CLIMATE AND THE OREGON ENERGY GRID: MODELING CLIMATE-DRIVEN DEMAND AND EMISSIONS UNDER 21ST CENTURY WARMING SCENARIOS

by

SAMUEL R. GOLDEN

A THESIS

Presented to the Department of Data Science
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

May 2026

An Abstract of the Thesis of

Samuel Golden for the degree of Bachelor of Science
in the Department of Natural Sciences to be taken June 2026

Title: Climate and the Oregon Energy Grid: Modeling Climate-Driven Demand and Emissions Under 21st Century Warming Scenarios

Approved: *Andrew Muehleisen, Ph.D.*
Primary Thesis Advisor

The state of Oregon has warmed approximately 2.2°F over the past century, with continued warming of roughly 0.4 to 0.9°F per decade projected across the state's climate regions through 2100. The PacifiCorp West (PACW) Balancing Authority, which serves electricity to the geographic majority of Oregon, is directly exposed to that warming through its temperature-driven demand and emissions profile. This thesis quantifies how rising temperatures will reshape daily PACW demand and total CO_2 emissions through the warming ranges plausible by 2050 and 2100, and how those climate-driven changes compare in magnitude to a structural decarbonization produced by PacifiCorp's July 2022 retirement of coal generation. Using ten years of daily PACW operational data from EIA's Hourly Electric Grid Monitor merged with PRISM climate data extracted at five Oregon stations, I develop a hybrid modeling framework that combines an ordinary least squares regression on temperature variables with tree-based machine learning models on residual variation, validated using a chronological train-test split and applied identically across all five climate stations. I find that the coal retirement reduced average daily PACW CO_2 emissions by approximately 12,000 metric tons per day, controlling for demand and weather, an effect that is statistically dominant in every specification tested and robust to climate-station choice. Across warming scenarios from 0°F to +15°F, demand follows a

shallow U-curve with a minimum near +4°F, and total CO_2 emissions decline by 5.6% to 15.8% across cities at moderate warming due to a dispatch-mix asymmetry I term as the carbon-lopsided grid: PACW's summer generation portfolio is materially cleaner per MW than its winter portfolio. The structural decarbonization from the coal retirement is approximately an order of magnitude larger than the deepest climate-driven reduction at any plausible warming level. The policy implication is that on a grid in active structural transition, structural decisions are the lever with the largest demonstrated effect on emissions, and the climate signal is a secondary contributor.

Acknowledgements

I would like to thank my primary thesis advisor, Andrew Muehleisen, for their guidance, support, and thoughtful feedback throughout this project. I would like to thank Dr. Larissa Stiglich for serving as my Clark Honors College representative within my thesis committee. I am also grateful to Dr. Kristen Rahilly for helping me find climate projection resources that strengthened the direction of my research. Finally, I would like to thank Jake Searcy and the whole of the Data Science department at University of Oregon for introducing me to machine learning models and helping me see their possibilities as tools for research and analysis.

Table of Contents

Chapter 1: Introduction 10

1.1: Motivation 10

1.2 Research Question 11

1.3 Contribution and Significance 12

1.4 Thesis Overview 13

Chapter 2: Background and Literature Review 14

2.1 Temperature-Electricity Demand Relationships 14

2.2 Pacific Northwest Climate and Grid Context 15

2.3 The 2022 PacifiCorp Coal Retirement 17

2.4 Previous Studies of Climate-Driven Grid Projections 20

2.5 Gaps in the Literature and Contribution of This Study 21

Chapter 3: Data 24

3.1 Overview 24

3.2 Electricity Data: EIA Grid Monitor (PACW) 25

3.3 Climate Data: PRISM 30

3.4 Data Cleaning and Alignment 34

3.5 Feature Engineering 35

3.6 Exploratory Data Analysis 38

Chapter 4: Methods 41

4.1 Overview 41

4.2 The Hybrid Modeling Architecture 41

4.3 Demand Model Specification 43

4.4 Emissions Model Specification 45

4.5 Train-Test Validation and Diagnostics 46

4.6 Warming Scenario Construction 48

4.7 Cross-City Comparison 51

4.8 Methodological Caveats 53

4.9 Evaluation Summary 55

Chapter 5: Results 57

5.1 Overview 57

5.2 Demand Model Findings 57

5.3 Emissions Model Findings 59

5.4 Warming Scenarios: Demand and Emissions 61

5.5 The Coal Exit and Its Quantification 63

5.6 Summary of Findings 64

Chapter 6: Discussion 66

6.1 Interpreting the U-Curve: A Heating-Dominated System 66

6.2 A Carbon-Lopsided Grid: The Emissions Asymmetry 69

6.3 Structural vs. Climate-Driven Decarbonization 72

6.4 Comparison to Existing Literature 74

6.5 Limitations 76

6.6 Future Work 78

Chapter 7: Conclusion and Reflection 81

7.1 Summary and Policy Takeaways 81

Bibliography 83

Supporting Materials

Dashboard-A Decade of Oregon Energy: https://huggingface.co/spaces/Sam-6464/oregon-energy-dashboard

List of Figures

1. Figure 2.3: CO_2 Emissions Before and After the Coal Exit p. 19
2. Figure 3.2: The PacifiCorp West Service Territory p. 29
3. Figure 3.2.1: Daily PACW Demand July 2015 - October 2025 p. 31
4. Figure 4.6: Climate Projections from the OCCRI Seventh Annual Oregon Climate Assessment p. 51
5. Figure 5.2: Demand Under Warming p. 59
6. Figure 5.3: Emissions Under Warming p. 61
7. Figure 5.4: Demand Through 2100 p. 63
8. Figure 5.4.1: Emissions Through 2100 p. 64
9. Figure 6.2: Cold and Hot Side Emissions Response p. 71

List of Tables

1. Table 3.6: Summary statistics for the merged PACW-Eugene daily dataset p. 40
2. Table 4.3: Feature partition for the hybrid demand model p. 44
3. Table 4.2: Hybrid XGBoost performance across the five climate stations p. 53
4. Table 6.1: Hot and Cold Side Demand Response p. 68

Chapter 1: Introduction

1.1: Motivation

In March 2026, the National Oceanic and Atmospheric Administration confirmed what most Oregonians had already experienced: the winter that had just ended tied 1934 for the warmest in 130 years of state record-keeping, with an average temperature of 38.7°F against a long-run norm of 32.4°F. Snowpack across the state hovered at 29% of normal, among the worst readings since the early 1980s, and this pattern was not unique to Oregon. Nationally, the 2025-2026 winter was the second-warmest winter in 131 years of contiguous-U.S. records; nine western states tied or set their own all-time warmest-winter records. The two warmest winters in the United States have now both occurred within the last three years.

As an Oregonian myself, to say this data is unsettling would be a gross understatement. Growing up in San Francisco, California, I came of age watching wildfire smoke turn the sky orange in the daytime as the words “climate change” rang through my ears from the news. The instinct to want to preserve what is around us has been instilled in me for years. By the time I got to college, that instinct had narrowed into curiosity, and I began asking questions. One of such questions, which I found most compelling, was: what will happen to the systems that we depend on such as power, water, and food, as the climate they were built for changes?

This thesis is an attempt to answer that question, scoped to a system I now live with directly: the Pacific Northwest electricity grid. It just so happens that as I’ve conducted this research, we’ve experienced a record-breaking winter.

The grid is a particularly compelling place to start because grid decisions are slow, and data is easily quantifiable and accessible. A power plant operates forty to sixty years after it is

built. Transmission infrastructure lasts longer. Capacity additions take roughly a decade from planning to commissioning. The system that exists today is for the most part, the system that was planned in the 1980s and 1990s, which means the system that exists in 2050 and 2070 needs to be planned right now. Any disconnect between the climate the grid was built for, and the climate it actually has to operate in, compounds over those time scales. My analysis: quantifying that disconnect, even imperfectly, is more useful than waiting for it to arrive.

When I started this project, I was initially consumed by a central worry – the snowball effect. My intuition was that as warming makes summers hotter, it will drive up demand, which will drive up emissions, which in turn drives more warming, and the cycle begins. A terrifying possibility. Luckily, I turned out to be more wrong than right for the Oregon system that this thesis studies. The actual relationship between warming and emissions on Oregon's grid is far more nuanced, more system-specific, and in many ways more hopeful than my snowball framing suggested. That this thesis ended up partially undoing my own starting hypothesis is, in retrospect, the part of the project I am most glad to be reporting on.

1.2 Research Question

The central question this thesis sets out to answer is: how do rising temperatures reshape daily electricity demand and total CO_2 emissions on the PacifiCorp West Balancing Authority, through the warming ranges that are plausible by 2050 and 2100? The thesis also asks a secondary, comparative question that emerged from the data rather than from the original framing: how have past structural grid changes altered the future of Oregon energy and emissions?

To answer the first requires a credible model of how PACW emissions respond jointly to temperature and to demand, fit in a way that respects the major structural change in the grid's

generation mix over the analytical period. The methodology developed in Chapter 4 attempts to build such a model. However, the answer to the second comparative question turns out to be the most policy-relevant finding of the project and is developed in Chapter 6.

1.3 Contribution and Significance

This thesis contributes a methodological framework and a set of empirical findings that together fill a gap in the regional climate-grid literature. On the methods side, I develop a hybrid modeling framework that combines an ordinary least squares (OLS) regression for the temperature component of demand and emissions with tree-based machine learning models (Random Forest and XGBoost) for the residual variation that OLS cannot capture. The hybrid design is built specifically to support warming scenario projections and is documented in Chapter 4.

On the empirical side, the thesis produces three findings worth naming up front. First, the PACW grid is heating-dominated rather than cooling-dominated: cold winter days drive both more demand and more emissions per megawatt-hour than hot summer days do, because winter dispatch leans on dirty energy imports while summer dispatch leans on clean hydroelectric, wind, and solar generation. Second, in July 2022 the PacifiCorp Region retired coal generation which reduced average daily CO_2 emissions on the PACW grid by approximately 11,950 to 12,790 metric tons per day after controlling for demand and weather, an effect that is statistically dominant in every model specification tested and that is robust to the choice of climate-station signal. Third, plausible warming through 2050 produces a partial decarbonization effect on the PACW system rather than the snowball pattern initially expected: the warming-driven reduction in winter heating dispatch (a dirty marginal stack) outweighs the warming-driven increase in

summer cooling dispatch (a clean marginal stack), at least at the warming levels currently projected by the Oregon Climate Change Research Institute.

1.4 Thesis Overview

Chapter 2 places the thesis in the context of existing literature on temperature-driven electricity demand, Pacific Northwest grid behavior, the supply-side dynamics of heat waves and hydropower, and the structural significance of the 2022 PacifiCorp coal retirement.

Chapter 3 documents the data: hourly PACW operational data from EIA's Grid Monitor, daily PRISM climate data extracted at five locations across Oregon, and the cleaning, alignment, and feature engineering that produced the merged analytical dataset.

Chapter 4 develops the methodology in a detailed and reproducible manner, following chronological order.

Chapter 5 reports the modeling results of demand and emissions across the full warming range, including the demand response curve, the emissions inflection behavior, and the structural-break magnitude.

Chapter 6 interprets the findings about the relative magnitude of climate-driven versus structural decarbonization, the implications for grid planning horizons, the comparison to other states and national findings, and the directions in which future work could extend the framework developed here.

Chapter 7 closes with a summary of the policy takeaways.

Chapter 2: Background and Literature Review

2.1 Temperature-Electricity Demand Relationships

The relationship between temperature and electricity demand is one of the most empirically robust patterns in energy economics. Demand responds nonlinearly to temperature, typically with a U-shaped profile: consumption rises sharply at low temperatures as heating loads activate, falls to a minimum at moderate or comfort-zone temperatures, and rises again at high temperatures as cooling loads activate. The exact shape and slope of this curve varies meaningfully by region, by sector, and by the appliance and adaptation patterns of the underlying customer base, but the U-shape itself appears almost universally in datasets that span both heating and cooling regimes.

The U-shape was characterized for U.S. residential consumption by Aroonruengsawat and Auffhammer (2011), who used household-level billing panel data across California's climate zones to estimate flexible temperature response functions and found pronounced geographic heterogeneity in both the location of the demand minimum and the steepness of the cold- and hot-side responses. Auffhammer and Aroonruengsawat (2012) extended this work into projection territory, simulating residential consumption changes through 2100 under different climate, price, and population scenarios and finding consumption increases of 1–6% under varied adaptation assumptions. Yao (2021) generalized the result globally using satellite night-light data as a consumption proxy and identified a U-shape with a global minimum near 14.6°C, with urban and industrial regions exhibiting slightly higher thresholds than rural ones.

The cooling side of the U-curve has received particular attention in the literature on climate adaptation, because air-conditioning adoption is the single most important moderator of how warming translates into demand. As AC use rises, the same temperature increase produces a

larger demand increase, and the long-run elasticity of demand to temperature can be several times its short-run value. Auffhammer, Baylis, and Hausman (2017) projected that warming will sharply increase both the frequency and the severity of peak load events across the United States, with the largest effects in regions where AC adoption is currently low and rising. Kumar et al. (2020) showed that the temperature-demand relationship in California becomes substantially more sensitive to heat at the upper end of the demand distribution than at the mean, indicating that peak load events respond more strongly to extreme heat than total consumption does. Deroubaix et al. (2021) integrated heating- and cooling-degree-day projections globally and found that cooling demand increases under warming consistently outweigh heating demand reductions, although the projections carry substantial model uncertainty. Emonekwe and Emodi (2022), using quantile regression on G7 economies, found that residential electricity demand responds more strongly to cold days than to hot days in heating-dominated economies, providing useful methodological precedent for the asymmetric specifications used later in this thesis.

The methodological consensus across these studies is clear: temperature-demand relationships should be modeled with flexible functional forms, typically piecewise or quadratic specifications, with separate treatment for heating and cooling regimes. Linear specifications are universally rejected as inadequate.

2.2 Pacific Northwest Climate and Grid Context

The Pacific Northwest occupies an unusual position in the temperature-demand literature. Historically, the region's demand has been heating-dominated rather than cooling-dominated, with limited summer load sensitivity attributable to low residential air-conditioning penetration relative to most of the United States. Winter demand has traditionally driven both peak load events and seasonal capacity planning. This pattern is visible directly in the PACW data

analyzed in this thesis (described in Chapter 3) and is consistent with regional descriptions in the Northwest Power and Conservation Council's Seventh Power Plan (2016), which projects regional demand under evolving load patterns and climate considerations.

Two trends are reshaping this historical pattern. The first is climate warming itself. The Oregon Climate Change Research Institute's Seventh Oregon Climate Assessment reports that Oregon's annual mean temperature increased by 2.2°F over the past century and projects further increases of approximately 5°F by 2074 and as much as 7.6°F by 2100 under continued high-emissions pathways. The same report projects that snowpack across the state will decline by at least 50% by the end of the century, which has direct implications for hydroelectric generation timing and availability. The 2025-26 winter, which tied 1934 as Oregon's warmest on record, illustrated several of these dynamics in real time: snowpack reached only 29% of normal, and Governor Kotek declared drought emergencies across nine counties by April 2026 (NOAA, 2026; Urness, 2026).

The 2021 Pacific Northwest heat dome, which produced unprecedented temperatures across western Oregon and Washington and contributed to hundreds of excess deaths regionally (Philip et al., 2022), accelerated AC adoption in a population that had previously gone without it. Regional grid operators have reported new summer peak records since 2021, and utilities have expanded demand response programs to manage evening peaks. The historical assumption that PACW would remain a heating-dominated system in perpetuity is no longer safe, even if the system remains heating-dominated as of the 2015–2025 analytical window of this thesis.

On the supply side, hydroelectric generation has historically anchored the Pacific Northwest grid, with Bonneville Power Administration's federal hydro system providing the regional backbone and individual utilities like PacifiCorp operating their own portfolios of

hydro, thermal, and renewable generation. Hydropower output is climate-sensitive in ways that magnify the demand-side dynamics: late-summer flows decline in dry years, transmission and thermal generation efficiency degrade in extreme heat, and the entire system experiences compounding stress during heat-wave events. The result is that climate change affects both demand and supply on the PACW grid simultaneously, and through different mechanisms.

