectively, when compared to inventory models which are based on the number of animals. However, top-down approaches such as aircraft flyovers are also susceptible to variable estimates. Golston et al. (2020) show that ammonia and methane measurements depend on time of day, highlighting the need to account for biases in daytime-only measurements, as is common with most aircraft observations. They also demonstrate large subfarm, site-to-site, and diurnal variability.

ties among CAFOs in northeastern Colorado, underscoring the importance of taking measurements across these scales to obtain representative emission factors.

Given the difficulty in measuring CAFO emissions, and the variable estimates produced by the abovementioned methods, there is no consensus on what the emissions actually are. In an ideal situation, CAFOs would be enclosed and fitted with monitors that would measure pollutants before they are released outside, similar to power plants that have monitors installed within their smokestacks. In the absence of this ideal scenario, the best approach might be to conduct a meta-analysis by deriving a large number of estimates and calculating the average. This paper contributes new estimates of CAFO ammonia and methane emissions, paving the way for a future meta-analysis.

We separately analyze ammonia and methane emissions using two of the four mentioned methods. Ammonia levels are measured using ground-based ammonia monitors and by taking advantage of wind direction. This combines traditional measurement approaches with causal identification techniques from econometrics. Despite being limited by the presence of very few monitors relative to the number of CAFOs, we find positive associations between wind (blowing from CAFOs to monitors) and measured ammonia levels. However, this association is only observed for CAFOs within 10 km of a monitor, with no significant correlation for CAFOs located further away. For methane, we use satellite observations that cover all CAFOs in the dataset, although the large grid size (5.5 km by 7 km) hampers precise estimation. Nevertheless, we find a small but significantly positive association between CAFOs and methane levels, with each additional CAFO increasing atmospheric methane concentrations by 0.04%.

The rest of the paper proceeds as follows: Section 2 outlines the identification strategy and the data. Section 3 highlights the main findings, and Section 4 encompasses the discussion, conclusion, and suggestions for future research.

3.2 Data and Identification Strategy

3.2.1 Data

Our analysis requires data on CAFO locations and ammonia levels in the vicinity. We first obtain CAFO information from EPA's Integrated Compliance Information System (ICIS), which contains latitude and longitude coordinates, animal count and type, land size of the facility, and other facility characteristics. However, this dataset does not cover all states, so we supplement it with CAFO information from individual states at the facility level. Some states have facility-level data for multiple years while others are cross-sectional. Although the data content varies from state to state, they all contain facility coordinates and addresses. To avoid double counting, we only retain EPA data for states for which separate data were not obtained.²⁷ We compile the state level data to assemble a novel dataset of 37,841 CAFOs across much of the contiguous US, as shown in Figure 3.1. Most CAFOs are of a single type (beef, dairy, hog, or poultry), although many are mixed as well. There is also substantial variation in CAFO type and density across states. We do not have data from all states. Notably, for example, North and South Dakota are missing or incomplete. The total number of CAFOs differs from the EPA figure of 21,539 establishments for several reasons. The EPA does not collect data from all states, and even within the states they do survey, their list is not comprehensive. Moreover, there are differences in the definition of CAFOs by states and the EPA. The EPA defines large CAFOs solely based on the number of animals, and medium CAFOs based on both animal numbers and whether the operations discharges into a water body. However, whether an operation discharges into water is determined by a state-level permitting authority. Some states like Georgia and Kentucky require a CAFO permit only if the CAFO is discharging, regardless of the size. So CAFOs that meet the EPA criteria for a large CAFO may not be considered a CAFO by some states if they do not discharge. Moreover, state-

²⁷To obtain state-level data, we gathered information from several sources. First, we obtained some data from the University of Iowa's cafomaps.org website. Second, we emailed state officials, usually in the water quality divisions as they collect location information for water discharge permitting, to request CAFO location information. Third, some states host this data on public websites. We attempted to obtain data from all states in the contiguous US, however some states either did not collect the data or we could not find a person responsible for collecting and maintaining the data.

level permitting authorities can designate operations as CAFOs regardless of their size if they are considered to be significant contributor of pollutants. This criteria can also vary from state to state. The EPA uses actual number of animals to define CAFO sizes, whereas many states still use animal units.²⁸ However, the state thresholds for animal units are mostly aligned with the EPA definition.

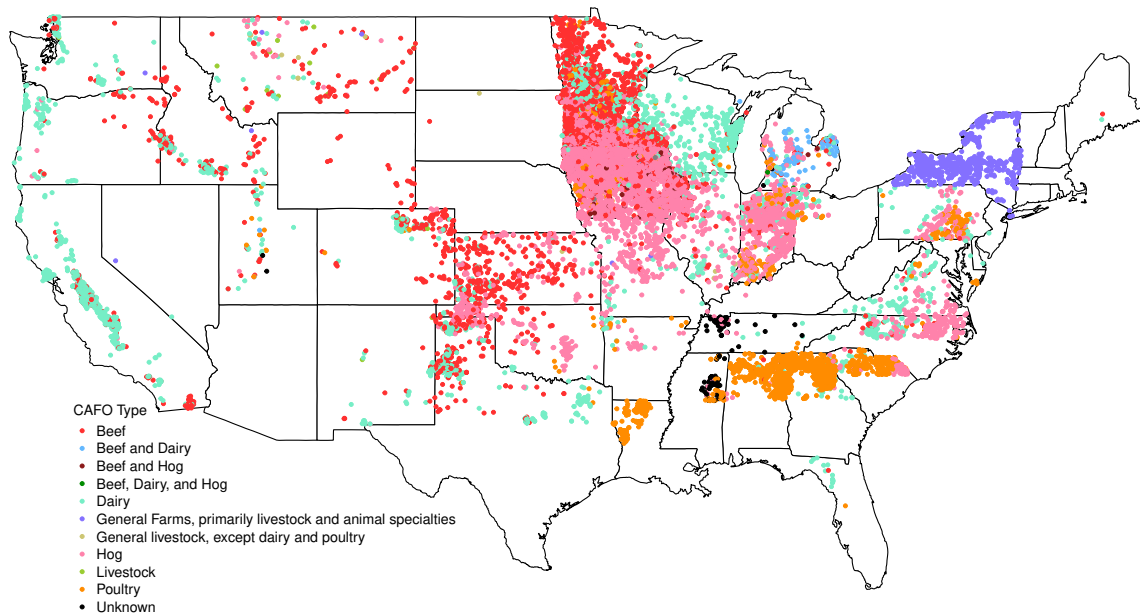


Figure 3.1: CAFO locations and types

We obtain ammonia levels from the Ammonia Monitoring Network (AMN) under the National Atmospheric Deposition Program. The AMN provides total ammonia concentrations, measured in micrograms per cubic meter ($\mu\text{g}/\text{m}^3$), every 14 days from 129 sites across 34 states. We exclude observations with an invalid quality rating. Monitors were installed at different times, resulting in an unbalanced panel of ammonia concentrations. The time frame of the ammonia component of this study is from the beginning of monitoring at each site (with the earliest year being 2007) to 2023. Since weather conditions influence ammonia concentrations, we also collect data on

