

The Causal Effect of Carbon Pricing on Industrial Energy Consumption in Canada

by

Amirali Kamalirad

A Thesis Submitted to the Faculty of Graduate and Postdoctoral Studies of
The University of Manitoba
in partial fulfillment of the requirements of the degree of

MASTER OF ARTS

Department of Economics
University of Manitoba
Winnipeg

Copyright © 2026 by Amirali Kamalirad

Abstract

This thesis examines the effect of carbon pricing policies on industrial energy consumption across Canadian provinces. Carbon pricing has become a key element of Canada’s climate policy framework, but its impact on industrial energy use is less clear than its expected effect on emissions. The Canadian context is useful for studying this question because provinces introduced these policies at different times and through different designs. Using a province-year panel dataset, this study applies a staggered Difference-in-Differences framework based on the [Callaway and Sant’Anna \(2021\)](#) method, with controls for GDP, population, and climate conditions measured by heating and cooling degree days. The results do not show clear evidence that carbon pricing caused a statistically significant reduction in industrial energy consumption. Some estimates suggest changes after implementation, especially across treatment cohorts, but the pattern is not consistent or robust enough to conclude a causal effect. Overall, based on the findings, the relationship between carbon pricing and aggregate industrial energy use in Canada remains empirically uncertain, especially in a setting where provinces differ in terms of policy design and industrial structure.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Dr. Ryan Godwin, for his guidance, support, and valuable feedback throughout my thesis. I am especially grateful for his relaxed and encouraging manner, which helped me stay confident and motivated during the process. I would also like to thank my committee members, Dr. Oguzoglu and Dr. Saberian, for their helpful comments and suggestions. Their feedback improved the direction, clarity, and quality of this work.

I am also grateful to the Department of Economics at the University of Manitoba for providing the academic environment and support necessary to complete this thesis.

Finally, I would like to thank my mother and my wife for their constant belief in me, patience, and support throughout this journey.

To my late father, whose love and memory live on within me.

Contents

Abstract	i
Acknowledgements	ii
1 Introduction	1
2 Institutional Background	4
3 Literature Review	7
4 Data and Variable Construction	11
4.1 Data Sources and Sample	11
4.2 Energy Consumption, GDP, and Population	12
4.3 Carbon Pricing Treatment Variable	14
4.4 Weather Data and Climate Variables	14
4.5 Final Panel and Descriptive Statistics	16
5 Empirical Strategy	19
5.1 Overview of the Empirical Strategy	19
5.2 Treatment Timing and Target Parameter	20
5.3 Identification Assumptions	20
5.4 Covariates and Baseline Implementation	21
5.5 Aggregation of Treatment Effects	22

5.6	Assessment of Pre-Treatment Patterns	24
5.7	Inference	25
6	Results	26
6.1	Overview of the Baseline Estimates	26
6.2	Identification Diagnostics and Pre-Treatment Patterns	27
6.3	Aggregate Treatment Effects	27
6.4	Dynamic Treatment Effects	28
6.5	Treatment-Cohort Heterogeneity	29
6.6	Summary of Main Results	31
7	Robustness Checks	33
7.1	Purpose of the Robustness Analysis	33
7.2	One-Year Anticipation Specification	34
7.3	Estimator Sensitivity	36
7.4	Log-Weather Specification	37
7.5	Restricted Sample Check	38
7.6	Alternative ETWFE Specification	39
7.7	Summary and Comparison	41
8	Conclusion	42
8.1	Summary of the Study	42
8.2	Main Findings	43
8.3	Policy Implications	43
8.4	Limitations	44
8.5	Future Research and Concluding Remarks	45
A	Additional Robustness Results	50
A.1	One-Year Anticipation Cohort Results	50

A.2	Doubly Robust Specification	51
A.3	Log-Weather Specification	51
A.4	Supplementary ETWFE Results	51

List of Tables

2.1	Main Carbon Pricing Policies and Treatment Timing Across Provinces	5
4.1	Main Variables and Data Sources	12
4.2	Treatment Cohorts Used in the Empirical Analysis	15
4.3	Summary Statistics	18
6.1	Aggregate Treatment Effect Estimates	28
6.2	Treatment-Cohort Effects	30
7.1	Robustness Check: Aggregate Effects with One-Year Anticipation . .	34
7.2	Estimator Sensitivity: Baseline and Doubly Robust Estimates	36
7.3	Robustness Check: Aggregate Effects with Logged Weather Controls .	38
7.4	Restricted Sample Check: Baseline and 1995–2018 Estimates	39
7.5	Alternative ETWFE Specification: Average Treatment Effect	40
A.1	Cohort Effects with One-Year Anticipation	50