2.3 The 2022 PacifiCorp Coal Retirement

In July 2022, PacifiCorp retired the coal-fired generation assets that had historically anchored a substantial share of in-house electricity production within its PacifiCorp West (PACW) balancing authority. The retirement was part of a multi-year transition documented in PacifiCorp's Integrated Resource Plans (PacifiCorp, 2021), which laid out scheduled end-of-life dates for the company's coal units across its six-state operating area and committed to retiring twenty of twenty-four coal units by 2038. The shift reflects a broader pattern visible across much of the western United States, in which coal generation is being phased out and replaced with a mix of natural gas, renewables, and inter-regional imports as state policy, federal regulation, and economics push utilities toward lower-emission generation portfolios.

For PACW specifically, the consequence of the retirement was a sharp and immediate change in the grid's emissions profile. The dirtiest segment of the in-house generation stack was removed without an immediate one-for-one replacement, leaving the system to meet customer demand through a combination of remaining hydro, wind, and solar generation, supplemented by larger volumes of imported electricity from neighboring balancing authorities. Total emissions associated with consumed electricity, which include the embedded emissions of imports, fell substantially as well, although not to zero, because some of the imports themselves carry meaningful carbon content from the generation mixes of neighboring systems.

This retirement is clearly visible in the operational data the EIA publishes for PACW (described in Chapter 3) as a clean before-and-after step on July 7, 2022. Across that single date boundary, daily CO_2 emissions from PACW-generated electricity drop from a multi-year mean of roughly 19,550 metric tons per day to a post-break mean of approximately 7,300 metric tons per day. Coal-fired generation in megawatts drops to zero on the same date and remains there. Daily demand and consumed electricity remain stable, ruling out a demand-side cause. Net imports from neighboring balancing authorities surge: the daily Import Share rises from a pre-break mean of approximately 37% to a post-break mean of approximately 75%. The pattern across these variables is internally consistent with a physical retirement and inconsistent with reporting changes, unit conversions, or other artifacts that might otherwise produce a discontinuity in a single emissions series. See the figure below.

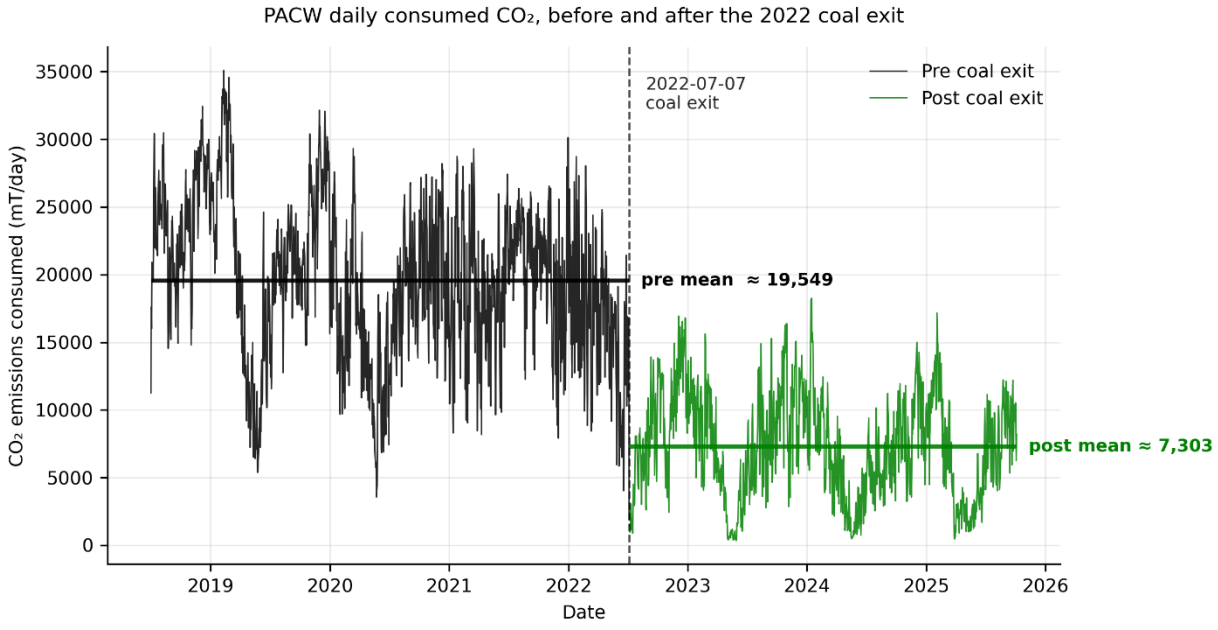


Figure 2.3: CO_2 Emissions Before and After the Coal Exit

The coal retirement is not the central question this thesis sets out to answer. The original research focus, as stated in Chapter 1, is the relationship between rising temperatures and PACW demand and emissions, not the structural decarbonization of the grid. The retirement entered the analysis because it surfaced organically while I was building the data pipeline, and because once it was visible it could not be responsibly ignored: any model that failed to account for it would conflate climate-driven and policy-driven emissions changes in ways that would compromise every downstream finding. What began as a data-quality investigation therefore became a methodological pillar of the thesis, and as Chapter 6 develops, the comparative magnitude of the coal-retirement effect against the climate-driven effect turned out to be one of the most policy-relevant findings of the project.

This retirement matters for this thesis for two reasons. First, it provides a natural before-and-after comparison of grid emissions under materially different generation mixes, with the same demand-side patterns and the same climate driving both regimes. Few balancing authorities offer such a clean within-system comparison of pre- and post-decarbonization emissions behavior in their recent operational data, and the PACW case is unusually well-documented for this type of analysis. Second, it must be modeled explicitly in any emissions specification that uses both pre- and post-retirement observations, because failing to account for it would let any model conflate climate-driven and policy-driven emissions changes. The methodological consequences of the break, including the construction of the post_july_2022 indicator variable used throughout the demand and emissions models, are documented in Chapter 4.

To support visual exploration of the structural change and the broader generation-mix transition it represents, I built an interactive dashboard alongside the formal modeling work. The dashboard allows the daily decomposition of PACW generation by fuel source, the rise in net

imports across the transition, and the corresponding shift in carbon intensity to be examined at a finer temporal grain than the static figures in this thesis can support. The dashboard is hosted here and is referenced again where relevant in subsequent chapters.

2.4 Previous Studies of Climate-Driven Grid Projections

A handful of empirical studies have attempted to project regional electricity systems forward under warming scenarios using statistical models fit to historical data. The most directly relevant for this thesis are the California-focused projections of Auffhammer and Aroonruengsawat (2012) and Kumar et al. (2020), and the national-scale projections of Auffhammer, Baylis, and Hausman (2017).

Auffhammer and Aroonruengsawat (2012) is the methodological template that most closely parallels the work in this thesis. They fit flexible temperature response functions to California household billing data and used the resulting models to project residential consumption under different SRES emissions scenarios through 2100. Their projections incorporated assumptions about population growth, electricity prices, and adaptation behavior, and they reported consumption increases of 1–6% by end-of-century depending on the scenario configuration. Their methodology is sound and widely cited, but their analysis is residential-only, focuses on consumption rather than emissions, and does not address structural changes in the generation mix.

Auffhammer, Baylis, and Hausman (2017) extended the projection framework to peak demand at the national scale, finding that climate change will substantially increase both the frequency and the severity of peak load events across the United States. Their analysis relied on aggregated load data across multiple balancing authorities and produced national-scale

conclusions, but it did not characterize individual systems with the resolution that grid planners require.

Kumar et al. (2020) examined how California summer demand responds to temperature variation across the demand distribution rather than only at the mean, using quantile regression to characterize the asymmetric sensitivity of peak versus average loads. Their finding that the upper tail of the demand distribution is substantially more temperature-sensitive than the mean is methodologically relevant for any study of warming-induced peak stress.

Deroubaix et al. (2021) provided a global-scale degree-day analysis covering 1990–2030 and characterized the trade-off between rising cooling demand and falling heating demand under warming. Their result that cooling trends exceed heating declines in aggregate is a finding worth situating against this thesis, which finds the opposite pattern in the heating-dominated PACW system.

What none of these studies provide is a Pacific-Northwest-specific analysis at daily resolution that integrates demand and emissions, that handles a structural change in the generation mix, and that uses a methodology robust enough to project credibly past the historical training range. The closest regional analysis, the Northwest Power and Conservation Council's Seventh Power Plan (2016), is a long-range planning document rather than a peer-reviewed empirical study, and it was published before the 2022 PacifiCorp coal retirement that anchors this thesis.

2.5 Gaps in the Literature and Contribution of This Study

The literature reviewed above leaves three gaps that this thesis addresses directly.

The first is regional. There is limited Pacific-Northwest-specific empirical evidence linking temperature to PACW demand and emissions at fine temporal resolution. The studies that do

exist for the region are either older (predating the 2022 coal retirement), focused on a different system (BPA's federal hydro system rather than PACW's investor-owned territory), or limited to demand without addressing emissions. This thesis fills that gap by analyzing 3,747 daily observations from PACW spanning July 2015 through October 2025, with both demand and emissions modeled jointly.

The second is methodological. Most existing climate-grid projection studies use either pure parametric regressions, which are interpretable but limited in flexibility, or pure machine-learning models, which are flexible but cannot extrapolate past their training range. This thesis develops a hybrid OLS-plus-tree-ensemble architecture specifically designed to handle warming projections that push inputs toward and past the historical edge, combining the extrapolation behavior of OLS on temperature variables with the flexibility of tree-based models on residual variation. The chronological train-test validation procedure, the diagnostic that surfaced an information-leakage problem in earlier random-split versions of the model, and the architectural switch from a multi-step emissions cascade to a direct CO_2 specification are themselves methodological contributions documented in detail in Chapter 4.

The third is relevant to policy. The existing literature has largely treated the question of how warming reshapes grids as a question about demand growth, with emissions treated as a downstream consequence calculated against fixed or slowly-evolving emission factors. This thesis argues that for a system experiencing structural change in its generation mix, which describes most U.S. balancing authorities undergoing decarbonization in the 2020s and 2030s, the structural change itself can dominate the climate-driven changes in magnitude. The 11,950 to 12,790 metric tons of CO_2 per day attributable to the 2022 PACW coal retirement, controlling for demand and weather, is large relative to any plausible warming-driven emissions change in the

same system at near-term warming levels. This finding reframes the policy-relevant comparison: for grids in transition, the pace of structural decarbonization may matter more for emissions outcomes than the pace of climate-driven demand response. Together these gaps motivate a study that is regional, methodologically rigorous, and integrative across demand and emissions. The chapters that follow develop that study in detail.

Chapter 3: Data

3.1 Overview

This thesis draws on two primary data streams: hourly electricity system data for the PacifiCorp West (PACW) Balancing Authority and daily climate data from five stations distributed across Oregon. Together, these datasets support an integrated analysis of how temperature shapes electricity demand and carbon emissions, and how that relationship is likely to evolve under projected climate warming.

The combined dataset covers 3,747 daily observations from July 1, 2015 through October 8, 2025, a span of approximately 10.3 years. This window is significant for two reasons. It captures the June 2021 Pacific Northwest heat dome, which produced record-setting temperatures and provided training data near the upper tail of the observed temperature distribution. It also spans the July 2022 retirement of PacifiCorp's coal generation assets, creating a natural before-and-after comparison of how Oregon's electricity system responds to temperature stress under materially different generation mixes.

Electricity data comes from the U.S. Energy Information Administration's Grid Monitor and includes hourly demand, generation by fuel type, interchange with neighboring balancing authorities, and per-fuel CO_2 emissions. Climate data comes from the PRISM Climate Group at Oregon State University, pulled at five point locations representing the major climate regimes within and adjacent to the PACW service territory: Eugene and Albany in the southern and mid-Willamette Valley, Portland in the lower Columbia region, Bend in the high desert east of the Cascades, and Medford in the Rogue Valley of southern Oregon.

Of these five stations, three sit fully inside the PACW balancing authority: Albany, Bend, and Medford. Each represents climate at points whose electricity load is part of the

demand series being modeled. Eugene is the only station fully outside PACW, served instead by Eugene Water & Electric Board, a municipal utility. However, its climate is extensively comparable to Albany and the greater Willamette Valley. Portland is a mixed case: PACW serves a substantial portion of the metropolitan area through the Pacific Power retail brand, including parts of downtown, northern and eastern Portland, while Portland General Electric serves the remainder.

Eugene and Portland were retained in the analysis despite their partial or full position outside PACW because they capture distinct regional climates and because they were used in the early exploratory phases of the modeling work, before the within-PACW Albany and Medford stations were added. The motivation for the five-station set, and the eventual selection of Albany as the primary climate signal for the production model, is developed in Section 4.10. Hourly electricity data were aggregated to daily resolution and merged with the daily PRISM records on normalized dates, producing a single time-indexed data frame for each climate station.

An initial dataset of Oregon's average retail electricity prices was included in the prospectus but ultimately set aside: modeling retail price pass-through would have required a distinct analytical framework, and I chose to focus the thesis on the demand and emissions questions where the data supported firmer conclusions.

The sections that follow describe each data stream in detail, document the cleaning and alignment decisions that produced the merged analytical dataset, and outline the feature engineering steps that transformed the raw variables into the inputs used for modeling.

3.2 Electricity Data: EIA Grid Monitor (PACW)

Electricity data for this analysis come from the U.S. Energy Information Administration's (EIA) Hourly Electric Grid Monitor, which publishes operational data for all balancing

authorities in the United States via Form EIA-930. A balancing authority (BA) is the entity responsible for maintaining real-time equilibrium between electricity supply and demand within a defined service territory, and for coordinating interchange with neighboring systems. The PacifiCorp West Balancing Authority, designated PACW, is the unit of analysis for this thesis.

3.2.1 The PACW Balancing Authority

PacifiCorp operates two distinct balancing authorities within the Western Electricity Coordinating Council (WECC) footprint: PacifiCorp East (PACE), which serves portions of Utah, Wyoming, and southeastern Idaho, and PacifiCorp West (PACW), which serves portions of Oregon, northern California, and southeastern Washington. The two areas are connected only through a limited amount of third-party transmission, and each operates as an independent balancing authority under North American Electric Reliability Corporation (NERC) standards. PACW corresponds to the operational footprint of PacifiCorp's Pacific Power subsidiary, headquartered in Portland.

Within Oregon, PACW serves most of the state outside the Portland metropolitan area, see Figure 3.2. Portland itself is slightly complicated. The city is served mainly by a separate balancing authority operated by Portland General Electric (PGE). However, PACW does maintain a service area within the eastern and northern portions of Portland. The PacifiCorp service area map reproduced as Figure 3.2 shows this split clearly.

As a system, PACW is mid-sized by U.S. standards. Pacific Power, the utility operating within the PACW territory, reports serving approximately 243 communities and roughly 800,000 customers across the three states. It produces roughly 8.1 TWh of electricity annually, ranking 45th of 65 U.S. balancing authorities by net generation. Its generation portfolio is also considerably cleaner than the national average but relies heavily on imports: hydroelectric

generation accounts for about 10% of the mix, wind 8%, solar 5%, and imports fill the rest. Coal generation, which contributed a meaningful share of the stack before 2022, has since been retired from the PACW system. The territory is interconnected via 500 kV, 230 kV, and lower-voltage transmission lines with a large number of neighboring balancing authorities, most significantly the Bonneville Power Administration (BPA), which operates the backbone transmission network of the Pacific Northwest.

PACW is an appropriate unit of analysis for a climate-and-grid study of Oregon for three reasons. First, the majority of its load and territory lies within Oregon, making it the single best available proxy for a state-level electricity analysis that respects operational grid boundaries. Second, it is a hydro-anchored system with substantial seasonal supply variability, meaning that both demand and emissions responses to temperature reflect the physics of a real Pacific Northwest grid rather than an aggregated national average. Third, the July 2022 retirement of PacifiCorp's coal generation creates a clean structural break in the data, enabling before-and-after analysis of how the same demand patterns produce different emissions under different generation mixes.