²⁸An animal unit (AU) is a standard unit of measurement used to compare different types of livestock based on their feed consumption and waste production.

precipitation, humidity, solar radiation, wind speed, and wind direction from GridMet. Table 3.1 shows the summary statistics of ammonia and weather variables in the dataset.

Table 3.1: Summary Statistics

Variable	Unit	Observations	Mean	SD	Min	Max
NH_3	$\mu g/m^3$	11635	1.49	2.22	0.00	48.13
pet	mm	11635	3.11	1.74	0.16	10.63
pr	mm	11635	1.68	1.55	0.00	21.52
rmax	%	11635	85.16	12.02	19.27	100.00
rmin	%	11635	43.91	11.46	4.19	82.59
srad	W/m^2	11635	182.90	78.08	30.49	360.03
tmmx	$^{\circ}F$	11635	62.30	18.66	-0.99	106.26
vs	m/s	11635	9.28	2.42	2.36	19.03

Note: pet is reference grass evapotranspiration; pr is precipitation; rmax and rmin are maximum and minimum near-surface relative humidity; tmmx is maximum near-surface air temperature; vs is wind speed

Given that ammonia is a local pollutant and does not travel far, we subset the CAFO data to include only CAFOs within 100 km of a monitor. Out of 129 monitors, only 68 have nearby CAFOs. This results in 6116 CAFOs for which we can quantify ammonia emissions.²⁹ Figure 3.2 shows the location of the monitors and nearby CAFOs. The lack of monitors causes many CAFOs in states like Minnesota and Iowa, which have a large number of these facilities, to be excluded.

Figure 3.3 illustrates the average monthly concentrations of ammonia for monitors located near and away from CAFOs. The data exhibits pronounced seasonality, with peak levels typically occurring during the summer months. According to Burns et al. (2023), this seasonal variation is attributed to favorable conditions in summer that enhance the dispersion of ammonia. Monitors near CAFOs consistently record higher ammonia levels compared to those situated farther away.

In contrast to ammonia concentrations, which are collected over a 14-day period, the methane concentrations for this study are obtained daily from satellite flyovers. We retrieve bias-corrected methane mixing ratios (in parts per billion) from the Sentinel-5P TROPOMI satellite, using grids

²⁹Note, in future work, we will obtain satellite based measurements of ammonia concentrations and compare the estimates with the monitor estimates.




Figure 3.2: Ammonia monitors and nearby CAFOs




Figure 3.3: Mean Ammonia Concentrations

of 5.5 km by 7 km, for the year 2023. The satellite data can be affected by noise due to cloud cover and other factors. To ensure the highest quality data, we retain only those observations with a quality assurance value greater than 0.5, in accordance with the manual. The satellite data covers the entire US, so we are able to retain all 37,841 CAFOs in the original dataset. We then calculate the number of CAFOs in each grid. 6.7% of grids have one or more CAFOs, and among the grids that have CAFOs, the average number of CAFOs is 3.2. Similar to the case of ammonia, grids

with CAFOs have higher methane concentrations on average for almost all months (Figure 3.4). We also incorporate GridMet data as weather conditions can affect methane levels.

Table 3.2: Summary Statistics (at the grid cell level)

Variable	Observations	Mean	SD	Min	Max
Methane ratio	4856927	1905.61	22.09	1371.16	2658.37
Number of CAFOs	4856927	0.22	1.26	0.00	38.00
pet	4856927	3.71	2.14	0.00	15.40
pr	4856927	0.01	0.08	0.00	7.56
rmax	4856927	77.96	17.97	8.60	100.00
rmin	4856927	27.20	13.51	1.00	100.00
srad	4856927	193.11	68.96	0.00	415.60
tmmx	4856927	67.70	19.78	-14.71	128.03
vs	4856927	8.38	3.25	0.00	40.71

Figures 3.5 and 3.6 show average ammonia and methane levels respectively. Most of the monitors record ammonia levels that are in the lower end of the range. However, there are a few sites, particularly in northern Colorado, where higher levels of ammonia are recorded. In contrast, methane levels are in the middle of the range throughout much of the US, while higher levels are observed in the southeastern region and parts of California.




Figure 3.4: Mean Methane Mixing Ratio

Note: There are no observations in May and June with a quality assurance value greater than 0.5.




Figure 3.5: Average ammonia levels at monitor sites near CAFOs




Figure 3.6: Average methane levels

3.2.2 Identification Strategy and Empirical Specifications

As mentioned in the previous section, there are 68 ammonia monitors that have CAFOs within 100 km. These monitors record ammonia from CAFOs and other sources accumulated over 14 days. In order to identify CAFO-related ammonia emissions, we exploit wind direction. We calculate the fraction of days within the 14-day period during which the wind blows from CAFOs to monitors. Wind direction is categorized into 10 degree wedges. By calculating the fraction of days the wind blows from CAFOs to monitors, we can isolate ammonia levels when the wind blows from CAFOs relative to ammonia levels when the wind does not blow from CAFOs, or ambient ammonia levels. However, this method cannot differentiate between ammonia emissions originating from individual CAFOs when multiple CAFOs are situated in the same wind direction, nor can it identify CAFO specific ammonia emissions as emissions disperse in proportion to the distance from the source. We account for the potential diminishing effect of distance on ammonia concentrations by creating fractions of CAFOs within concentric radii of 10 km, 10-20 km, 20-30 km, and beyond 30 km from the monitors.