List of Figures

4.1	Log Industrial Energy Consumption by Province, 1995–2024	13
4.2	Selected Weather Stations After Quality Screening	17
6.1	Dynamic Effects of Carbon Pricing on Log Industrial Energy Consumption	29
6.2	Average Treatment Effects by Treatment Cohort	31
7.1	Dynamic Effects with One-Year Anticipation	35
7.2	Event-Study Estimates from the Alternative ETWFE Specification	40
A.1	Cohort Effects with One-Year Anticipation	51
A.2	Dynamic Effects from the Doubly Robust Specification	52
A.3	Cohort Effects from the Doubly Robust Specification	52
A.4	Dynamic Effects with Logged Weather Controls	53
A.5	Cohort Effects with Logged Weather Controls	53
A.6	Cohort Effects from the Alternative ETWFE Specification	54
A.7	Calendar-Year Effects from the Alternative ETWFE Specification	55

Chapter 1

Introduction

Carbon pricing has become one of Canada's main policies for reducing greenhouse gas emissions and leading toward a lower-carbon economy. The idea is that by making carbon-intensive activities more expensive, governments create a motivation for firms and households to use energy more efficiently, reduce fossil fuel consumption if possible, and adjust production and consumption decisions. However, in practice, the policy impact depends on how firms respond to the price signal and how they change their energy use.

Studying the industrial sector is both useful and challenging for this question. Industrial activity, based on the data, includes manufacturing, mining, oil and gas extraction, and other production processes that rely heavily on energy inputs. At the same time, many other factors can change the amount of energy consumption. For example, output levels, economic conditions, weather variation, energy intensity, and provincial industrial composition can influence energy demand. Thus, evaluating carbon pricing requires a framework that can separate policy effects from broader changes in provincial economies.

There is a good opportunity for this study in Canada. Some provinces adopted their own systems before the federal framework became more important, while oth-

ers became subject to the federal backstop or introduced systems that met federal requirements later. This staggered adoption creates useful variation for policy evaluation, although it makes the analysis more complicated because all the policies mentioned may not have the same impact.

This thesis specifically asks whether carbon pricing caused a reduction in industrial energy consumption after accounting for other factors. It uses a provincial panel dataset and applies a staggered Difference-in-Differences model based on the [Callaway and Sant’Anna \(2021\)](#) method. This approach fits the context and allows treatment effects to differ between treatment cohorts and over time, rather than forcing all provinces and all years to share one common treatment effect.

Industrial energy consumption is the main outcome variable, and carbon pricing adoption is the main treatment variable. The empirical models control for provincial GDP, population, and heating and cooling degree days. GDP and population capture changes in economic scale, while weather variables account for temperature-related variation in energy demand. The analysis also examines dynamic treatment effects to see whether the response changes as provinces remain exposed to carbon pricing for longer periods.

The results do not provide strong evidence that carbon pricing caused a statistically significant reduction in aggregate provincial industrial energy consumption. Some estimates show changes after intervention, but the general pattern is not consistent, so a causal effect cannot be strongly supported. This does not mean that carbon pricing does not affect industrial behaviour. It means that any effect on aggregate industrial energy consumption is difficult to identify, especially when policies differ in coverage and design and when provinces have very different industrial structures.

The thesis also examines the sensitivity of the findings using additional specifications. These include allowing for a one-year anticipation period before policy im-

plementation and estimating an alternative version of the [Callaway and Sant’Anna \(2021\)](#) model using a different estimator. These checks are used to assess whether the main conclusions depend on treatment timing or estimation choices.

This thesis has two main contributions. First, it studies industrial energy consumption, an outcome that has received less direct attention than emissions, household energy use, or transportation fuel demand. Second, it applies a modern staggered Difference-in-Differences approach to Canadian provincial carbon pricing, where policy adoption occurred at different times and through different systems. Overall, the study highlights why treatment timing, provincial heterogeneity, and economic conditions matter when evaluating climate policy.

The next chapters build on this introduction. Chapter 2 explains the institutional background of carbon pricing in Canada, Chapter 3 reviews the related literature, Chapter 4 describes the data and variable construction, and Chapter 5 presents the empirical strategy. Chapter 6 reports the main results, Chapter 7 examines robustness checks and sensitivity analysis, and Chapter 8 concludes with the main findings, policy implications, limitations, and directions for future research.