Figure 3.2: The PacifiCorp West Service Territory

3.2.2 Variables Retained

The raw dataset contains hourly observations from July 2015 through October 2025 and includes a wide range of operational variables. The variables retained for this thesis fall into three categories. Demand variables include hourly demand, adjusted demand, and demand forecast, along with net generation, total interchange with neighboring balancing authorities, and imputed values for days on which EIA's standard reporting was incomplete. Generation variables disaggregate net generation by fuel source: coal (COL), natural gas (NG), hydro (WAT), solar (SUN), wind (WND), and other (OTH). Emissions variables include total CO_2 emissions for generated, consumed, and imported electricity, along with per-fuel CO_2 emissions for coal,

natural gas, oil, and other sources, and CO_2 emissions intensity for both generated and consumed electricity.

3.2.3 Temporal Aggregation

Hourly data were aggregated to daily resolution to align with the daily climate data and to match the time scale on which temperature-demand relationships are most interpretable. For energy and emissions quantities, the daily aggregation was the sum of the 24 hourly observations (equivalently, the hourly mean multiplied by 24), yielding daily totals in MWh of demand and generation and in metric tons of CO_2 rather than per-hour rates. Peak Demand was computed separately as the daily maximum of hourly demand, preserving information about within-day load extremes that would otherwise be lost in the aggregation.

3.2.4 Data Cleaning

A small number of data quality issues were identified during initial inspection and addressed before modeling. Six specific dates in 2016 were removed entirely due to reporting anomalies in the raw EIA data: January 10, March 2, October 13, and November 15 through 17. On these dates the reported values were clearly erroneous, including zeros and, in some cases, values in the order of millions of megawatts, neither of which has any plausible physical interpretation for a balancing authority of PACW's size. With no documentation available to reconstruct the true values, these days were dropped from the analysis.

Beyond those removals, two threshold-based corrections were applied to the demand series. Daily mean demand values at or above 100000 MW were flagged as implausible and converted to missing. The single lowest observed daily demand value was also flagged and converted to missing on similar grounds. Daily peak hourly demand values at or above 10,000

MW were flagged and converted to missing by the same logic, as 10,000 MW is more than double the all-time system peak. All flagged values were then filled using linear interpolation in both directions, which preserves the continuity of the time series without introducing artificial structure.

After these corrections, the daily PACW dataset contained a continuous record of demand, generation by fuel type, and emissions variables from July 2015 through October 2025, suitable for merging with daily climate data at the city level.

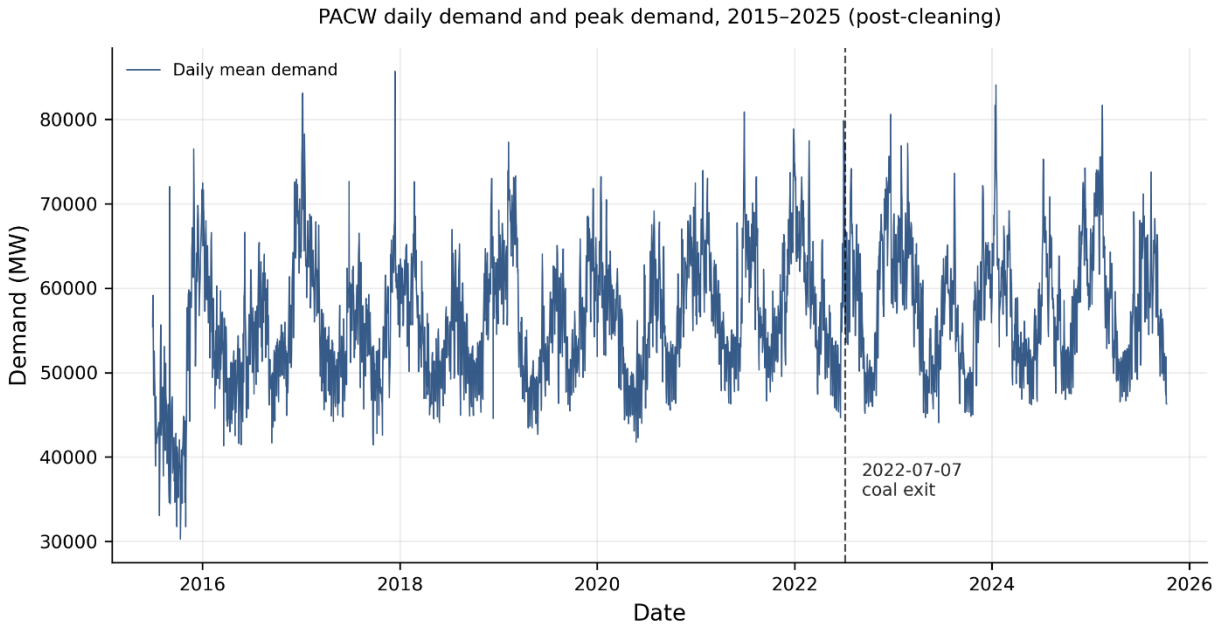


Figure 3.2.1: Daily PACW Demand July 2015 – October 2025

3.3 Climate Data: PRISM

Daily climate data for this analysis come from the PRISM Climate Group at Oregon State University. PRISM, which stands for Parameter-elevation Regressions on Independent Slopes Model (Daly et al., 2008), is a hybrid statistical-geographic climate dataset that produces

interpolated daily values of temperature, precipitation, and other meteorological variables across the continental United States at a spatial resolution of approximately 4 kilometers. Rather than reporting raw observations from individual weather stations, PRISM uses a regression-based interpolation model that incorporates elevation, coastal proximity, terrain orientation, and other physiographic factors to produce a spatially continuous climate surface. The dataset used in this thesis is AN91d, PRISM's daily all-networks product covering 1981 through the present, with a PRISM "day" defined as the 24-hour period ending at 12:00 UTC on the date shown (approximately 4:00 AM PST / 5:00 AM PDT)

3.3.1 Station Selection

PRISM data can be extracted at any set of latitude-longitude coordinates within the continental United States. The selection of extraction points evolved over the course of this thesis. The initial three stations (Eugene, Bend, and Portland) were chosen at the start of the analysis as the population centers representing Oregon's three major climate regimes. As the modeling work progressed, two additional stations (Albany and Medford) were added to address a limitation that became visible during the early exploratory phase: Eugene is served by Eugene Water & Electric Board rather than PACW, and most of Portland is served by Portland General Electric rather than PACW. Albany and Medford were added because they sit unambiguously inside the PACW service territory, and Medford in particular extends the geographic coverage southward into the Rogue Valley, which the original three-station set did not capture. The full station set used in the thesis comprises five locations:

Eugene (44.0408°N, 123.0870°W, elevation 453 ft), representing the southern Willamette Valley climate. Eugene was the primary station for the initial exploratory phase of the modeling work and continues to appear throughout the thesis as a comparison station, but it is served by

Eugene Water & Electric Board rather than PACW. The surrounding mid-Willamette Valley communities of Corvallis, Albany, Lebanon, and Junction City are PACW customers and share Eugene's climate.

Albany (44.6365°N, 123.0959°W, elevation 213 ft), representing the mid-Willamette Valley climate inside the PACW service territory. Albany sits roughly 25 miles north of Eugene at similar elevation and shares the same broad maritime climate, but unlike Eugene it is a PACW customer and represents a city whose load is part of the demand series being modeled. Albany was added later in the modeling process and ultimately becomes the primary climate signal for the production model in Section 5.10.

Bend (44.0465°N, 121.3274°W, elevation 3,717 ft), representing the central Oregon high desert. Bend is inside the PACW service territory and provides a strong contrast to the Willamette Valley stations. Notably, it sits at nearly the same latitude as Eugene but more than 3,200 feet higher in elevation and on the opposite side of the Cascade rain shadow, providing a strong contrast in daily climate conditions along a single east-west transect.

Portland (45.5030°N, 122.6730°W, elevation 105 ft), representing the lower Columbia region and northern Willamette Valley. Portland is largely outside the PACW service territory (most of the Portland metropolitan area is served by Portland General Electric) but PACW does serve an eastern and northern portion of the city, and Portland's climate broadly characterizes the lower Columbia region. Portland was retained in the analysis because it was one of the original three stations and because its climate represents a distinct regional signal.

Medford (42.3265°N, 122.8756°W, elevation 1,382 ft), representing the Rogue Valley in southern Oregon. Medford is inside the PACW service territory, sits roughly 250 miles south of the Willamette Valley load centers, and has a meaningfully warmer and drier climate than the

Willamette stations. Medford was added later in the modeling process to test how a southern-Oregon climate signal would compare to the Willamette Valley stations.

Together, these five points capture the three major climate regimes that define Oregon's weather patterns, the wet maritime west, the drier interior valleys of the south, and the dry continental east, and provide multiple within-PACW stations along the I-5 corridor where PACW load is concentrated. The cross-station comparison developed in Section 5.10 uses all five stations as parallel single-station inputs to the same modeling architecture, with the goal of identifying which station best captures the dispatch dynamics of the system as a whole.

3.3.2 Variables Retained

The PRISM extraction for each of the five locations includes seven daily variables. Five of these were retained for modeling:

- Precipitation (ppt), in inches
- Minimum daily temperature (tmin), in degrees Fahrenheit
- Mean daily temperature (tmean), in degrees Fahrenheit
- Maximum daily temperature (tmax), in degrees Fahrenheit
- Mean daily dew point temperature (tdmean), in degrees Fahrenheit

The two remaining variables, minimum and maximum daily vapor pressure deficit (vpdmin, vpdmax) in hectopascals, were pulled with the rest of the PRISM output but not used in the final analysis. Wet bulb temperature, which serves a similar atmospheric-humidity role and is more directly interpretable in a temperature-demand context, was derived from tmean and tdmean using the Stull (2011) approximation and is described in Section 3.5.

The PRISM data for all five locations cover 1981 to present at daily resolution, giving substantially more historical data than the PACW electricity record. Only the dates overlapping

the PACW window (July 1, 2015 through October 8, 2025) were retained in the merged analytical dataset.

3.4 Data Cleaning and Alignment

With the PACW electricity data cleaned as described in Section 3.2 and the PRISM climate data delivered in a clean, continuous form directly from Oregon State University, the remaining preparation step was to align the datasets on a common temporal index and merge them into a single analytical data frame.

The PACW data began as an hourly time series starting July 1, 2015. After the temporal aggregation and cleaning described in Section 3.2, the working PACW dataset is indexed at daily resolution from July 2015 through October 2025. The PRISM data are delivered at daily resolution from January 1, 1981 through the extraction date, covering a much longer window than the electricity record. Both datasets use local Pacific dates, but the raw indices differed in format: the PACW index carried hour-level timestamps from the original hourly data, while the PRISM index was already daily. To ensure that dates aligned cleanly across sources, all five PRISM indices and the daily PACW index were explicitly cast as pandas datetime objects and normalized to midnight, producing a consistent daily key that matched exactly across datasets.

Each PRISM location was then merged separately onto the daily PACW dataset using a left join on the normalized date index. This produced five parallel merged data frames, one per climate station, each retaining all PACW rows and adding the corresponding daily climate variables from Eugene, Albany, Bend, Portland, or Medford. The left join preserved the electricity record as the anchor and ensured that no PACW observations were dropped during the merge. Because PRISM's record begins well before the PACW record, every PACW date found a matching PRISM row, so no climate values were missing on merge.

Finally, the merged datasets were trimmed to the shared analytical window of July 1, 2015 through October 8, 2025. This produced five daily data frames, each containing 3,747 observations across approximately 36 columns of electricity, generation, emissions, and climate variables. These merged data frames form the foundation for all subsequent feature engineering and modeling, described in the following sections.

3.5 Feature Engineering

Before modeling, a set of derived variables was constructed to capture temperature nonlinearities, humidity effects, calendar and behavioral patterns, and the structural shift associated with PacifiCorp's 2022 coal retirement. These engineered features complement the raw electricity and climate variables described in earlier sections and serve as direct inputs to both the OLS and the Random Forest / XGBoost residual models.

Temperature transformations. Electricity demand is well known to respond nonlinearly to temperature, rising at both hot and cold extremes relative to a comfort minimum. To capture this curvature within a linear regression framework, a temperature-squared variable was created from the daily mean temperature:

$$tmean^2 = tmean * tmean$$

Together with tmean entered linearly, this quadratic term allows the OLS model to fit the U-shaped demand-temperature response that is characteristic of residential and commercial electricity use.

Heating and cooling degree days. Two additional temperature transformations, heating degree days (HDD) and cooling degree days (CDD), decompose the daily temperature signal into

its heating-side and cooling-side components. Both are computed relative to a base temperature T_base:

$$HDD = \max(0, T_{base} - tmean)$$

$$CDD = \max(0, tmean - T_{base})$$

On any given day, at most one of HDD or CDD is nonzero. HDD accumulates "demand pressure" on cold days, CDD accumulates it on hot days, and together they provide a piecewise-linear representation of the demand-temperature relationship that complements the smooth quadratic term above.

The choice of base temperature warrants explicit justification. The standard base temperature used throughout the HVAC and residential energy literature is 65°F, a value chosen historically to approximate a thermal comfort threshold for U.S. building stock. For this thesis I set the base temperature at 58°F rather than 65°F. The choice was empirical: when I fit a simple OLS model of daily demand on daily mean temperature across the full PACW dataset, the fitted curve reached its minimum at approximately 58°F. In other words, PACW demand is lowest, and heating and cooling pressure balance, at a cooler temperature than the national convention assumes. This is consistent with the character of the PACW service territory, a Pacific Northwest climate with historically limited air conditioning penetration where cooling demand begins to rise at relatively cool outdoor temperatures. Setting the base at 58°F aligns the HDD and CDD variables with the demand minimum actually observed in the data rather than with a nationally generic threshold.

Wet bulb temperature. To capture the role of atmospheric humidity alongside temperature, wet bulb temperature was computed for each day from the daily mean temperature

(tmean) and the daily mean dew point temperature (tdmean) using the Stull (2011) approximation. The Stull formulation first estimates relative humidity from the temperature and dew point pair, then calculates wet bulb temperature using a closed-form empirical expression validated against psychrometric charts across a wide range of meteorological conditions. The underlying computation is performed in Celsius, with Fahrenheit values converted in and out as needed.

Wet bulb temperature is worth including because it captures something that dry bulb temperature alone misses: how hot the air actually feels to a human body and how hard a cooling system has to work. Physically, the wet bulb temperature is the lowest temperature that can be reached by evaporating water into the air at constant pressure. When humidity is low, sweat and other moisture evaporate quickly, efficiently carrying heat away from the body and from building interiors. A 90°F day in Bend's dry high-desert air is uncomfortable but manageable because evaporation does most of the cooling work. When humidity is high, evaporation slows down or stops, so the body cannot shed heat effectively and air conditioners have to remove both sensible heat and moisture from indoor air in order to maintain comfort. A 90°F day in a humid maritime climate places meaningfully more stress on both people and the electricity grid than a 90°F day in a dry one, and that difference shows up in cooling demand. Because the PACW territory spans both dry continental (Bend) and humid maritime (Eugene, Portland) climates, including a humidity-aware variable alongside dry bulb temperature gives the residual models a sharper view of the thermal conditions actually driving demand.

Wet bulb temperature is correlated with the underlying dry bulb and dew point variables from which it is computed, but it is retained in the feature set because it packages the interaction

between the two into a single physiologically meaningful variable that tree-based models can exploit directly.

Calendar variables. Three calendar-based variables were extracted directly from the date index. The variable month takes integer values from 1 through 12 and captures seasonal patterns in both demand and grid supply, for example the spring snowmelt period that affects hydropower availability independently of temperature. The variable day_of_week takes integer values from 0 (Monday) through 6 (Sunday) and captures weekly cycles in industrial and commercial electricity consumption. The variable weekend is a binary indicator equal to 1 on Saturdays and Sundays and 0 otherwise, capturing the sharp drop in non-residential load on weekends.

These calendar variables primarily serve the Random Forest and XGBoost residual models, which use them to capture the behavioral and seasonal structure in demand that the temperature variables alone cannot explain.