We estimate the impact of fraction days wind blows toward monitor i in period t on ammonia concentrations using the following specification:

$$\begin{aligned}
 NH3_{it} = & \beta_0 + \beta_1 \times (FractionDays \leq 10km)_{it} \\
 & + \beta_2 \times (FractionDays > 10 \text{ and } \leq 20km)_{it} \\
 & + \beta_3 \times (FractionDays > 20 \text{ and } \leq 30km)_{it} \\
 & + \beta_4 \times (FractionDays > 30 \text{ and } \leq 100km)_{it} \\
 & + \beta_5 \times W_{it} + \gamma_{imy} + \varepsilon_{it}
 \end{aligned} \tag{3.1}$$

Here, each *FractionDays* variable represents the fraction of days the wind blows from CAFOs to monitors within a specified radius within a 14-day period. W_{it} is a vector of weather controls, and γ_{imy} denotes a monitor-month-year fixed effect. This serves as our base model, which we adjust to examine other aspects of CAFO emissions, such as wind blowing from specific types

of CAFOs. We hypothesize that β_1 will be positive, with the subsequent coefficients decreasing, reflecting the potential diminishing effect of distance.

In our approach to methane, we take a different route by focusing on the variation in grid cells that either have or do not have CAFOs. Despite only a small fraction of grids containing CAFOs, this method enables us to estimate the association between CAFOs and atmospheric methane concentrations. Potential confounders include whether a grid is within regions of oil and gas extraction, which are significant sources of methane (Pollack et al., 2022). The environmental science literature addresses the challenge of isolating CAFO emissions from other sources, such as oil and gas activities. McCabe et al. (2023) propose two methods: a subtraction method, which uses ammonia as a tracer to identify a CAFO plume before employing an ethane-to-methane ratio to exclude oil and gas methane emissions, and a multivariate regression with independent gas tracers to attribute emissions to different sectors. These methods are suitable for local-level analysis. McCabe et al. (2023) focus on Northern Colorado using aircraft flyover data. In contrast, our analysis is national in scope with a substantial number of observations. Therefore, we employ the following specification to estimate the effect of number of CAFOs on methane levels in grid i during day t

$$\begin{aligned}
 CH4_{it} = & \beta_0 + \beta_1 \times \text{Number of CAFOs}_{it} \\
 & + \beta_2 \times \text{Shale}_i + \beta_3 \times \text{Season}_t \\
 & + \beta_4 \times W_{it} + \gamma_{sm} + \varepsilon_{it}
 \end{aligned} \tag{3.2}$$

where $Shale_i$ indicates whether grid i is located in a shale play, and $Season_t$ is a categorical variable for season. W_{it} is a vector of weather controls, and γ_{sm} denotes state-month fixed effects. Given that methane is a global pollutant, the methane levels recorded in a grid containing CAFOs may originate from sources outside that grid. However, because we include numerous grids without CAFOs, they provide an adequate control for background methane concentrations. Likewise, we account for whether a grid is located in a shale region, with grids outside of shale areas serving as controls. We hypothesize that both β_1 and β_2 will be positive, as an increase in the number of CAFOs and being located in a shale region are likely to increase methane concentrations.

3.3 Results

We first present and discuss the results from the ammonia component of this study. Table 3.3 shows the estimates of Equation (3.1) with site-year-month and site-year-season fixed effects.

Table 3.3: Results for the Base Regression Model

Dependent Variable:	Ammonia concentration	
Model:	(1)	(2)
<i>Variables</i>		
frac10	1.447** (0.6040)	1.362*** (0.4466)
frac1020	-0.4026 (0.3331)	-0.4178* (0.2446)
frac2030	0.1490 (0.2684)	0.0623 (0.2364)
frac30100	0.0866 (0.0724)	0.1219* (0.0691)
pet	0.2896*** (0.0984)	0.4338*** (0.0712)
pr	-0.0057 (0.0106)	-0.0220** (0.0087)
rmax	0.0080*** (0.0028)	0.0084*** (0.0024)
rmin	-0.0088*** (0.0029)	-0.0092*** (0.0025)
srad	-0.0042*** (0.0014)	-0.0052*** (0.0009)
tmmx	0.0177*** (0.0036)	0.0026 (0.0031)
vs	-0.0237* (0.0128)	-0.0300*** (0.0073)
<i>Fixed-Effects</i>		
site-year-month	Yes	No
site-year-season	No	Yes
<i>Fit statistics</i>		
Observations	11,635	11,635
R ²	0.90618	0.83199

Clustered standard-errors in parentheses

Signif Codes: ***: 0.01, **: 0.05, *: 0.1

Note: pet is reference grass evapotranspiration; pr is precipitation; rmax and rmin are maximum and minimum near-surface relative humidity; tmmx is maximum near-surface air temperature; vs is wind speed

The fraction of days that the wind blows to monitors from CAFOs within 10 km is positively associated with ammonia levels at the monitor. Given that ammonia levels are recorded over 14-day periods, the coefficients for each radii can be divided by 14 to obtain the impact of an additional day of wind blowing from a CAFO to a monitor. Thus, for CAFOs within 10 km, one additional day of wind blowing from a CAFO to a monitor increases ammonia concentrations by $0.103 \mu\text{g}/\text{m}^3$ (6.7% increase above the mean) when estimating Equation (3.1) with site-year-month fixed effects. However, there is no significant association between ammonia concentrations and wind blowing from CAFOs beyond 10 km. This aligns with the notion that ammonia does not travel far, so we might not expect to see the monitors record ammonia from further away. Table 3.3 also illustrates how weather variables such as temperature affect ammonia levels.