Structural break indicator. Finally, a binary structural break variable, post_july_2022, was constructed to mark the retirement of PacifiCorp's coal generation assets. The indicator takes a value of 1 for all dates on or after July 7, 2022, and 0 for all prior dates. This variable enters both the demand model, where its effect is modest, and the direct CO_2 emissions model, where it is the single strongest predictor in the regression. It is central to the before-and-after comparison of grid carbon emissions developed in Chapters 5 and 6.

3.6 Exploratory Data Analysis

Before any modeling, I inspected the merged dataset to confirm that variables behaved sensibly, to identify the dynamic range over which later models would be fit, and to surface any patterns or anomalies worth noting in the data itself. Table 3.6 summarizes the key variables

used in modeling across the full analytical window of July 1, 2015, through October 8, 2025. Statistics are reported for the Eugene-merged dataset, which was the primary input to the single-station demand and emissions models as they were developed, explained in Chapter 5. Albany, Bend, Portland, and Medford climate variables are comparable in structure though different in magnitude, and those stations are referenced in the cross-city comparison.

Table 3.6: Summary statistics for the merged PACW-Eugene daily dataset (n = 3,747 days, July 2015 through October 2025).

Variable	**Units**	**Min**	**Mean**	**Max**	**Std.**
Demand (*daily mean*)	MW	30,264.0	56,640.96	85,697.04	7,601.76
Peak Demand (*daily max*)	MW	1261.0	2785.36	5446.00	428.77
Net Generation	MW	-11,798.88	33,379.44	258,850	18,346
CO_2 Emissions (*Consumed*)	Metric tons	344.6	14,069.8	35,093.5	7,937.5
CO_2 Emissions Intensity	lbsCO_2/MWh	0.01	0.54	1.27	0.30
Mean Temperature (*tmean*)	°F	24.60	53.56	87.10	11.40
Minimum Temperature (*tmin*)	°F	14.50	42.62	65.50	8.54
Maximum Temperature (*tmax*)	°F	27.50	64.50	108.90	15.31
Mean Dew Point (*tdmean*)	°F	13.70	44.81	64.20	7.68
Precipitation (*ppt)*	Inches	0.00	0.10	2.09	0.22
Wet Bulb Temperature	°F	20.30	48.25	71.38	8.70
HDD (base 58°F)	°F-days	0.00	7.43	33.40	7.77
CDD (base 58°F)	°F-days	0.00	2.99	29.10	5.03

Several patterns in the data merit note. Daily mean demand varies over a meaningful range across seasons, reflecting PACW's combined heating and cooling loads against a backdrop of an underlying baseline driven by non-thermal consumption. Temperature variability at Eugene is typical of the southern Willamette Valley, with winter daily means in the 30s and 40s (°F), summer daily means in the 60s and 70s, and occasional extreme days above 90°F including the late-June 2021 Pacific Northwest heat dome event, which produced the highest temperatures in the historical record for the region. HDD observations substantially outnumber CDD observations in both count and magnitude, confirming that the PACW climate is heating-dominated on an annual basis and that heating demand drives more of the observed temperature-load relationship than cooling demand does. CO_2 emissions intensity shows a visible regime shift around mid-2022, with consumed-electricity intensity falling substantially following the retirement of PacifiCorp's coal generation. This structural break is formally discovered and modeled in Chapter 5 via the post_july_2022 indicator and explored quantitatively in Chapter 6.

Chapter 4: Methods

4.1 Overview

This chapter describes the methodology used to model how electricity demand and total CO_2 emissions on the PacifiCorp West (PACW) grid respond to temperature, and how that response is projected to behave under warming scenarios. The chapter is intended to be self-contained: a reader should be able to reproduce the analysis from the descriptions provided here, and the data documentation in Chapter 3.

The analytical framework consists of two hybrid models, run sequentially. A hybrid demand model predicts daily PACW demand from temperature variables. A hybrid direct-CO_2 emissions model predicts daily total CO_2 emissions from demand, temperature, and a structural break indicator marking the 2022 PacifiCorp coal retirement described in Chapter 2. Under warming scenarios, predicted demand from the first model feeds the second model as an input, producing joint demand-and-emissions projections across a range of temperature increases. Each of the two models is built from the same underlying architecture: an ordinary least squares (OLS) regression captures the smooth, physically interpretable component of the relationship, and a tree-based machine learning model captures the residual variation that OLS misses. The full prediction is the sum of the two layers, explained in further detail below.

4.2 The Hybrid Modeling Architecture

Both the demand model and the emissions model in this thesis are built on the same hybrid architecture: an ordinary least squares (OLS) regression captures the smooth, parametric component of the relationship, and a tree-based machine learning model, Random Forest (RF) or

XGBoost (XGB), captures the residual variation that OLS cannot. The two layers are fit independently and combined additively at prediction time:

Prediction = (OLS prediction from temperature features) + (ML prediction of OLS residual from non-temperature features)

The motivation for this design is the central role of temperature extrapolation in the warming scenarios that are the eventual deliverable of the analysis. Tree-based machine learning models, including Random Forest and XGBoost, predict by routing inputs through a hierarchy of binary splits learned during training and aggregating the values stored in terminal leaves. When a new observation falls outside the range of any feature seen during training, the model has no terminal leaves corresponding to inputs beyond that range and instead returns the value learned at the closest training-edge leaf. The result is that tree-based predictions flatten into a constant level once inputs exceed the historical range, regardless of how far past the edge the inputs go.

For warming scenarios that push temperatures toward and past the upper edge of the PACW historical envelope, a pure tree-based model would suppress precisely the U-curve continuation that the scenarios are designed to characterize. OLS, by contrast, is a parametric model: its prediction at any input is a closed-form function of the fitted coefficients, and that function continues to behave smoothly when inputs move past the training range. A quadratic OLS model fit to demand on temperature continues to produce a U-shaped response curve at temperatures it has never seen, because the curve is mechanically generated from the fitted coefficients rather than from interpolation among historical observations.

The hybrid architecture is therefore a careful matching of model class to feature role. Temperature is the variable that warming scenarios push outside the historical envelope, so temperature features are placed in the OLS layer where extrapolation is well-behaved. Calendar

variables, humidity, precipitation, and the structural break indicator stay within their historical ranges in every scenario, so they can safely use the more flexible ML layer. The two layers are trained sequentially: OLS is fit first on the target variable directly, residuals are computed, and the ML model is then fit on the residuals using the non-temperature feature set. At prediction time the two outputs are summed.

The same architecture, hyperparameter settings, train-test split design, and residual feature set are applied to both the demand model and the emissions model, with the OLS feature partition differing between the two as documented in Sections 4.4 and 4.5.

4.3 Demand Model Specification

The hybrid demand model predicts daily PACW demand using the architecture described in Section 4.3. The OLS layer uses four temperature-derived variables, and the residual layer uses the remaining weather, calendar, and structural-break features. Table 4.3 shows the partition of features between the two layers.

Table 4.3: Feature partition for the hybrid demand model.

OLS Layer (*temperature*)	**Residual ML Layer** (*everything else*)
tmean	tmax
temp_squared	tmin
CDD (base 58°F)	tdmean
HDD (base 58°F)	wet_bulb_F
	precipitaion (*ppt*)
	day_of_week
	month

	weekend
	post_july_2022

The OLS layer carries the variables that the warming scenarios push outside the historical envelope. The four temperature variables in the OLS layer are not independent of one another; temp_squared is a deterministic function of tmean, and HDD and CDD are piecewise functions of the same, meaning the resulting design matrix has known multicollinearity. This is accepted as a cost of the specification because the goal of the OLS layer is not to identify clean individual coefficients but to produce a smooth, extrapolable temperature response curve. Multicollinearity affects the interpretation of individual OLS coefficients, not the predictive behavior of the fitted curve as a whole.

The residual layer receives nine features. Three of them (tmax, tmin, tdmean) are additional temperature-spread variables that were not used in the OLS layer; they could in principle have been added there but would have introduced still more multicollinearity without changing the smooth U-shape that the OLS curve is meant to capture. Placing them in the residual layer instead lets the ML model exploit their information about within-day temperature variability and humidity without compromising the OLS layer's interpretability. Wet bulb, precipitation, the three calendar variables, and the structural-break indicator complete the residual feature set.

Hyperparameters. Both residual models use moderately conservative settings, fixed at their values rather than tuned through cross-validation. The Random Forest residual model uses 300 trees and a minimum of three samples per leaf. The XGBoost residual model uses 300 estimators, a learning rate of 0.05, a maximum tree depth of 4, 80% row subsampling, and 80%

column subsampling. Both use a random state of 42 for reproducibility. The choice not to tune is acknowledged as a limitation in Section 4.9. The values were chosen as reasonable defaults from the broader applied machine learning literature, conservative enough to limit overfitting on a dataset of fewer than 4,000 daily observations, and consistent enough across model families to allow direct comparison of RF and XGB performance on the same underlying task.

4.4 Emissions Model Specification

The hybrid emissions model predicts daily total CO_2 emissions for consumed electricity using the same architecture and the same residual-layer feature set as the demand model, but with a different OLS specification that includes predicted demand as a regressor.

Direct prediction of total CO_2. The OLS layer is:

$$CO_2_predicted = \beta_0 + \beta_1 \cdot Demand + \beta_2 \cdot tmean + \beta_3 \cdot temp_squared + \beta_4 \cdot CDD + \beta_5 \cdot HDD + \beta_6 \cdot post_july_2022 + \varepsilon$$

Demand enters the OLS layer directly as a predictor alongside the four temperature-derived variables and the structural-break indicator. The residual layer uses the same nine non-temperature features as the demand model: tmax, tmin, tdmean, wet_bulb_F, precipitation, day_of_week, month, weekend, and post_july_2022.

This is a direct specification rather than a cascade. An earlier version of the emissions model was structured as a two-stage cascade in which a separate hybrid model predicted CO_2 emissions intensity (lbs CO_2 per MWh), and total daily CO_2 was recovered by multiplying predicted intensity by predicted demand. The cascade was conceptually appealing because intensity has a clean physical interpretation as the dirtiness of the marginal megawatt-hour, but it produced problems in practice. The multiplicative architectures compound errors: an error of

10% in predicted demand and an error of 10% in predicted intensity could combine into more than 20% error in total CO_2, with no single error stream against which the architecture could be jointly validated. The direct specification eliminates the multiplicative compounding, produces a single error stream, while allowing the models to be more interpretable.

Non-negative clipping. The hybrid CO_2 predictions occasionally produce small negative values on low-output days, when the residual layer drives the prediction below zero. CO_2 emissions cannot be negative, and final predictions are clipped at zero before being reported or fed into scenario calculations. Clipping affects fewer than 2% of test-set rows and changes scenario means by less than one percent across the warming sweep, confirming that it is a defensible physical-realism constraint rather than a substantive distortion of the predictions. Both clipped and unclipped predictions are computed during scenario runs for transparency.

Hyperparameters. Identical to the demand model: 300 trees with minimum three samples per leaf for Random Forest; 300 estimators, learning rate 0.05, max depth 4, and 80% row/column subsampling for XGBoost. Both use a random state of 42.

4.5 Train-Test Validation and Diagnostics

Both hybrid models are evaluated using a chronological 80/20 train-test split rather than a random one. The merged dataset for each station is sorted by date, and the first 80% of observations form the training window while the final 20% form the test window. The OLS layer is fit on the training window, residuals are computed, and the residual ML layer is fit on the same training window using those residuals as the target. Predictions on the test window are evaluated using R^2, RMSE, and MAE, computed against the actual values of demand or total CO_2 in the test rows. The same procedure is applied identically to all five climate stations.

For the emissions model specifically, the training window starts later than the demand model's because EIA's CO_2 Emissions Consumed series begins reliably reporting in mid-2018; rows with missing CO_2 values are dropped before splitting, which shortens the available training history. All other elements of the validation procedure are identical between the two models.

Why chronological splitting? Typically, models are trained on a random split of the existing data. The implicit assumption behind random splitting is that observations are exchangeable, that any random subset of the data is statistically representative of any other, and that no systematic structure is lost when rows are reshuffled. For a time series with a structural break, this assumption does not hold. The PACW dataset contains the July 2022 coal retirement, which produces a sharp regime change in the emissions series and a smaller but real shift in some demand-side patterns. A random split places pre-coal-exit and post-coal-exit observations into both training and test sets, allowing the model to identify which regime each test row belongs to via the post_july_2022 indicator and predict accordingly. The structural break, which should be a forward-projection challenge for the model, becomes a free pattern-recognition cue that inflates the apparent test R^2 without delivering corresponding forward-projection skill.

An earlier version of this analysis used a random 80/20 split and reported strong-looking emissions test R^2 values that turned out to be substantially inflated. Refitting the same emissions architecture under a chronological split, with the structural break entirely contained inside the training window, dropped the test R^2 from approximately 0.77 to approximately 0.25, most of the apparent fit was leakage. The same mechanism affected the demand model in principle, but its consequences there were small because demand changes very modestly across the coal exit. The chronological split is therefore the production validation strategy for both models. Beyond catching the leakage problem, chronological splitting matches the structure of the warming

scenarios: the models are being asked to project forward in time under hypothetical inputs, and a forward-validation procedure is the only honest way to test that capability on historical data.

Diagnostic metrics reported. For each model and each station, the production validation reports the train R^2, test R^2, train RMSE, and test RMSE. The train-test R^2 gap is reported as an overfitting diagnostic: a small gap (under approximately 0.05) indicates that the model generalizes well to held-out data, while a large gap (over approximately 0.15) indicates substantial overfitting to training-window patterns. XGBoost is preferred over Random Forest as the production residual model in both the demand and emissions specifications because XGBoost consistently produces tighter train-test gaps on this dataset, suggesting better generalization. The RF results are retained as comparison baselines and reported alongside the XGBoost numbers throughout Chapter 5.

Feature importances. For each XGBoost residual model, the relative importance of each residual-layer feature is computed using XGBoost's default gain-based importance metric. These importances describe which non-temperature features the residual model relied on most to explain the variation that OLS missed; they do not describe the contribution of the OLS layer's temperature features, which are reported separately as OLS coefficient estimates. Feature importances and OLS coefficients are both reported in Chapter 5.

4.6 Warming Scenario Construction

The hybrid demand and emissions models produce predictions on the historical dataset. To project forward, the same models are evaluated on a modified version of the dataset in which temperature variables are uniformly shifted upward by a specified warming amount. Each scenario produces a paired set of predictions, one for daily demand and one for daily total CO_2

emissions, across all post-coal-exit days in the modeling window. The warming scenarios sweep from 0°F to +15°F in 0.5°F increments.

Procedure. For each warming level ΔT, every observation in the analytical dataset receives the same upward shift in its temperature variables: tmean, tmax, tmin, and tdmean are each increased by ΔT degrees Fahrenheit. Three derived quantities are then recomputed from the shifted inputs rather than shifted directly. Wet bulb temperature is recomputed from the shifted dry bulb and dew point using the Stull (2011) approximation, preserving the physical relationship between humidity and temperature that a flat shift would have broken. The temperature-squared term is recomputed from the shifted tmean. HDD and CDD are recalculated against the 58°F base from the shifted tmean. Calendar variables, precipitation, and the structural break indicator are unchanged across scenarios. The hybrid demand model is then applied to the shifted dataset to produce a predicted demand series under that warming level. The predicted demand series feeds the direct CO_2 model as the value of the Demand variable, alongside the shifted temperature variables and the structural-break indicator, to produce a predicted total CO_2 series for the same warming level. Predictions are clipped at zero. Resulting predictions are summarized at each warming level by their mean, 95th percentile, and maximum. The same procedure is applied identically across all five climate stations.

Anchoring to climate science. The mapping from warming-axis values to plausible 21st-century outcomes uses the shared socioeconomic pathway (SSP) scenario bands reported in the OCCRI Seventh Oregon Climate Assessment, which projects mean annual warming for the Pacific Northwest through 2050 of approximately 2.0–3.5°F under SSP2-4.5, 3.5–5.0°F under SSP3-7.0, and 4.5–6.5°F under SSP5-8.5. Through 2100, the upper bound under the most aggressive SSP5-8.5 pathway reaches approximately +14°F. These bands are overlaid on the

warming-axis scenario figures in Chapter 5 to anchor the interpretation. Scenarios beyond +10°F should be read as structural tests of model behavior at and slightly past the upper edge of OCCRI's projected envelope, not as forecasts of likely outcomes. The forecast of Oregon's climate and the SSP bands is visible in the figure below, taken directly from the OCCRI Seventh Annual Oregon Climate Assessment.

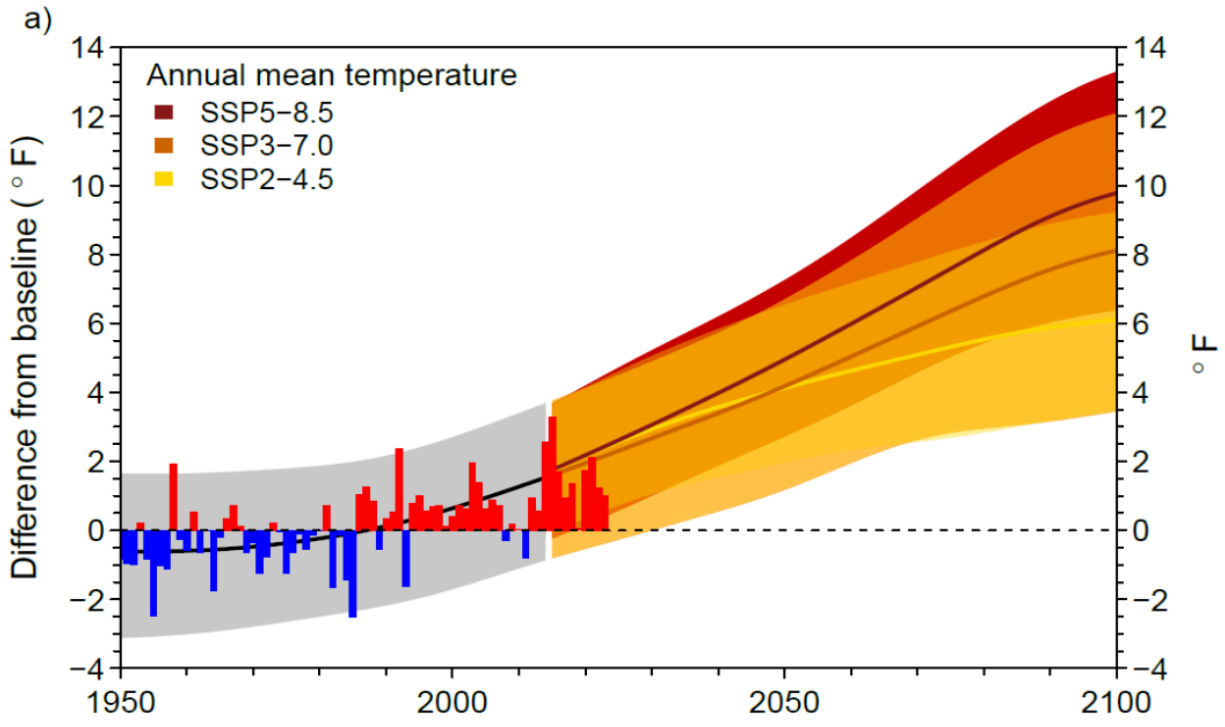


Figure 4.6: Climate Projections from the OCCRI Seventh Annual Oregon Climate Assessment

Scope and what could be done further. The scenarios in this thesis simulate uniform shifts in annual mean temperature: every day in the input dataset has its temperature variables shifted by the same fixed amount. Real warming is not a uniform additive shift. It is more pronounced in some seasons than others, increases the frequency of extreme heat events

disproportionately to mean changes, alters the distribution of daily temperature ranges, and interacts with precipitation patterns in ways that affect hydroelectric dispatch directly. For the scope of this thesis, I focus on uniform annual mean shifts as a way to characterize the structural response of the PACW system to temperature, not as a forecast of what particular days or years will look like under specific warming pathways. Future extensions of this framework, discussed in Chapter 6, are well-positioned to incorporate additional dimensions of warming.

4.7 Cross-City Comparison

The hybrid demand model and the direct-CO_2 emissions model were each fit independently using five different climate stations as the temperature signal: Eugene, Albany, Bend, Portland, and Medford. The motivation for fitting parallel single-station models rather than a single combined multi-station model is partly methodological and partly geographic. Methodologically, single-station specifications are more interpretable than multi-station ones: each model's OLS coefficients have a clean reading as the response of PACW demand or emissions to a single, geographically coherent climate signal, and feature importances in the residual layer are not muddied by correlated weather signals across stations. Geographically, no single station perfectly represents the load-weighted climate of the PACW service territory, and running parallel fits at multiple stations reveals which signals best track the dispatch dynamics of the system as a whole.

Each model fit follows the set file architecture, feature partition, hyperparameter settings, train-test split design, warming scenario procedure, and clipping rule. The only thing that varies between fits is the source of the climate variables. For each station's merged dataset, the same temperature-side variables (tmean, tmean^2, HDD, CDD against the 58°F base, wet bulb, tmax, tmin, tdmean) are constructed from that station's PRISM data and fed into an otherwise

identical demand-and-emissions pipeline. The OLS layers receive the same four temperature-derived variables for demand and the same six demand-and-temperature features for emissions. The residual layers receive the same nine non-temperature features in both models. Random Forest and XGBoost residual models use the same fixed hyperparameters (300 trees, three samples per leaf for RF; 300 estimators, learning rate 0.05, max depth 4, 80% subsampling for XGB). The chronological 80/20 split is applied identically. The warming scenario sweep runs from 0°F to +15°F in 0.5°F increments with post_july_2022 set to 1 across all rows. The clipping rule is identical: final CO_2 predictions are clipped at zero.

This template-based approach makes the five-station fits directly comparable. Any difference in performance, coefficient values, or warming scenario behavior across the five fits is attributable to the underlying climate signal at that station, not to any methodological choice that would otherwise vary between specifications. Table 4.7 reports the production hybrid XGBoost performance for both the demand model and the direct-CO_2 model across all five stations

Table 4.7. Hybrid XGBoost performance across the five climate stations, time-aware split.

City	Demand Train R^2	Demand Test R^2	Demand Test RMSE	CO_2 Train R^2	CO_2 Test R^2	CO_2 Test RMSE
Eugene	0.842	0.824	120.1	0.860	0.330	129.3
Albany	**0.840**	**0.840**	**120.9**	**0.858**	**0.516**	**105.2**
Bend	0.823	0.813	123.6	0.850	0.409	121.4
Portland	0.837	0.813	123.8	0.859	0.429	119.3
Medford	0.817	0.788	139.3	0.857	0.497	107.1

All five stations produce demand models that are substantially better than the temperature-only OLS baseline reported in Section 4.6, with hybrid XGBoost test R^2 values in the 0.79 to 0.84 range across the set. The cross-city spread on demand performance is modest, reflecting that the broad shape of the temperature–demand relationship is similar across Oregon's climate regimes and that the hybrid architecture is reasonably robust to climate-station choice for this target. The emissions models show wider cross-city variation: hybrid XGBoost test R^2 ranges from 0.330 (Eugene) to 0.516 (Albany). The variation reflects how well each station's climate signal captures the dispatch-mix dynamics that drive marginal emissions on the PACW grid, and it is the more diagnostic of the two metrics for selecting a primary specification.

Albany emerges as the strongest specification on both targets. It produces the highest demand test R^2 (tied with Eugene at 0.840) and the highest CO_2 test R^2 (0.516), the lowest CO_2 test RMSE (2,524 mT/day), and the smallest train-test gap on both models (zero gap on demand, 0.342 on CO_2). It is also the only one of the five stations that combines best-in-class statistical performance with full geographic placement inside the PACW service territory. Albany sits in the mid-Willamette Valley, surrounded on every side by PACW customer communities, and its climate signal pairs with electricity load that is genuinely part of the demand series being modeled. For these reasons Albany is adopted as the primary climate signal for the production demand and emissions models. The remaining four stations (Eugene, Bend, Portland, Medford) are retained as cross-station robustness checks, and the substantive differences in their warming scenario outputs are reported and interpreted in Chapter 5.

4.8 Methodological Caveats

Hyperparameters were fixed rather than tuned. The Random Forest and XGBoost residual models use reasonable default settings drawn from common applied machine learning

practice rather than values selected through formal cross-validation. Tuning these settings could plausibly improve test performance modestly, although the relatively tight train-test gaps observed for XGBoost across all five stations suggest that the unTuned defaults are not severely overfitting and that gains from tuning would likely be incremental rather than transformative.

Cross-validation beyond the chronological split was not performed. Each model is evaluated on a single chronological test window. A more robust validation approach would use rolling-origin or expanding-window cross-validation, in which the model is trained on successively earlier windows and tested on successively later ones, producing a distribution of test metrics rather than a single point estimate. The single-split design is a deliberate simplification that prioritizes interpretability and reproducibility over uncertainty quantification.

Warming scenarios apply a uniform additive shift to temperature variables. As discussed in Section 4.7, this approach captures the structural response of the system to mean temperature change but abstracts away from seasonal, year-to-year, and event-level variability that real climate change produces. Coupling the framework to a downscaled climate ensemble would produce richer projections at the cost of substantially more methodological infrastructure and is identified as a future-work direction.

The post-coal training window for the emissions model is short. Because EIA's CO_2 Emissions Consumed series begins reliably reporting in mid-2018 and the structural break occurs in mid-2022, the production emissions model is trained on roughly four years of pre-coal-exit data and only a fraction of a year to a year of post-coal-exit data within the training window itself. Test-set performance on the post-coal regime is therefore evaluated against a relatively small training history, and modest test R^2 values for the emissions model (in the 0.33 to 0.52

range across stations) reflect this constraint as much as they reflect any limitation of the architecture. As more post-coal-exit data accumulates over time, this limitation will ease.

Multicollinearity within the OLS layers is acknowledged. The four temperature-derived OLS variables in the demand specification, and the six demand-and-temperature variables in the emissions specification, are not statistically independent. The resulting design matrices have high condition numbers, which affects the precision of individual coefficient estimates but not the predictive behavior of the fitted curves used in scenario projection. Coefficient values are reported in Chapter 5 with appropriate caveats.

Despite these limitations, the methodology delivers what it was designed to deliver: a reproducible, interpretable, and forward-validated framework for characterizing how PACW demand and total CO_2 emissions respond to temperature, and how that response is likely to scale under plausible climate warming. The empirical findings produced by applying this framework are the subject of Chapter 5.

4.9 Evaluation Summary

This chapter has documented a reproducible methodology for projecting electricity demand and total CO_2 emissions on the PACW grid under hypothetical warming scenarios. The framework consists of two hybrid models, one for demand and one for direct total CO_2, each combining an OLS layer that captures the smooth temperature response with a tree-based residual layer that captures behavioral, calendar, and structural-break variation. Both models are validated using a chronological 80/20 train-test split, and both are applied identically across five Oregon climate stations to produce the cross-station comparison reported in Section 4.8. Albany is adopted as the primary climate signal for the production specification, with the remaining four stations retained as robustness checks.

A few honest acknowledgments about what this methodology does not provide are warranted before turning to the results in Chapter 5.

Chapter 5: Results

5.1 Overview

This chapter reports the empirical findings produced by applying the methodology developed in Chapter 4 to the merged PACW dataset across all five climate stations. The findings are presented as numbers, coefficients, and figures, with interpretation deliberately deferred to Chapter 6. Section 5.2 reports the demand model results. Section 5.3 reports the emissions model results, including the coefficient estimates and production performance metrics. Section 5.4 reports the warming scenarios for both demand and emissions across all five cities. Section 5.5 isolates the 2022 coal-exit coefficient as a headline finding in its own right. Section 5.6 summarizes.

5.2 Demand Model Findings

The hybrid demand model was fit on each of the five climate-station merged datasets using the architecture and chronological 80/20 split described in Chapter 4. Figure 5.2 projects each of the five cities' demand model findings under the constructed warming scenario described in the previous chapter.

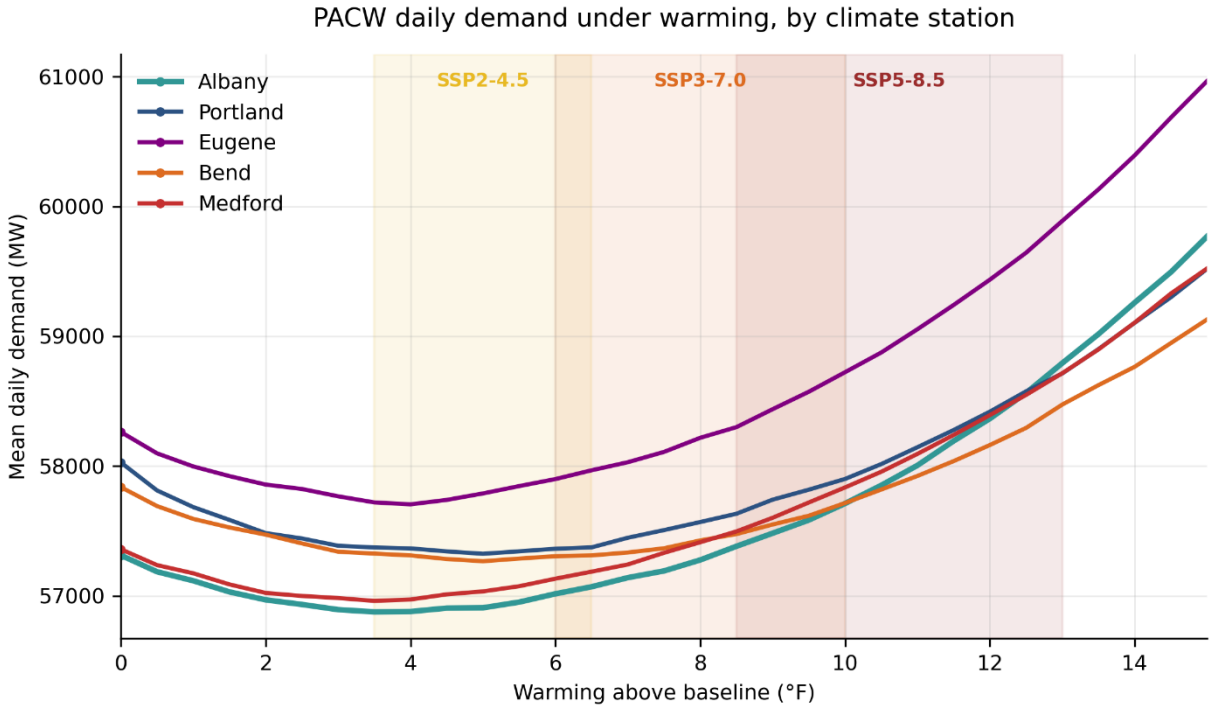


Figure 5.2: Demand Under Warming

Each climate station depicts a similar demand response. Initially, demand drops until reaching an inflection point at 3.5-4°F, and then demand rises consistently throughout additional warming. The feature coefficients at each stage of modeling depict these curves precisely.

OLS feature importances. The most notable and telling of the features from the OLS model are the HDD and CDD variables. The HDD coefficient is positive, large, and highly significant across all five cities, ranging from +64 MW per heating-degree-day at Bend to +85 MW per heating-degree-day at Eugene and Medford. The CDD coefficient is negative across four of five cities, ranging from −17 MW at Albany to −40 MW at Medford, and is approximately zero at Portland. The combination of large positive HDD and modest-to-negative CDD coefficients quantifies the heating-dominated character of PACW demand: cold days drive demand up sharply, while hot days drive it up only modestly in comparison.

Residual feature importances. XGBoost residual-layer feature importances are dominated by calendar variables across all five cities. The weekend indicator alone accounts for between 47% (Bend) and 63% (Medford) of residual importance, followed by day_of_week at 14–17%. The remaining importance is distributed across the temperature-spread variables (tmax, tmin, tdmean, wet_bulb_F), precipitation, and the post_july_2022 indicator, with post_july_2022 contributing 4–7% of residual importance for the demand model. The dominant role of weekend and day-of-week effects in the residual layer indicates that, beyond temperature, behavioral patterns in industrial and commercial consumption are the primary explanatory variables the model relies on.