We next explore whether wind speed affects ammonia levels. A higher wind speed could cause the ammonia to disperse more quickly, enter the upper atmosphere, or carry ammonia from CAFOs away from the monitors, thereby *reducing* the ammonia concentrations measured by the monitors. Conversely, stronger winds could transport ammonia from more distant CAFOs to the monitors, *increasing* the ammonia concentrations measured by further monitors. We assess the impact of wind by interacting wind speed with the various radii. Table 3.5 shows the results with the wind interactions for the models with site-year-month fixed effects. Including the interaction term substantially increases the coefficient on frac10, although the interaction somewhat diminishes the effect. This is consistent with the idea that as wind speed increases, it reduces the impact of wind blowing from CAFOs within 10 km on ammonia measured at the monitors. Columns 2 through 5 show the results for CAFO type. Once we include the wind interaction term, the fraction of days the wind blows from beef CAFOs becomes significant, unlike in the previous models. The coefficients for the other CAFO types increase as well and the interaction diminishes the impact somewhat, but the results are significant for beef and dairy only. Overall, including the wind interaction term with the various radii makes the coefficients on each radii positive in line with our a priori expectations. Given that our identification strategy exploits wind direction, incorporating the interaction with wind speed (which could affect how far ammonia travels) appears to be a bet-

Table 3.4: Results for Regressions for CAFO types

Dependent Variable: Model:	Ammonia Concentration			
	Beef	Dairy	Poultry	Hogs
<i>Variables</i>				
frac10	-15.66 (16.04)	2.960* (1.559)	10.04** (5.084)	0.1790 (0.1531)
frac1020	1.208 (1.148)	-1.068* (0.6233)	-0.6791 (1.268)	0.4385 (0.6024)
frac2030	0.7129 (0.9177)	-0.5719 (0.4780)	-0.2038 (0.7201)	0.4869 (0.3119)
frac30100	0.1984 (0.1437)	0.1371 (0.1118)	0.0079 (0.1829)	-0.0364 (0.1052)
pet	0.2978*** (0.1003)	0.3062*** (0.1007)	0.2921*** (0.1001)	0.2971*** (0.1002)
pr	-0.0060 (0.0113)	-0.0037 (0.0113)	-0.0074 (0.0110)	-0.0058 (0.0112)
rmax	0.0081*** (0.0028)	0.0083*** (0.0028)	0.0081*** (0.0028)	0.0081*** (0.0028)
rmin	-0.0086*** (0.0030)	-0.0089*** (0.0030)	-0.0086*** (0.0030)	-0.0087*** (0.0030)
srad	-0.0042*** (0.0015)	-0.0043*** (0.0015)	-0.0043*** (0.0015)	-0.0043*** (0.0015)
tmmx	0.0170*** (0.0037)	0.0170*** (0.0038)	0.0180*** (0.0037)	0.0177*** (0.0037)
vs	-0.0244* (0.0131)	-0.0243* (0.0131)	-0.0226* (0.0129)	-0.0243* (0.0131)
<i>Fixed-effects</i>				
sym	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	11,635	11,635	11,635	11,635
R ²	0.90605	0.90606	0.90627	0.90576

Clustered (sym) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

ter way to estimate the impact of CAFO on ammonia levels. Importantly, our model assumes a linear relationship between CAFOs and ammonia concentrations as well as the how wind speed moderates this relationship. A more accurate way of modeling this relationship would allow for nonlinear relationships that could account for more complex dispersion patterns. This will be explored in future work.

Table 3.5: Results for regressions with wind interaction

Dependent Variable:	Ammonia Concentration				
Model:	All	Beef	Dairy	Poultry	Hogs
<i>Variables</i>					
frac10	8.198** (3.620)	217.7* (122.1)	29.09** (14.40)	24.16 (16.26)	0.2460 (0.8438)
frac1020	0.2250 (1.414)	10.51 (7.037)	-0.9101 (2.842)	-8.612 (5.830)	3.289 (2.662)
frac2030	1.013 (1.252)	1.070 (4.741)	-1.287 (1.706)	1.117 (2.589)	1.377 (1.156)
frac30100	-0.5200 (0.3975)	0.0690 (0.5503)	-0.0116 (0.4600)	0.5315 (0.6548)	-0.6875* (0.0370)
pet	0.2846*** (0.0982)	0.2934*** (0.1021)	0.3033*** (0.1017)	0.2930*** (0.1000)	0.2985*** (0.0998)
pr	-0.0063 (0.0105)	-0.0059 (0.0107)	-0.0045 (0.0112)	-0.0080 (0.0111)	-0.0061 (0.0112)
rmax	0.0077*** (0.0028)	0.0086*** (0.0029)	0.0084*** (0.0029)	0.0082*** (0.0028)	0.0082*** (0.0028)
rmin	-0.0087*** (0.0029)	-0.0092*** (0.0030)	-0.0091*** (0.0030)	-0.0087*** (0.0030)	-0.0086*** (0.0030)
srad	-0.0042*** (0.0014)	-0.0041*** (0.0015)	-0.0042*** (0.0015)	-0.0043*** (0.0015)	-0.0043*** (0.0015)
tmmx	0.0180*** (0.0036)	0.0169*** (0.0037)	0.0172*** (0.0038)	0.0181*** (0.0037)	0.0175*** (0.0037)
vs	-0.0337 (0.0242)	-0.0204 (0.0153)	-0.0191 (0.0165)	-0.0207 (0.0137)	-0.0305* (0.0160)
vs × frac10	-0.7623** (0.3600)	-30.84** (15.26)	-3.213** (1.616)	-1.653 (1.716)	-0.0108 (0.1001)
vs × frac1020	-0.0655 (0.1351)	-0.9830 (0.6719)	0.0117 (0.2916)	0.8961 (0.5493)	-0.3101 (0.2473)
vs × frac2030	-0.0893 (0.1103)	-0.0363 (0.4458)	0.0941 (0.1569)	-0.1222 (0.2011)	-0.0948 (0.1124)
vs × frac30100	0.0649 (0.0408)	0.0128 (0.0544)	-0.0116 (0.0446)	-0.0563 (0.0608)	0.0684* (0.0370)
<i>Fixed-effects</i>					
sym	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	11,635	11,635	11,635	11,635	11,635
R ²	0.90651	0.90764	0.90654	0.90646	0.90581

Clustered (sym) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Lastly, we explore whether the impact of wind blowing from CAFOs is stronger in recent years (2021-2023) when there are more monitors in the sample. Table 3.6 reveals a more pronounced relationship between wind from CAFOs within 10 km and ammonia levels, as compared to Column 1 of Table 3.5. While the effect of beef CAFOs shows a slight decrease, the impacts of dairy and poultry CAFOs markedly increase, with the poultry coefficient also achieving statistical significance. These findings suggest varying impacts on ammonia concentrations over different years. Possible explanations for this are: 1) the panel of ammonia data is more complete in the latter years with more monitors, 2) the CAFO data is more accurate in the latter years. The CAFO data is constant over time and CAFOs may have come online later in sample.

Having estimated the impact of CAFOs on ambient ammonia, we now turn to estimating the impact on atmospheric methane. Table 3.7 presents the results from estimating Equation (3.2) both with and without state-month fixed effects in the first two columns, showing very similar outcomes. The analysis is conducted at the grid level. Thus, a one-unit increase in the number of CAFOs within a grid corresponds to a 0.04% rise in atmospheric methane in that grid. Similarly, being located in a shale play corresponds to 0.11% higher methane levels in that grid. Both of these results align with our a priori expectations. Table 3.7 also shows how weather and season affect methane concentrations. The seasonal variations indicated in the table is consistent with environmental science literature such as Khalil and Rasmusses (1983) who report that methane concentrations peak during the fall and are lowest during the summer in the northern hemisphere.

As discussed in the identification section, we have two types of grids: those with CAFOs, henceforth referred to as CAFO grids, and those without, referred to as non-CAFO grids. In theory, if methane is uniformly distributed in the atmosphere because it is a global pollutant, then grid cells with above average levels of methane should be identifying a methane source in a nearby location. The difference in atmospheric methane concentrations between these two groups, while controlling for weather, seasonality, and being located in a shale play, can be attributed to the presence of CAFOs. However, one concern arises from the global nature of methane. Methane originating from a CAFO in a CAFO grid can travel to a nearby non-CAFO grid, raising methane

Table 3.6: Regression Results for Recent Years (2021-2023)

Dependent Variable:	Ammonia Concentration				
Model:	All	Beef	Dairy	Poultry	Hogs
<i>Variables</i>					
frac10	23.95** (9.663)	216.5* (121.9)	66.20** (27.77)	36.17** (17.21)	-1.794 (2.384)
frac1020	-0.5967 (2.977)	25.76* (13.96)	-2.778 (5.156)	-10.37 (9.347)	4.366 (4.966)
frac2030	1.773 (3.275)	0.8314 (10.09)	-3.214 (4.261)	14.58 (12.59)	3.896 (2.555)
frac30100	-0.6550 (0.6329)	0.3380 (1.212)	1.428 (0.9814)	0.3947 (1.636)	-1.005 (0.6960)
vs	-0.0649** (0.0291)	-0.0533** (0.0241)	-0.0458* (0.0253)	-0.0527** (0.0215)	-0.0792*** (0.0262)
vs × frac10	-2.364** (1.023)	-30.66** (15.22)	-7.387** (3.218)	-2.953 (1.808)	0.2233 (0.2788)
vs × frac1020	-0.0128 (0.2963)	-2.475* (1.336)	0.1194 (0.5368)	1.153 (1.041)	-0.3851 (0.4506)
vs × frac2030	-0.2180 (0.2967)	0.0684 (0.9472)	0.2315 (0.4417)	-2.012 (1.530)	-0.3591 (0.2241)
vs × frac30100	0.0971* (0.0582)	0.0188 (0.1146)	-0.1059 (0.0897)	-0.0383 (0.1392)	0.1427** (0.0717)
<i>Fixed-effects</i>					
sym	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	3,157	3,157	3,157	3,157	3,157
R ²	0.93596	0.93760	0.93657	0.93499	0.93350

Clustered (sym) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Note: Weather variables except wind speed are not shown due to space constraints.

levels in that grid. This results in a spillover of methane into the non-CAFO grid group, which is essentially considered the control group, leading to bias.

The spillover effect can be partially addressed by excluding non-CAFO grids located near a CAFO grid. Therefore, we remove non-CAFO grids within 100 km of a CAFO grid, resulting in a control group less susceptible to spillover effects. The results with these non-CAFO grids excluded are shown in Columns 3 and 4. The effect of the number of CAFOs on methane levels is very similar to the results in the previous columns, thus mitigating the spillover concern. Overall, although

the estimated coefficients are small, our empirical results suggest higher methane concentrations in CAFO grids, in line with our hypothesis.

Table 3.7: Regression Results

Dependent Variable:	Log Methane Ratio			
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
Constant	7.537*** (8.34×10^{-5})		7.535*** (0.0001)	
Number of CAFOs	0.0004*** (1.06×10^{-5})	0.0003*** (4.54×10^{-5})	0.0005*** (1.21×10^{-5})	0.0002*** (5.32×10^{-5})
Shale Grid	0.0011*** (4.06×10^{-5})	0.0012*** (0.0002)	0.0011*** (7.53×10^{-5})	0.0007*** (0.0002)
Spring	-0.0074*** (2.16×10^{-5})		-0.0063*** (3.82×10^{-5})	
Summer	-0.0087*** (2.67×10^{-5})		-0.0095*** (4.53×10^{-5})	
Winter	-0.0033*** (2.04×10^{-5})		-0.0017*** (3.5×10^{-5})	
pet	-0.0027*** (1.3×10^{-5})	-0.0003 (0.0002)	-0.0027*** (2.06×10^{-5})	-0.0005* (0.0003)
pr	-0.0083*** (8.53×10^{-5})	-0.0042*** (0.0008)	-0.0088*** (0.0002)	-0.0040*** (0.0011)
rmax	2.22×10^{-5} *** (7.62×10^{-7})	3.56×10^{-5} *** (8.47×10^{-6})	3.31×10^{-5} *** (1.28×10^{-6})	3.61×10^{-5} *** (1.43×10^{-5})
rmin	9.61×10^{-5} *** (1.23×10^{-6})	-4.84×10^{-6} (1.82×10^{-5})	8.54×10^{-5} *** (2.02×10^{-6})	-2.38×10^{-5} (2.83×10^{-5})
srad	2.21×10^{-5} *** (2.59×10^{-7})	3.69×10^{-5} *** (4.96×10^{-6})	1.62×10^{-5} *** (4.51×10^{-7})	2.8×10^{-5} *** (6.13×10^{-6})
tmmx	0.0003*** (1.17×10^{-6})	-2.4×10^{-5} (1.96×10^{-5})	0.0003*** (1.94×10^{-6})	1.78×10^{-6} (2.9×10^{-5})
vs	0.0002*** (3.15×10^{-6})	-0.0001** (4.27×10^{-5})	0.0002*** (5.63×10^{-6})	-0.0001* (5.82×10^{-5})
<i>Fixed-effects</i>				
state-month	No	Yes	No	Yes
<i>Fit statistics</i>				
SE: Clustered by	grid	state-month	grid	state-month
Observations	4,856,927	4,856,927	1,614,715	1,614,715
R ²	0.18374	0.44611	0.19046	0.46520

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Note: Columns (1) and (2) are the estimation results for the full sample.