5.3 Emissions Model Findings

The hybrid emissions model was fit on each of the five climate-station merged datasets using the architecture and chronological 80/20 split described in Chapter 4, and here the OLS layer of the emissions model produces coefficient estimates that describe how demand and weather variables jointly predict total CO_2. Figure 5.3 depicts each of the five cities' emissions model findings under the same warming present in the demand model projections.

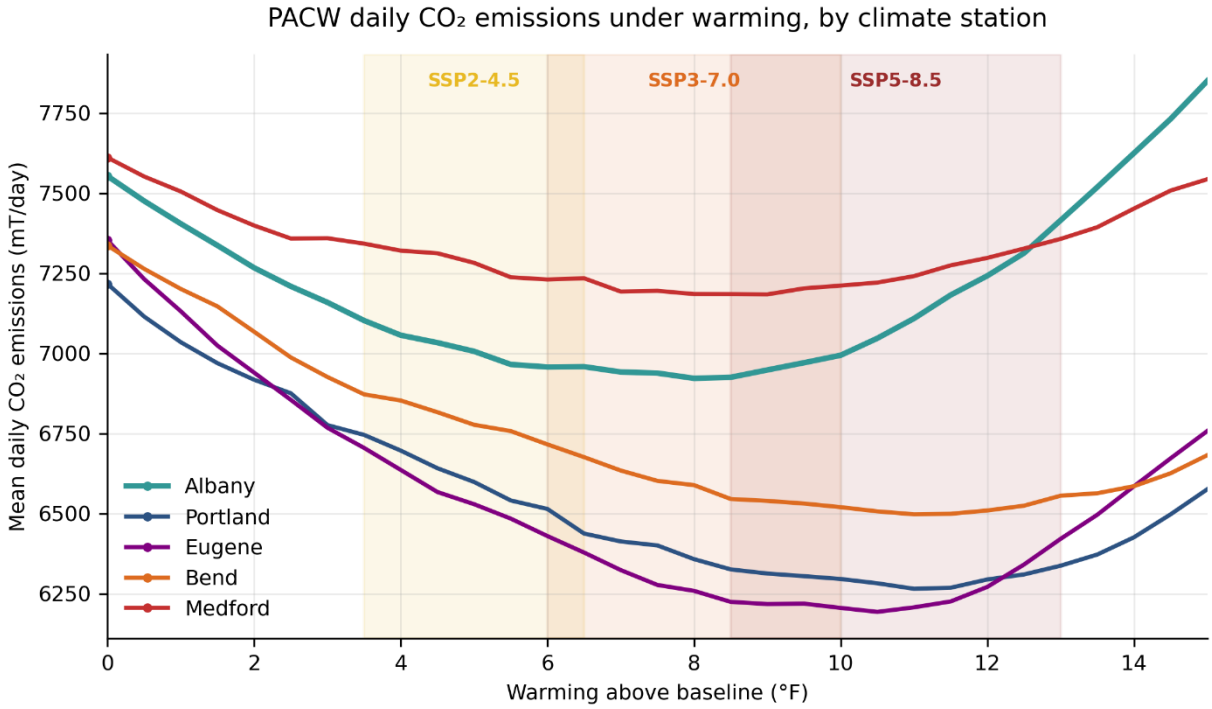


Figure 5.3: Emissions Under Warming

Again, each city depicts a similar emissions response to warming. Initially, emissions decline similar to the demand response, however they continue to do so until a much later inflection point at 9-10°F of warming. The feature coefficients at each stage of modeling depict these curves precisely. Recall the OLS specification of the emissions model:

$$CO_2_predicted = \beta_0 + \beta_1 \cdot Demand + \beta_2 \cdot tmean + \beta_3 \cdot temp_squared + \beta_4 \cdot CDD + \beta_5 \cdot HDD + \beta_6 \cdot post_july_2022 + \varepsilon$$

OLS feature importances. Most notably, the Demand coefficient is positive, small, and highly significant across all five cities, ranging from +0.26 metric tons CO_2 per MWh (Albany) to +0.37 metric tons CO_2 per MWh (Medford). The HDD coefficient is positive across all five cities and reaches statistical significance in four of them (Eugene, Albany, Bend, Portland), with

magnitudes ranging from +35 mT/day (Medford) to +394 mT/day (Eugene). The CDD coefficient is negative in four of five cities (Eugene, Albany, Bend, Medford) and positive at Portland, with statistical significance only at Bend.

Test R^2 ranges from 0.330 (Eugene) to 0.516 (Albany). The cross-city spread is wider than the demand model as discussed in Chapter 4 and is the more diagnostic metric for selecting Albany as the primary climate signal in Chapter 4.

Residual feature importances. XGBoost residual-layer feature importances for the emissions model are dominated by month, which accounts for between 39% (Eugene, Bend, Portland) and 44% (Albany, Medford) of residual importance across all cities. The remaining importance is distributed across the temperature-spread variables (tmax, tmin, tdmean), precipitation, and wet bulb temperature, each contributing roughly 8–11%. Calendar variables weekend and day_of_week contribute small shares (5–8% combined). The dominance of month indicates that beyond demand and the temperature-related OLS variables, the residual model relies primarily on seasonal patterns to explain variation in total CO_2.

5.4 Warming Scenarios: Demand and Emissions

As the previous sections have discussed, the hybrid demand and emissions models were applied to the post-coal-exit dataset under uniform temperature shifts from 0°F to +15°F in 0.5°F increments, with post_july_2022 set to 1 across all rows. This next section discusses them in further detail and depicts the grid response over time in figures 5.4 and 5.4.1.

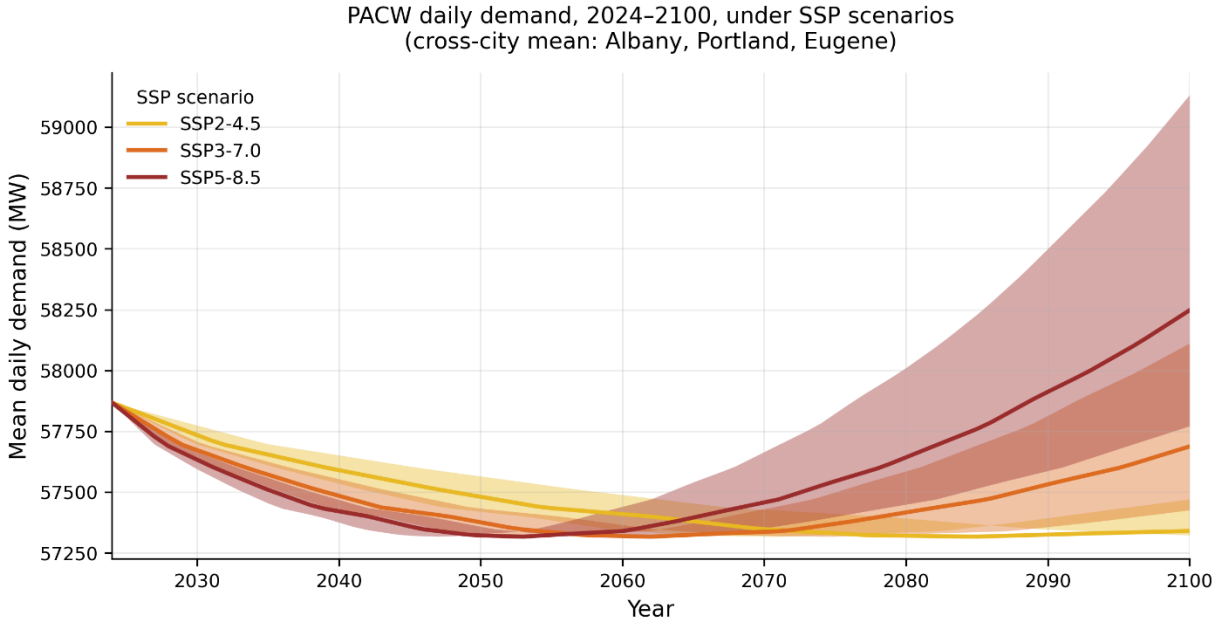


Figure 5.4: Demand Through 2100

The demand response across the three most predictive models (Albany, Portland, Eugene) follows a consistent qualitative pattern: a shallow decline through to 2060, followed by a more nuanced response in further years. Warming under SSP2-4.5 would see demand remain below baseline, whereas warming under SSP3-7.0 and SSP5-8.5 could see demand rise back above baseline by at most +4.6% (Eugene). The depth of the demand minimum is small in absolute terms (roughly 15–30 MW).

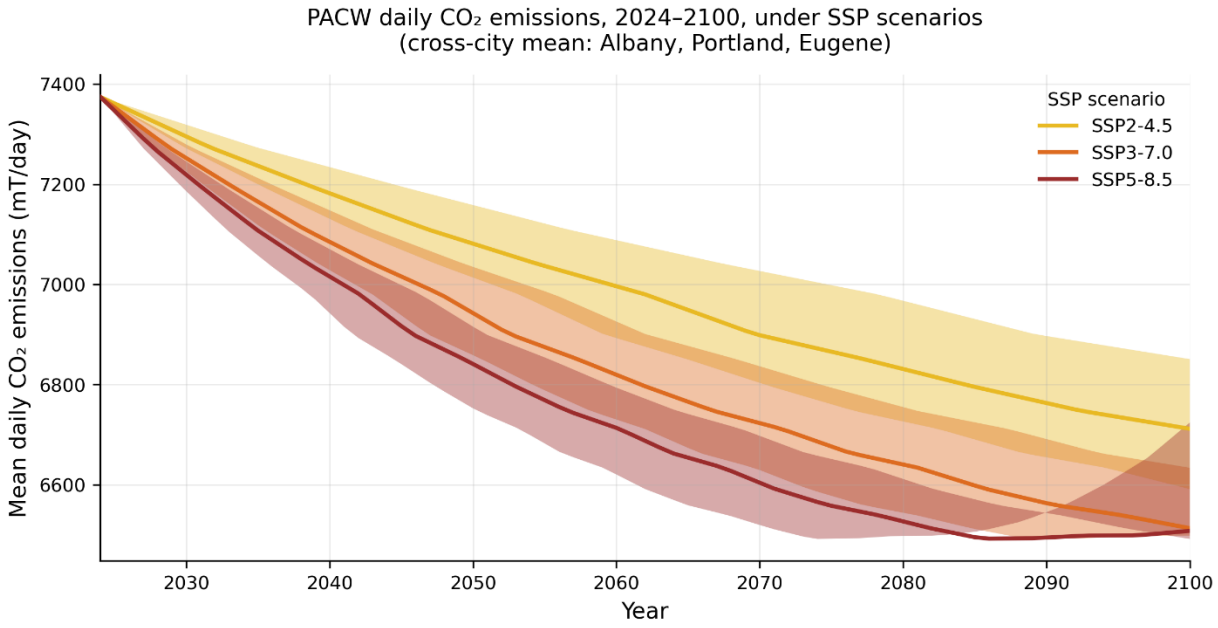


Figure 5.4.1: Emissions Through 2100

The emissions response shows wider cross-city variation as seen in Figure 5.3, but a more consistent trend over time. The three best predictive models (Albany, Portland, Eugene) show a clear emissions decline through to 2100 regardless of warming. Variation only arises in the amount of the decline; more significant warming implies more significant emissions loss. The inflection point visible in figure 5.3 is still present here in figure 5.4.1. However, it's only visible under extreme warming given SSP5-8.5 by 2090 or 2100.

5.5 The Coal Exit and Its Quantification

The single most stable empirical finding in the thesis briefly merits note, the OLS coefficient on the post_july_2022 structural-break indicator in the direct-CO_2 specification. Recall figure 2.3.

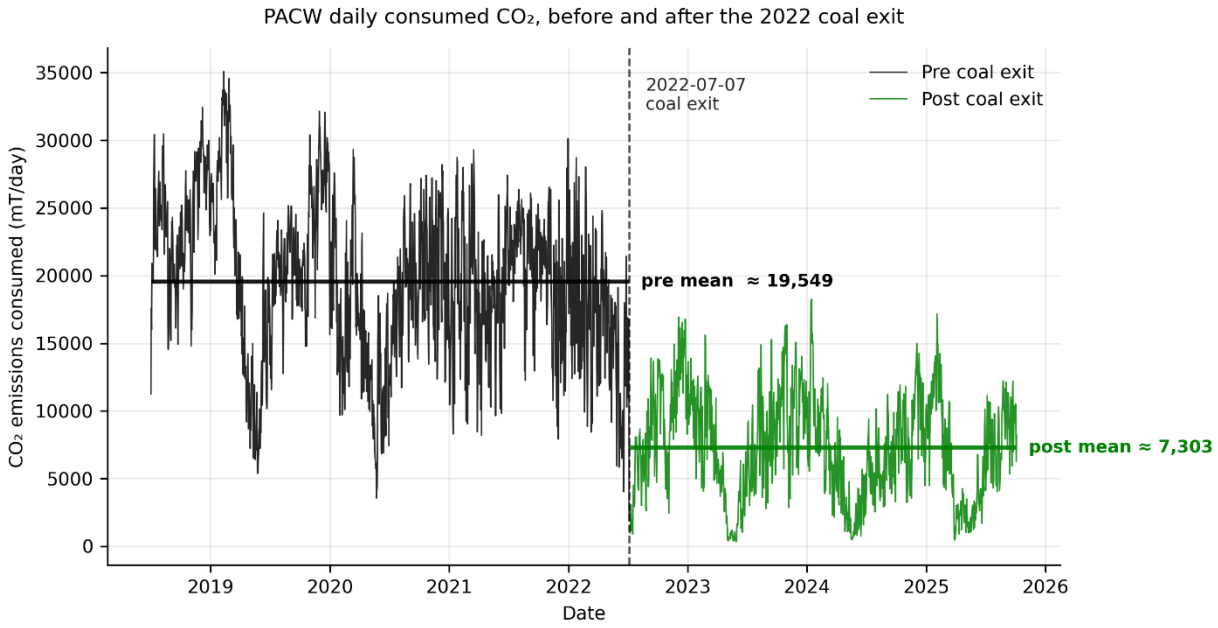


Figure 2.3: CO_2 Emissions Before and After the Coal Exit

Holding demand and weather constant, the 2022 coal retirement reduced average daily CO_2 emissions on the PACW grid by between 11,950 and 12,790 metric tons per day across all five specifications. The post_july_2022 coefficient is robust to the choice of climate-station signal, and it is the largest-magnitude coefficient in every emissions regression by an order of magnitude.

5.6 Summary of Findings

The empirical findings from Chapters 4's methodology, applied across five climate stations, can be summarized as follows.

The PACW demand response to temperature is **heating-dominated.** HDD coefficients in the demand model are large, positive, and highly significant across all five cities (+64 to +85 MW per heating-degree-day), while CDD coefficients are modest or negative in four of five cities. The hybrid demand model achieves test R^2 values between 0.79 and 0.84 across cities, with Albany producing the highest test R^2 (0.840).

The emissions model captures a single dominant structural finding: the post_july_2022 coefficient is between −11,950 and −12,790 metric tons CO_2 per day across all five city specifications, with t-statistics above 48 and p-values below 0.001 in every case. Holding demand and weather constant, the 2022 coal exit reduced average daily PACW CO_2 emissions by approximately 12,000 metric tons per day, and this estimate is robust to climate-station choice.

Under uniform warming scenarios, all five city models project a shallow demand minimum at +3.5°F to +5°F warming (0.7%–1.2% below baseline), followed by demand growth of +2.2% to +4.6% at +15°F. The emissions response shows wider cross-city variation, but a consistent trend over time: all five cities project an emissions decline until an inflection at high warming levels (+8°F to +11°F), with reductions of 5.6% (Medford) to 15.8% (Eugene) at the inflection.

The substantive interpretation of these findings, including the negative-CDD finding, the comparison between climate-driven and structural decarbonization, and the policy implications of the cross-city variation, is the subject of Chapter 6.

Chapter 6: Discussion

6.1 Interpreting the U-Curve: A Heating-Dominated System

The PACW demand model fits a U-shaped response to temperature: demand is minimized at moderate temperatures and rises in both colder and hotter conditions. However, the shape of that U-curve, and the asymmetry between its cold and hot sides, tells us something specific about how the PACW system is used.