Columns (3) and (4) are results after dropping non-CAFO grids within 100 km of a CAFO grid.

3.4 Discussion, Conclusions, and Future Research

This study investigates the impact of CAFOs on ammonia and methane concentrations across the United States, leveraging ground-level monitoring data for ammonia and satellite-based observations for methane. Our findings highlight significant associations between CAFO proximity and elevated ammonia levels, as well as heightened methane concentrations in grid areas containing CAFOs. These results illustrate the localized environmental impacts of intensive animal agriculture, highlighting both the spatial distribution and seasonal variability of these pollutants.

Ammonia levels exhibit a clear positive relationship with the fraction of days wind blew from CAFOs within 10 km of monitoring sites. This proximity effect was robust even after accounting for seasonal variations, weather conditions, and site-specific factors. The study's findings suggest that ammonia emissions from CAFOs influence local air quality, particularly during periods when weather conditions facilitate pollutant dispersion. However, the lack of significant associations beyond the 10 km radius underscores the localized nature of ammonia transport, aligning with previous research highlighting its short atmospheric lifespan.

Further analysis by CAFO type revealed nuanced impacts, with dairy and poultry operations showing more pronounced associations with elevated ammonia levels compared to beef operations. These disparities likely stem from varying management practices and animal husbandry techniques across different livestock types. Nevertheless, the study was limited in its ability to account for specific operational factors such as manure management practices, which could have provided deeper insights into these observed differences.

In examining methane concentrations, our study found that areas with higher numbers of CAFOs exhibited slightly elevated methane levels compared to non-CAFO areas, even after controlling for weather conditions and seasonal variations. This finding supports the hypothesis that methane emissions from CAFOs contribute to regional atmospheric methane burdens. Importantly, the exclusion of non-CAFO grids near CAFO-grids helped mitigate potential spillover effects, strengthening the validity of our findings.

Despite robust findings, some limitations have to be considered. First, the study's reliance on aggregated data from ground monitors and satellite observations imposes constraints on detailed spatial and temporal resolutions. Variability in monitor installation dates and sampling frequencies created an unbalanced panel, which may have influenced the study's ability to capture temporal trends accurately.

Second, while wind direction was used as a proxy to attribute ammonia emissions to CAFO sources, this approach assumes uniform emission characteristics across all CAFOs within a given radius and does not differentiate between individual facilities. Future research could benefit from incorporating more precise emission estimation methods, such as atmospheric modeling or inclusion of pollutant dispersion models, to improve spatial specificity and source attribution.

Third, the study's focus on methane emissions from CAFOs did not fully account for potential interactions with other methane sources, such as oil and gas extraction activities in shale regions. Although we included shale play indicators in our models, there is substantial heterogeneity among shale plays and even within a shale play. These differences are not captured in our model. Furthermore, our study does not include locations of oil and gas wells, thereby affecting the precision of the results.

Our findings can inform policy interventions regarding agricultural management practices. Given the localized nature of ammonia emissions, strategies aimed at mitigating CAFO-related pollution could benefit from targeted regulations tailored to specific livestock types and regional environmental conditions. Improved monitoring and reporting standards for ammonia emissions from CAFOs could also enhance regulatory compliance and environmental stewardship efforts.

The estimates of the association between CAFOs and pollutants can be converted to social costs using some back-of-the-envelope calculations. For CAFOs within 10 km of an ammonia monitor, an additional day of wind blowing from a CAFO to the monitor increases ammonia concentration by about $0.103 \mu\text{g}/\text{m}^3$. With an air volume of 10 cubic km, this converts to 1.03 tonnes of ammonia. There is little literature on the social costs of ammonia, although Gilmore et al. (2019) report emission-weighted damages from ground level ammonia, with a range of \$38,000 to \$49,000 per

tonne. Our estimates thus indicate damages that range between \$39,140 and \$50,470 per day. Calculating the damages from methane emissions involves complex factors such as temperature, air pressure, and specific methane types, which we leave to future work.

There are several next steps for this research. We will obtain satellite-based measurements of ammonia concentrations and compare these estimates with those from ground monitors. In addition to modeling the distance between CAFOs and monitors using various radii, another approach is to model distance as a continuous variable and include a squared term to account for diminishing impacts. Several different specifications, including incorporating CAFO size data, can also be explored. We will also work on identifying ammonia and methane emissions and tying the results to health and economic damages.

A key objective is to investigate the negative and insignificant relationship between ammonia levels and wind blowing from beef CAFOs. This can be addressed by examining specific cases, such as a region with only beef CAFOs.

For the methane component of the study, we have so far only considered the number of CAFOs in each grid cell. Investigating how CAFO size and type are associated with atmospheric methane would enhance our understanding of CAFO methane emissions.

Finally, we can exploit the spatial variation in CAFOs as shown in Figure 3.1. For instance, in several southern states, poultry and hog farms are concentrated in the northern parts, but not in the southern parts. Similarly, there is east-to-west variation in California. Leveraging these spatial differences can help us better understand the relationship between CAFOs and pollutant levels.

In conclusion, this study contributes to the growing body of literature on CAFO emissions and their environmental and climate impacts. By assembling a novel dataset of CAFO locations across the US, and estimating the relationship between CAFO presence and pollutant levels, we add to the literature focusing on quantifying CAFO-related ammonia and methane. Our results can be supplemented with additional research that utilize more precise techniques, such as animal-level monitors and aircraft flyovers.