A finding in section 5.2 deserves direct attention: the CDD coefficient is negative in four of the five city specifications. Read in isolation, this seems to imply that hot days reduce electricity demand, which is wrong. Hot days do drive demand up. The four temperature-derived OLS variables (mean temperature, temperature squared, HDD, CDD) work as a system, not in isolation. The pair of mean temperature and temperature squared together produces a smooth, symmetric quadratic curve that rises in both directions from a single minimum. That parabola does most of the work of fitting the U-shape. The HDD and CDD variables are corrections that allow the model to add or subtract steepness on one side of the curve without affecting the other. However, there is significant multicollinearity between HDD, CDD, and the temperature variables. In a system whose hot-side response is shallower than the symmetric parabola alone would predict, the CDD coefficient comes out negative as a correction. The full curve, generated by all four variables together, still rises on hot days, but is pulled down by the negative CDD value. This happens specifically because of the multicollinearity present.

Once the full fitted curve is evaluated rather than read off the coefficient signs, the per-degree asymmetry is more nuanced than the "heating-dominated" framing alone implies. Evaluating each city's fitted curve at reference temperatures of 20°F (cold) and 95°F (hot), with 58°F as the comfort minimum, the per-degree slopes are roughly:

City	Cold Side Slope (MWh/day per °F)	Hot Side Slope (MWh/day per °F)	Ratio (cold/hot)
Eugene	1039.2	1060.8	0.98
Albany	984.0	964.8	1.02
Bend	580.8	964.8	0.60
Portland	955.2	708.0	1.35
Medford	1017.6	789.6	1.29

Table 6.1: Hot and Cold Side Demand Response

Eugene and Albany show essentially symmetric per-degree response: a degree colder than 58°F drives demand up by roughly the same MW as a degree hotter. Portland and Medford show the cold-side dominant pattern: per-degree response is 29% to 35% larger on the cold side than on the hot side. Bend is the surprise, it shows hot-side dominance, with per-degree cooling response 66% larger than per-degree heating response. However, as previously established Bend is a weak predictor of the system as a whole and sits in a climate largely separate from the majority of the PACW region.

The cross-city variation in per-degree response complicates a simple "heating-dominated" framing for the fitted curves. But the framing remains correct for the system as a whole. Across the 10.3-year analytical window Albany logs 2,304 cold-side days against only 1,435 hot-side days, with cumulative annual HDD exposure averaging more than twice annual CDD exposure. Oregon spends substantially more time in heating-side conditions than in cooling-side conditions.

This distinction matters for interpreting the warming scenarios. As warming progresses, days that were previously cold-side days move toward 58°F or cross into cooling-side territory.

The shift reduces the cumulative annual HDD exposure by more than it raises cumulative annual CDD exposure, because there are simply more cold-side days available to be re-categorized. This is the mechanism behind the shallow demand minimum visible in Section 5.4: at +3.5°F to +5°F warming, the loss of cumulative heating exposure slightly exceeds the gain in cooling exposure, producing a small net demand decline before further warming pushes the system into a regime where cooling growth dominates.

The single Bend exception in the per-degree comparison deserves a brief note. Bend's cooling-side per-degree response is 66% larger than its heating-side response, opposite the pattern at the other four stations. This is consistent with Bend's position as a high-desert station east of the Cascade Range, where summer temperatures rise more sharply with each degree of dry-side warming and where the historical training window contains more variation in hot-extreme conditions than in cold-extreme conditions. The cumulative annual HDD/CDD ratio computed at Bend would also tell a different story than at Albany, because Bend is colder on average. The cross-city variation in per-degree response is itself a finding worth noting, and is discussed further in Section 6.4.

The framing carried into the rest of this discussion is therefore the more empirically honest one. PACW is heating-dominated in cumulative annual exposure and in the climate that the system sits in, with cold-side days outnumbering hot-side days by roughly 60% to 40% across the year. The per-degree response of the fitted demand curves varies meaningfully by climate station, with the Willamette Valley stations showing roughly symmetric response and the others diverging in different directions. The system-level claim that cold days drive PACW demand more than hot days do is consistent with both pieces, symmetric per-degree response

combined with substantially more cold-day exposure produces more total cold-side demand pressure on the system as a whole.

6.2 A Carbon-Lopsided Grid: The Emissions Asymmetry

The asymmetry that defines PACW demand has a counterpart on the emissions side, but the mechanism is different. Where the demand-side asymmetry comes from the climate spending more time below 58°F than above it, the emissions-side asymmetry comes from how the grid actually meets that demand, which generators are dispatched on cold versus hot days, and how clean or dirty those generators are.

To interpret what the emissions model is saying about how warming affects daily CO_2 output, it's critical to think about what each variable in the regression is really doing. The OLS layer of the direct-CO_2 model includes Demand directly as a regressor. That means the other coefficients, tmean, temp_squared, HDD, CDD, and post_july_2022, are describing how emissions vary holding demand constant. They are not telling us about absolute emissions on cold versus hot days. They are telling us about the dispatch mix: how dirty the megawatt-hours generated to meet demand actually are, and how that dirtiness shifts as temperature changes.

To make the asymmetry concrete, the fitted city emissions curve was evaluated at three reference temperatures, a HDD of 20°F, a CDD day of 20°F and at the comfort point 58°F. Holding demand at the modeling-dataset mean (2,404 MW) and post_july_2022 set to 1, figure 6.2 depicts this evaluation.

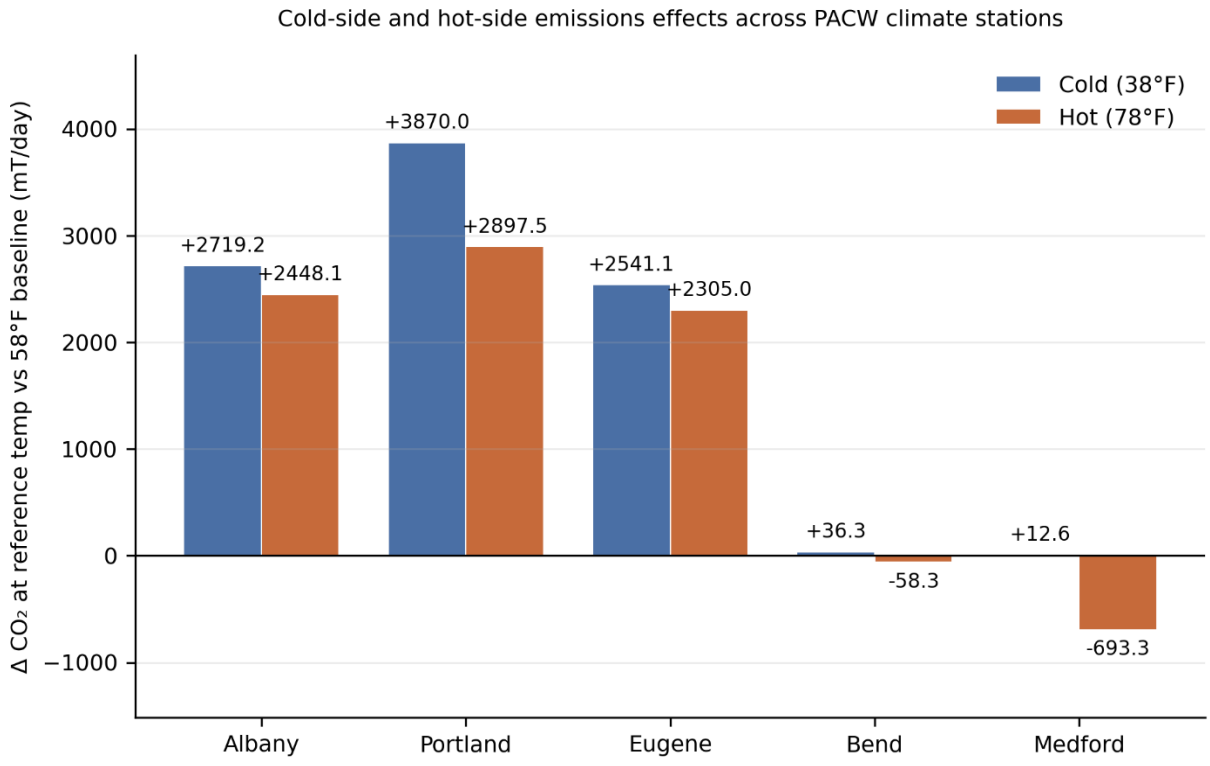


Figure 6.2: Cold and Hot Side Emissions Response

In Albany, both extremes raise predicted emissions above the comfort baseline, by +2,719 mT/day on the cold side and +2,448 mT/day on the hot side, but the cold-side rise is meaningfully steeper per degree. The Eugene specification produces a similar pattern (+2,541 mT/day cold, +2,305 mT/day hot), and the Portland specification produces an even larger gap (+3,870 cold, +2,898 hot). Across the three Willamette Valley and lower Columbia stations, the picture is consistent: both sides of the temperature curve raise per-MW emissions, but cold-side variation does so more aggressively than hot-side variation.

This dispatch-mix asymmetry has a name worth using directly: PACW is a **carbon-lopsided grid**. The grid is asymmetric, dirtier per MW on cold days than on hot days, and dirtier on either kind of extreme than at the comfort point. The mechanism is consistent with PACW's generation portfolio. Summer dispatch on the PACW grid is anchored by hydroelectric output

during runoff season, supplemented by wind generation and by solar reaching its annual peak. The marginal megawatt-hour during summer dispatch is largely drawn from this clean stack. Winter dispatch is different. Hydroelectric output is constrained by reservoir refilling timing and environmental flow requirements; solar reaches its annual minimum; wind is variable and uncorrelated with cold extremes. The marginal megawatt-hour during cold-weather dispatch is more likely to be drawn from imports purchased from neighboring balancing authorities, which on average draw from dirtier portfolios than PACW's own generation provides. The HDD coefficient is positive because cold days draw more imports per MW; the CDD coefficient acts as correction that softens the symmetric parabola into the asymmetric curve the data actually contains.

The two stations whose climate signal is most decoupled from the actual load center, Bend and Medford, show muted versions of the same pattern. At the same reference temperatures and demand, Bend's cold-side and hot-side effects are within a few dozen mT/day of zero, and Medford's are similarly small. These muted curves do not undermine the carbon-lopsided framing; they are properties of the climate-load coupling at those specific stations and are discussed further in Section 6.4. The substantive carbon-lopsided story is most clearly visible in the Willamette Valley and lower Columbia specifications where the climate signal best tracks PACW load.

The carbon-lopsided framing has a direct consequence for what warming does to total grid emissions. As warming progresses, days that were previously cold-side days move toward the comfort zone or cross into cooling-side conditions. Each degree of warming reduces the cumulative annual HDD exposure, which the OLS coefficient says is associated with dirtier per-MW dispatch, so reducing HDD exposure cleans up the per-MW intensity. Each degree of

warming also raises cumulative annual CDD exposure, which (because the symmetric parabola is being softened by a negative CDD term) is associated with a much shallower per-MW rise than the cold-side reduction is producing in the opposite direction. The net effect is that warming pulls the per-MW dispatch in a cleaner direction over much of the warming sweep, which combined with the modest demand response described in 6.1 produces emissions declines through to 2100. Section 6.3 takes this finding further by comparing the magnitude of these climate-driven emissions changes to the magnitude of the structural emissions reduction produced by the 2022 coal exit.

6.3 Structural vs. Climate-Driven Decarbonization

Two distinct decarbonization signals are visible in the PACW dataset and in the model outputs documented in Chapter 5. The first is structural: the 2022 coal retirement reduced average daily CO_2 emissions on the grid by roughly 12,000 mT/day, holding demand and weather constant. The second is climate-driven: warming through the OCCRI projection range produces additional emissions reductions through the dispatch-mix mechanism described in Section 6.2. This section compares the two and discusses their significance.

The structural signal is the cleaner of the two to quantify and has been discussed extensively. The retirement of PacifiCorp's coal generation removed approximately 12,000 metric tons of daily CO_2 from PACW's average emissions baseline, after accounting for demand and temperature effects. The retirement compressed the per-MW carbon intensity of the system by roughly two-thirds.

The climate-driven signal is more complicated and much smaller. Across the five cities, the decline is consistent but greatest at Eugene: a 15.8% decline at most, falling from a baseline of 7,353 mT/day to a minimum of 6,193 mT/day. The decline at Albany, the primary

specification, is shallower: from 7,553 mT/day to 6,922 mT/day. In absolute terms, the largest climate-driven emissions reduction at any plausible warming level is roughly 1,160 mT/day at Eugene, or roughly 631 mT/day at Albany.

Setting the two effects against each other directly, the coal retirement removed approximately 12,000 mT/day, and the climate-driven reduction removes at most 1,160 mT/day. There is no plausible combination of climate scenarios and modeling choices that makes the climate-driven decarbonization comparable in magnitude to what the coal retirement already produced.

This finding has direct implications for how to think about grid decarbonization on a system like PACW. The temperature-driven dispatch-mix story documented in Section 6.2 is real and physically meaningful, warming genuinely does shift PACW's per-MW intensity downward through the mechanism of cleaner summer dispatch and reduced reliance on cold-weather imports. But the magnitude of that shift is small relative to what was already accomplished by retiring a single category of generation asset. Climate change is doing the second-order work of making a cleaner grid slightly cleaner. Policy did the first-order work of making a dirty grid clean.

It is worth being precise about what this comparison does and does not say. It does not say that climate change is good for PACW emissions in some general sense, or that the climate-driven signal is unimportant. It does say that, on a grid in active structural transition, the structural decisions made by utilities and regulators are the single most important decarbonization lever with the largest demonstrated effect, and the climate signal is at most a secondary contributor. The implication for grid planning is that the policy-relevant question is not only "how should we adapt the grid to climate-driven emissions changes?" but "how do we

replicate the kind of structural decarbonization observed in PACW, *everywhere*?" The answer to that question depends regionally on which generators remain on the dirty end of dispatch curve and what it would take to retire or displace them.

A second implication worth surfacing is about the warming scenario projections themselves. The headline emissions findings that warming through 2050 to 2100 produces meaningful emissions reductions on the PACW grid, are real, but they should not be over-read as a policy argument that climate change "helps" decarbonization. Without the coal retirement, the same warming applied to the same demand pattern would produce a different and substantially dirtier result. The climate-driven decarbonization observed in this thesis is contingent on the structural change that preceded it.

The practical reading of this chapter is therefore that on a grid like PACW, the most consequential lever for emissions reduction is the one already pulled. Future structural decisions, which natural gas units stay or leave, which import contracts are renewed, which storage and renewable additions enter the dispatch order, will determine the next order of magnitude of decarbonization. Climate-driven changes will continue to nudge the system, but the direction and magnitude of those nudges depend on a generation portfolio that is itself shaped by policy.

6.4 Comparison to Existing Literature

The findings of this thesis sit alongside a body of empirical literature on temperature-driven electricity demand, summarized in Chapter 2. Three points of comparison are worth surfacing directly.

The U-shaped temperature-demand response documented across all five city specifications is consistent with the broader empirical literature. Aroonruengsawat and Auffhammer (2011), working with California household billing data, found a U-shape with

pronounced geographic heterogeneity in both the demand minimum and the response slopes. Yao (2021) generalized the pattern globally using satellite night-light data and identified a U-shape with a global minimum near 14.6°C (approximately 58°F), strikingly close to the empirically derived 58°F base used in this thesis. The cross-city variation in per-degree response observed in the PACW data, Eugene and Albany roughly symmetric, Bend hot-side dominant in the fitted curve, Portland and Medford cold-side dominant, is consistent with the heterogeneity Aroonruengsawat and Auffhammer documented across California's climate zones. The mechanism is similar in both cases: local climate and local building stock interact with a centralized load to produce different per-degree responses at different observation points.