Bibliography

- [1] Luís Aguiar-Conraria, Gilmar Conceigao, and Maria Joana Soares. How far is gas from becoming a global commodity? *The Energy Journal*, 43(4):179–198, 2022.
- [2] J. D. Angrist, G. W. Imbens, and D. B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.
- [3] Robert Baron, Paul Bernstein, W David Montgomery, and Sugandha Tuladhar. Macroeconomic impacts of lng exports from the united states. *Economics of Energy & Environmental Policy*, 4(1):37–52, 2015.
- [4] M. Bernedo, P. J. Ferraro, and M. Price. The persistent impacts of norm-based messaging and their implications for water conservation. *Journal of Consumer Policy*, 37(3):437–452, 2014.
- [5] S. P. Bhanot. Isolating the effect of injunctive norms on conservation behavior: New evidence from a field experiment in california. *Organizational Behavior and Human Decision Processes*, 163:30–42, 2021.
- [6] H. S. Bloom. Accounting for no-shows in experimental evaluation designs. *Evaluation Review*, 8(2):225–246, 1984.
- [7] J. Boomhower and L. Davis. Do energy efficiency investments deliver at the right time? *American Economic Journal: Applied Economics*, 12(1):115–139, 2020.
- [8] Paul Brehm. Natural gas prices, electric generation investment, and greenhouse gas emissions. *Resource and Energy Economics*, 58:101106, 2019.
- [9] D. A. Brent, J. H. Cook, and S. Olsen. Social comparisons, household water use, and participation in utility conservation programs: Evidence from three randomized trials. *Journal of the Association of Environmental and Resource Economists*, 2(4):597–627, 2015.

- [10] D. A. Brent, C. Lott, M. Taylor, J. Cook, K. Rollins, S. Stoddard, and J. Cook. Are normative appeals moral taxes? evidence from a field experiment on water conservation, 2017. Louisiana State Department of Economics Working Papers Series.
- [11] D. A. Brent, C. Lott, M. Taylor, J. Cook, and S. Stoddard. What causes heterogeneous responses to social comparison messages for water conservation? *Environmental and Resource Economics*, 77(3):503–537, 2020.
- [12] J. Burkhardt, K. Gillingham, and P. K. Kopalle. Experimental evidence on the effect of information and pricing on residential electricity consumption, 2019. Working Paper 25576, National Bureau of Economic Research.
- [13] Alyssa M Burns, Gabriel Chandler, Kira J Dunham, and Annmarie G Carlton. Data gap: Air quality networks miss air pollution from concentrated animal feeding operations. *Environmental Science & Technology*, 57(49):20718–20725, 2023.
- [14] Michael A Clark, Nina GG Domingo, Kimberly Colgan, Sumil K Thakrar, David Tilman, John Lynch, Inês L Azevedo, and Jason D Hill. Global food system emissions could preclude achieving the 1.5 and 2 c climate change targets. *Science*, 370(6517):705–708, 2020.
- [15] Claudia Copeland. Animal waste and water quality: EPA regulation of concentrated animal feeding operations CAFO. *Water Pollution Issues and Developments*, (Ed. SV Thomas.) pp, pages 51–73, 2008.
- [16] Melanie Cozad and Jacob LaRiviere. Fuel price increases and the timing of changes in household driving decisions. *Journal of Environmental Economics and Management*, 65(2):194–207, 2013.
- [17] Joseph A Cullen and Erin T Mansur. Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution. *American Economic Journal: Economic Policy*, 9(3):106–133, 2017.

- [18] Nina GG Domingo, Srinidhi Balasubramanian, Sumil K Thakrar, Michael A Clark, Peter J Adams, Julian D Marshall, Nicholas Z Muller, Spyros N Pandis, Stephen Polasky, Allen L Robinson, et al. Air quality–related health damages of food. *Proceedings of the National Academy of Sciences*, 118(20):e2013637118, 2021.
- [19] Scott J Eilerman, Jeff Peischl, J Andrew Neuman, Thomas B Ryerson, Kenneth C Aikin, Maxwell W Holloway, Mark A Zondlo, Levi M Golston, Da Pan, Cody Floerchinger, et al. Characterization of ammonia, methane, and nitrous oxide emissions from concentrated animal feeding operations in northeastern colorado. *Environmental Science & Technology*, 50(20):10885–10893, 2016.
- [20] Harrison Fell and Daniel T Kaffine. The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions. *American Economic Journal: Economic Policy*, 10(2):90–116, 2018.
- [21] P. J. Ferraro, J. J. Miranda, and M. K. Price. The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment. *American Economic Review*, 101(3):318–322, 2011.
- [22] P. J. Ferraro and M. K. Price. Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *The Review of Economics and Statistics*, 95(1):64–73, 2013.
- [23] K. S. Fielding, A. Spinks, S. Russell, R. McCrea, R. Stewart, and J. Gardner. An experimental test of voluntary strategies to promote urban water demand management. *Journal of Environmental Management*, 114:343–351, 2013.
- [24] Levi M Golston, Da Pan, Kang Sun, Lei Tao, Mark A Zondlo, Scott J Eilerman, Jeff Peischl, J Andrew Neuman, and Cody Floerchinger. Variability of ammonia and methane emissions from animal feeding operations in northeastern colorado. *Environmental Science & Technology*, 54(18):11015–11024, 2020.

- [25] Christoph Halser and Florentina Paraschiv. Fuelling the energy transition: The effect of german wind and pv electricity infeed on ttf gas prices. In *Quantitative Energy Finance: Recent Trends and Developments*, pages 135–179. Springer, 2024.
- [26] Catherine Hausman and Ryan Kellogg. Welfare and distributional implications of shale gas. Technical report, National Bureau of Economic Research, 2015.
- [27] Matthew N Hayek and Scot M Miller. Underestimates of methane from intensively raised animals could undermine goals of sustainable development. *Environmental Research Letters*, 16(6):063006, 2021.
- [28] Jinhyok Heo, Peter J Adams, and H Oliver Gao. Public health costs of primary pm_{2.5} and inorganic pm_{2.5} precursor emissions in the united states. *Environmental science & technology*, 50(11):6061–6070, 2016.
- [29] J Scott Holladay and Jacob LaRiviere. The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy. *Journal of Environmental Economics and Management*, 85:205–227, 2017.
- [30] Chenghan Hou and Bao H Nguyen. Understanding the us natural gas market: A markov switching var approach. *Energy Economics*, 75:42–53, 2018.
- [31] M. M. Jaime Torres and F. Carlsson. Direct and spillover effects of a social information campaign on residential water-savings. *Journal of Environmental Economics and Management*, 92:222–243, 2018.
- [32] Joseph P Janzen, Aaron Smith, and Colin A Carter. Commodity price comovement and financial speculation: the case of cotton. *American Journal of Agricultural Economics*, 100(1):264–285, 2018.
- [33] Reid Johnsen, Jacob LaRiviere, and Hendrik Wolff. Fracking, coal, and air quality. *Journal of the Association of Environmental and Resource Economists*, 6(5):1001–1037, 2019.

- [34] MAK Khalil and RA Rasmussen. Sources, sinks, and seasonal cycles of atmospheric methane. *Journal of Geophysical Research: Oceans*, 88(C9):5131–5144, 1983.
- [35] Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American economic review*, 99(3):1053–1069, 2009.
- [36] Lutz Kilian and Daniel P Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied econometrics*, 29(3):454–478, 2014.
- [37] Lutz Kilian and Xiaoqing Zhou. The econometrics of oil market var models. NBER Working Paper 26795, National Bureau of Economic Research, 2020.
- [38] Christopher R Knittel, Konstantinos Metaxoglou, and André Trindade. Environmental implications of market structure: Shale gas and electricity markets. *International Journal of Industrial Organization*, 63:511–550, 2019.
- [39] Xi Lu, Jackson Salovaara, and Michael B McElroy. Implications of the recent reductions in natural gas prices for emissions of co2 from the us power sector. *Environmental science & technology*, 46(5):3014–3021, 2012.
- [40] S. Lurbé, J. Burkhardt, C. Goemans, D. Manning, and L. Hans. Further evidence on social comparison and residential water use. *Water Resources and Economics*, 41:100214, 2023.
- [41] J. S. Mankin, I. Simpson, A. Hoell, R. Fu, J. Lisonbee, A. Sheffield, and D. Barrie. Noaa drought task force report on the 2020–2021 southwestern u.s. drought, 2021. NOAA Drought Task Force, MAPP.
- [42] Megan E McCabe, Ilana B Pollack, Emily V Fischer, Kathryn M Steinmann, and Dana R Caulton. Isolating methane emissions from animal feeding operations in an interfering location. *Atmospheric Chemistry and Physics*, 23(13):7479–7494, 2023.

- [43] Maria C Mirabelli, Steve Wing, Stephen W Marshall, and Timothy C Wilcosky. Race, poverty, and potential exposure of middle-school students to air emissions from confined swine feeding operations. *Environmental health perspectives*, 114(4):591–596, 2006.
- [44] M. Moerbeek and S. V. Schie. What are the statistical implications of treatment non-compliance in cluster randomized trials: a simulation study. *Statistics in Medicine*, 38(26):5071–5084, 2019.
- [45] M. Nemati, S. Buck, and H. Soldati. Heterogeneous effects of real-time consumption analytics on residential water consumption, 2016. Available at SSRN 2877233.
- [46] Ilana B Pollack, Megan E McCabe, Dana R Caulton, and Emily V Fischer. Enhancements in ammonia and methane from agricultural sources in the northeastern colorado front range using observations from a small research aircraft. *Environmental Science & Technology*, 56(4):2236–2247, 2022.
- [47] Zach Raff and Andrew Meyer. Cafos and surface water quality: evidence from wisconsin. *American Journal of Agricultural Economics*, 104(1):161–189, 2022.
- [48] Valerie A Ramey. Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, 2:71–162, 2016.
- [49] M. E. Renwick and R. D. Green. Do residential water demand side management policies measure up? an analysis of eight california water agencies. *Journal of Environmental Economics and Management*, 40(1):37–55, 2000.
- [50] E. D. Sherwin, R. M. Meyer, and I. M. L. Azevedo. Limitations of econometric evaluation of nonrandomized residential energy efficiency programs: A case study of northern california rebate programs. *Environmental Data Science*, 1:e1, 2022.
- [51] Christopher A Sims. Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, pages 1–48, 1980.

- [52] Stacy Sneeringer. Does animal feeding operation pollution hurt public health? a national longitudinal study of health externalities identified by geographic shifts in livestock production. *American Journal of Agricultural Economics*, 91(1):124–137, 2009.
- [53] Stacy Sneeringer. A national, longitudinal study of the effects of concentrated hog production on ambient air pollution. *American Journal of Agricultural Economics*, 92(3):821–835, 2010.
- [54] B. Stitzel and C. Rogers. No tips please: Informational messaging for water conservation, 2023.
- [55] James H Stock and Matthew Zaragoza-Watkins. The market and climate implications of us lng exports. Technical report, National Bureau of Economic Research, 2024.
- [56] Richard W Todd, N Andy Cole, R Nolan Clark, Thomas K Flesch, Lowry A Harper, and Bok H Baek. Ammonia emissions from a beef cattle feedyard on the southern high plains. *Atmospheric Environment*, 42(28):6797–6805, 2008.
- [57] United Nations Environment Programme and Climate and Clean Air Coalition. Global methane assessment: Benefits and costs of mitigating methane emissions. Report, United Nations Environment Programme, Nairobi, 2021.
- [58] U.S. Environmental Protection Agency. NPDES CAFO Regulations Implementation Status Report 2022. <https://www.epa.gov/npdes>.
- [59] U.S. Environmental Protection Agency. *Environmental Assessment of Proposed Revisions to the National Pollutant Discharge Elimination System Regulation and the Effluent Guidelines for Concentrated Animal Feeding Operations*. 2001.
- [60] U.S. Environmental Protection Agency, Office of Water. *National Water Quality Inventory: 1996 Report to Congress*. Number EPA841-R-97-008. U.S. Environmental Protection Agency, Washington, D.C., 1997.
- [61] USGCRP. Fifth national climate assessment, 2023. U.S. Global Change Research Program.

[62] NT Vechi, J Mellqvist, J Samuelsson, B Offerle, and C Scheutz. Ammonia and methane emissions from dairy concentrated animal feeding operations in california, using mobile optical remote sensing. *Atmospheric Environment*, 293:119448, 2023.