The heating-dominated character of the PACW system aligns with Emonekwe and Emodi (2022), who used quantile regression on G7 economies and found that residential demand in heating-dominated economies responds more strongly to cold days than to hot days. The empirical pattern they identified at the country level is visible at the balancing-authority level in the PACW data, with the additional nuance that the dominance varies by climate-station signal. Where they find consistent country-level cold-day dominance, this thesis finds a more textured picture in which the cumulative-exposure framing (more cold days than hot days, driving more total annual heating demand) is robust across all stations even when the per-degree response varies. The two findings are complementary.

The carbon-lopsided framing developed in Section 6.2 is, to the best of my knowledge, novel as an empirical result for a U.S. balancing authority. The closest precedent in the literature is the broader observation that marginal emissions during peak periods exceed average annual emissions factors, particularly when peakers and imports are disproportionately drawn upon during high-demand events. This thesis extends the marginal-emissions framing in two specific

ways: by tying it to the cold-side rather than the conventionally-discussed hot-side of the temperature distribution, and by quantifying it at the OLS coefficient level on a single grid system. The result that PACW is dirtier per MW on cold days than on hot days inverts the standard "peakers fire on hot days" story that dominates the U.S. literature, because PACW's peakers and imports fire on cold days instead. This is a regionally specific finding that reflects PACW's hydroelectric anchor and its low historical AC penetration, and it would not generalize to grids whose summer peak demand exceeds their winter peak demand.

What is not directly comparable to the existing literature is the magnitude finding from Section 6.3, that the structural decarbonization from the 2022 coal retirement dwarfs any climate-driven emissions change at any plausible warming level. There is no comparable peer-reviewed study that puts a specific dollar-or-pound number on the relative magnitude of policy-driven versus climate-driven emissions changes on a single balancing authority. The closest analogues are integrated assessment models that compare projected national emissions trajectories under different policy scenarios, but these operate at a different level of aggregation and rarely produce the kind of cleanly-attributed coefficient that this thesis does. The 10x ratio between structural and climate signals is, in this sense, a regional contribution that does not have a direct comparison in the existing literature.

6.5 Limitations

This thesis makes substantive empirical claims, and several limitations on those claims should be made explicit.

Climate-station coupling and station selection. The cross-city comparison developed throughout this thesis demonstrates that station choice matters meaningfully for both demand and emissions modeling outcomes. The most defensible specifications are the ones whose climate

signal is closely coupled to the load center, Albany, Eugene, and to a lesser extent Portland, and the muted patterns observed at Bend and Medford reflect climate-load decoupling rather than model failure. This has a direct methodological implication for any future study attempting to replicate this framework on another balancing authority: geographic placement inside the BA territory is necessary but not sufficient. The selected climate signal must also be tightly coupled to where the load actually lives. For PACW, that meant the I-5 corridor in the Willamette Valley. For other balancing authorities, the corresponding choice will depend on the geography of the underlying load distribution. This is a lesson worth carrying forward, and a limitation worth flagging here: the production specification adopted in this thesis is defensible for PACW specifically, and not automatically transferable.

Uniform warming scenarios. As discussed in Section 4.7, the warming scenarios apply a uniform additive shift to all temperature variables on every day in the input dataset. This approach captures the structural response of the system to mean temperature change, but it abstracts away from seasonally non-uniform warming patterns, the disproportionate increase in extreme heat events under warming, changes in daily temperature range, and supply-side hydrological consequences of changed precipitation patterns. A more sophisticated treatment would couple this framework to a downscaled climate ensemble. The directional findings of this thesis are likely robust to such an extension; the specific magnitude estimates are less so.

Short post-coal training window. The direct CO_2 model is trained on roughly four years of pre-coal-exit data and only one to two years of post-coal-exit data within the chronological training window. This produces test R^2 values for the emissions model in the 0.33 to 0.52 range, which are modest. As more post-coal-exit data accumulates, the model could be retrained with a more balanced window and would likely produce better test performance.

No formal hyperparameter tuning. The Random Forest and XGBoost residual models use fixed default settings. Cross-validated tuning could plausibly improve test performance modestly, although the relatively tight train-test gaps observed for XGBoost suggest the gains would be incremental.

Demand and emissions models are not jointly trained. The two-stage scenario procedure feeds predicted demand from the demand model into the emissions model as an input. Errors in the demand model propagate into the emissions model. The architecture is more transparent than a single end-to-end model would be, and the propagation is bounded by the small Demand coefficient in the emissions OLS, but it is a real source of compound uncertainty in the scenario projections.

The dispatch-mix story is mechanistic, not directly verified. The carbon-lopsided framing in Section 6.2 explains the observed HDD and CDD coefficient signs through a hypothesized mechanism, cleaner summer dispatch, dirtier cold-weather imports. The mechanism is consistent with the structure of PACW's generation portfolio and with the empirical Import Share data discussed in Chapter 2, but the model itself does not directly verify which generators are dispatched on which days. A more granular analysis using EIA's hourly generation-by-fuel data could test the mechanism explicitly. This thesis takes the mechanism as the most plausible physical explanation for the coefficient pattern, but it remains a hypothesis rather than a directly tested claim.

6.6 Future Work

Several extensions of this framework would address the limitations above and produce additional findings worth pursuing.

Coupling to downscaled climate projections. The most natural extension is replacing the uniform warming sweep with a downscaled climate ensemble that produces seasonally and event-distributed temperature inputs. This would propagate ensemble uncertainty into the demand and emissions projections and produce probability-weighted ranges rather than single point estimates. The OCCRI reports and other research across the Pacific North West does provide projections at the daily resolution that this framework would require, and the existing pipeline could be adapted to ingest those projections directly.

Mechanism verification using hourly generation data. EIA's Hourly Electric Grid Monitor includes generation broken out by fuel source at hourly resolution. A direct test of the dispatch-mix mechanism described in Section 6.2 would compute marginal generation-by-fuel during different temperature regimes and confirm whether cold-day marginal generation is in fact dirtier than hot-day marginal generation per MW. If confirmed, this would convert the carbon-lopsided framing from a hypothesized mechanism to a directly verified one. If not, the OLS coefficient pattern would require an alternative explanation.

Extension to other balancing authorities. The hybrid OLS-plus-tree-ensemble architecture, the chronological validation procedure, and the cross-station comparison framework developed in this thesis are not specific to PACW. They could be applied to any U.S. balancing authority with EIA Grid Monitor coverage and PRISM climate data extraction. Cross-system comparison would test whether the findings of this thesis generalize: do other balancing authorities also exhibit asymmetric U-curves? Do they also show a structural-vs-climate-driven decarbonization gap? Are heating-dominated grids systematically different from cooling-dominated grids in their climate response? A multi-system extension is the natural next step in establishing whether PACW is representative or exceptional.

Joint demand-and-emissions model. The two-stage architecture used here was a deliberate choice for interpretability, but a single jointly-trained model with appropriate uncertainty propagation could reduce the compound-error issue noted in Section 6.5. A multi-task neural network or a Bayesian hierarchical model could produce joint posterior distributions over demand and emissions under warming, which would be more honest about the propagated uncertainty than the current two-stage framework.

Peak vs. average distinction. The mean-demand and mean-CO_2 projections are reported throughout this thesis, with peak and 95th-percentile values shown only briefly in the warming scenario tables. A more focused analysis of how warming affects peak load and peak emissions specifically, particularly the shift from heating-driven peaks to cooling-driven peaks under warming, would be directly relevant to grid capacity planning, where peaks rather than averages determine reserve margin and infrastructure requirements.

Updated post-coal regime characterization. As more post-coal data accumulates, the structural break coefficient will become more precisely estimated and the post-coal scenario projections will become more credible. Re-running this analysis with a 2027 or 2028 data cut would produce a substantially better-anchored characterization of the post-coal grid.

These directions are not exhaustive. The framework developed here is intended to be extensible, and the production specification (Albany hybrid XGBoost, time-aware split, 0–15°F warming sweep) is documented in sufficient detail that any of these extensions could be built on top of it without fundamental architectural changes.

Chapter 7: Conclusion and Reflection

7.1 Summary and Policy Takeaways

This thesis set out to characterize how rising temperatures will reshape daily electricity demand and total CO_2 emissions on the PacifiCorp West (PACW) Balancing Authority. Three findings answer that question.

The PACW grid is heating-dominated. The system spends the majority of its days below the demand minimum, and this is the most significant factor or characterizing how the grid reacts to warming. Across the warming sweep, demand follows a shallow U-shape, where initially the loss of heating exposure barely outweighs the gain in cooling exposure, but eventually that relationship flips, and gradual demand growth is observed through higher warming levels. The system the grid was built to serve is fundamentally a winter system, and it remains one across every warming scenario tested.

The PACW grid is also carbon-lopsided. Eugene, Albany, and Portland, the three stations most tightly coupled to PACW load, all show cold-side effects of emissions greater than hot side effects. The mechanism is the dispatch mix: PACW's summer stack (hydro, wind, solar) is materially cleaner than its winter stack (constrained hydro, gas, and imports from neighboring authorities). When warming reduces cold-side exposure and raises hot-side exposure, the per-MW intensity of dispatch shifts in a cleaning direction. Combined with the modest demand response, the result is that warming through 2050 and 2100 produces emissions decline.

But these climate-driven effects are not necessarily the headline. The 2022 retirement of PacifiCorp's coal generation reduced average daily PACW CO_2 emissions by an order of magnitude larger than the deepest climate-driven reduction projected at any plausible warming level. The structural decarbonization that policy and infrastructure decisions have already

produced dwarfs the climate-driven decarbonization that warming might add. Climate change is doing the second-order work of making a cleaner grid slightly cleaner. Policy did the first-order work of making a dirty grid clean.

The policy takeaway from this comparison is direct. On a grid in active structural transition, the generation mix decisions are the lever with the largest demonstrated effect on emissions. The relevant question for grid planners and policymakers is now not only "how should we adapt the grid to climate-driven emissions changes?" but "how do we replicate the kind of structural decarbonization observed in PACW, *everywhere*?" The answer depends on which generators remain on the dirty end of regional dispatch curves and what it would take to retire or displace them. The PACW result demonstrates that the magnitude available from this kind of structural change is large enough to dominate any projected climate-driven changes. Other balancing authorities, particularly those still operating coal capacity at scale, have access to the same lever. The findings of this thesis are therefore most useful not only as projections of the PACW system specifically, but as evidence of what one structural decision can accomplish when measured cleanly against the alternatives.

The future emissions trajectory of grids like PACW will not be primarily determined by climate change. It will be determined by which units retire, which contracts are renewed, and which storage and renewable capacity enters the dispatch order. Climate-driven changes will continue to nudge the system, but the direction and magnitude of those nudges depend on a generation portfolio that is itself shaped by policy. The policy work and the climate work are not separate projects; the structural decisions made now determine how much room climate change has to do its second-order work in the decades ahead.

Bibliography

Aroonruengsawat, A., & Auffhammer, M. (2011). Impacts of climate change on residential electricity consumption: Evidence from billing data. In G. D. Libecap & R. H. Steckel (Eds.), *The economics of climate change: Adaptations past and present* (pp. 311–342). University of Chicago Press. http://www.nber.org/chapters/c11991.pdf

Auffhammer, M., & Aroonruengsawat, A. (2012). Simulating the impacts of climate change, prices and population on California's residential electricity consumption. *Climatic Change, 109*(S1), 191–210. https://doi.org/10.1007/s10584-011-0299-y

Auffhammer, M., Baylis, P., & Hausman, C. H. (2017). Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the United States. *Proceedings of the National Academy of Sciences, 114*(8), 1886–1891. https://doi.org/10.1073/pnas.1613193114

Daly, C., Halbleib, M., Smith, J. I., Gibson, W. P., Doggett, M. K., Taylor, G. H., Curtis, J., & Pasteris, P. A. (2008). Physiographically-sensitive mapping of temperature and precipitation across the conterminous United States. *International Journal of Climatology, 28*(15), 2031–2064.

Deroubaix, A., et al. (2021). Large uncertainties in trends of energy demand for heating and cooling. *Nature Communications, 12*(1), Article 4413. https://doi.org/10.1038/s41467-021-25504-8

Emonekwe, C. C., & Emodi, N. V. (2022). Temperature and residential electricity demand for heating and cooling in G7 economies: A method of moments panel quantile regression approach. *Climate, 10*(10), Article 142. https://doi.org/10.3390/cli10100142

Fleishman, E. (Ed.). (2025). *Seventh Oregon climate assessment*. Oregon Climate Change Research Institute, Oregon State University. https://doi.org/10.5399/osu/1181

GridInfo. (2026). *PacifiCorp West | U.S. balancing authority*. https://www.gridinfo.com/balancing-authority/pacificorp-west

Kumar, R., et al. (2020). Asymmetrical response of California electricity demand to summer-time temperature variation. *Scientific Reports, 10*, Article 10904. https://doi.org/10.1038/s41598-020-67695-y

Li, Y., et al. (2018). Climate change and residential electricity consumption in the annual global context. PubMed Central. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6329981/

National Oceanic and Atmospheric Administration. (2026, March). *U.S. winter 2025–2026 climate report*. NOAA National Centers for Environmental Information.

Northwest Power and Conservation Council. (2016). *Seventh power plan*. https://www.nwcouncil.org/energy/powerplan/7/plan/

PacifiCorp. (n.d.). *PacifiCorp fact sheet*. https://www.pacificorp.com/content/dam/pcorp/documents/en/pacificorp/about/PacifiCorp_Fact_Sheet.pdf

PacifiCorp. (2017). *2017 integrated resource plan, Appendix F: Flexible reserve study*. https://www.pacificorp.com/content/dam/pcorp/documents/en/pacificorp/energy/integrated-resource-plan/2017-irp/2017-irp-support-and-studies/2017_IRP_AppendixF_Flexible_Reserve_Study_2017_FINAL.pdf

PacifiCorp. (2021). *2021 integrated resource plan, Volume I.* https://www.pacificorp.com/energy/integrated-resource-plan.html

PacifiCorp, TNA Merchant Power, Inc., & Chehalis Power Generating, LLC. (2008). *Order authorizing merger and disposition of jurisdictional facilities* (Docket No. EC08-82-000, 124 FERC ¶ 61,046). Federal Energy Regulatory Commission.

Pacific Power. (n.d.-a). *Areas we serve*. https://www.pacificpower.net/community/service-area.html

Pacific Power. (n.d.-b). *Pacific Power service area map* [PDF]. https://www.pacificpower.net/content/dam/pcorp/documents/en/pacificpower/PP_ServiceAreaMap-Simple.pdf

Pacific Power. (2026, April 9). *Pacific Power completes construction on Link River Dam West Side canal* [Press release]. https://flashalert.net/id/PacificPower

Philip, S. Y., et al. (2022). Rapid attribution analysis of the extraordinary heat wave on the Pacific coast of the US and Canada in June 2021. *Earth System Dynamics, 13*(4), 1689–1713. https://doi.org/10.5194/esd-13-1689-2022

Portland General Electric. (n.d.). *PGE service area*. https://portlandgeneral.com/about/info/service-area

PRISM Climate Group. (n.d.). *Descriptions of PRISM spatial datasets for the conterminous United States*. https://prism.oregonstate.edu/documents/PRISM_datasets.pdf

PRISM Climate Group. (2025). *PRISM gridded climate data* [Dataset AN91d]. Oregon State University. https://prism.oregonstate.edu/

Roberts, M. J., et al. (2021). *Using temperature sensitivity to estimate shiftable electricity demand: Implications for power system investments and climate change* (arXiv preprint No. 2109.00643). arXiv. https://arxiv.org/abs/2109.00643

Stull, R. (2011). Wet-bulb temperature from relative humidity and air temperature. *Journal of Applied Meteorology and Climatology, 50*(11), 2267–2269.

Urness, Z. (2026, March 9). Oregon ties record for warmest winter in 2025-26. Why was it so hot? *Statesman Journal*.

U.S. Energy Information Administration. (n.d.). *Hourly electric grid monitor*. https://www.eia.gov/electricity/gridmonitor/

Yao, J. (2021). *Electricity consumption and temperature: Evidence from satellite data* (IMF Working Paper No. 21/22). International Monetary Fund. https://www.imf.org/en/Publications/WP/Issues/2021/02/05/Electricity-Consumption-and-Temperature-Evidence-from-Satellite-Data-50031