Chapter 2

Institutional Background

Carbon pricing in Canada was developed through a mix of provincial initiatives and federal minimum standards. Before the federal framework became the main policy, some provinces had already introduced their own systems. Alberta applied an early industrial pricing system for large emitters in 2007. British Columbia introduced a broad carbon tax in 2008. Quebec also developed its own framework before joining a cap-and-trade system. These early choices created important differences among provinces, both in when carbon pricing began and in how it was designed.

The federal role became more important after the Pan-Canadian Framework and the Greenhouse Gas Pollution Pricing Act. The federal framework had two main parts: a fuel charge on fossil fuels and an Output-Based Pricing System for large industrial emitters. Provinces and territories could still design their own systems, but those systems had to meet minimum national standards. Where a provincial system did not meet the federal standard, the federal backstop could apply.

This policy framework was important because not all system impacts are the same. For example, a carbon tax directly increases the cost of covered fossil fuels, while a cap-and-trade system sets an emissions limit and allows firms to trade allowances. These systems can encourage firms to reduce emissions intensity while

decreasing competitiveness concerns. Thus, the effect of carbon pricing on total industrial energy consumption may be less direct than its effect on emissions.

The timing of provincial carbon pricing policies is a key element of this thesis. British Columbia’s carbon tax began in 2008. Alberta’s early system applied to large industrial emitters. Quebec’s cap-and-trade system began in 2013 and was linked with California’s system in 2014. Ontario introduced a cap-and-trade program in 2017 but later cancelled it, making it a special case for empirical analysis. Several other provinces became subject to the federal backstop or introduced systems aligned with federal requirements in 2019. A summary of the main policy timing used in this thesis is provided in Table 2.1.

Table 2.1: Main Carbon Pricing Policies and Treatment Timing Across Provinces

Province	Main system	Key timing	Notes
Alberta	Industrial large-emitter pricing	2007	Early system focused on large emitters and emissions intensity
British Columbia	Carbon tax	2008	Broad fuel-based carbon tax applied to fossil fuels
Saskatchewan	Federal backstop / provincial industrial system	2019	Partial federal application with provincial industrial components
Manitoba	Federal backstop	2019	Later-treated province under the federal framework
Ontario	Cap-and-trade / federal backstop	2017, 2019	Cap-and-trade was implemented and later repealed
Quebec	Cap-and-trade	2013	Provincial system linked with California
New Brunswick	Federal backstop / provincial system	2019	Later-treated case with later provincial adjustments
Nova Scotia	Cap-and-trade	2019	Provincial cap-and-trade system, later replaced by an OBPS
Newfoundland and Labrador	Provincial performance standards system	2019	Provincial carbon pricing system for large industry
Prince Edward Island	Federal OBPS / provincial fuel pricing	2019	Federal industrial pricing with provincial fuel pricing elements

This setting creates both an opportunity and a limitation for causal analysis. The opportunity comes from staggered adoption. Provinces did not adopt carbon pricing simultaneously, so the analysis can compare changes in industrial energy consumption across treated and comparison provinces over time. The limitation is that treatment is not identical across provinces. A carbon tax, a cap-and-trade system, and an output-based pricing system may not create the same incentives for firms. In the main empirical model, these design differences are not estimated as separate policy-type effects. Instead, they shape how the estimated average effect of carbon pricing adoption should be interpreted.

Industrial energy consumption, as the only available and reliable data, is also a difficult outcome to evaluate. It depends not only on carbon pricing, but also on production levels, GDP, population, weather conditions, industrial composition, and fuel availability. For example, an oil-producing province may follow a different energy-use pattern than a province with a larger manufacturing base or a cleaner electricity mix. These differences make economic and other controls necessary in the model and also highlight the importance of careful interpretation.

It is also important to note later policy changes. In 2025, the federal fuel charge was set to zero, while the federal Output-Based Pricing System remained in effect. Since the main sample ends before this change, the 2025 reform should be treated as a later development rather than part of the main treatment period.

Overall, the institutional background helps explain the empirical design. It shows why a staggered Difference-in-Differences framework is appropriate, why treatment timing must be defined carefully, and why the estimated effects should be considered as average effects of carbon pricing adoption across a heterogeneous provincial context, rather than as the effect of one uniform national policy.

Chapter 3

Literature Review

The empirical literature on carbon pricing has expanded as more provinces have implemented their carbon pricing regulations. This literature mainly focuses on whether carbon pricing has changed the behaviour of firms and households in a measurable way. In theory, these policies should reduce fossil fuel use and emissions by increasing the cost of carbon-intensive activity. In practice, the estimated effects depend on a wide range of factors, including policy design, sectoral coverage, tax rates, revenue recycling, and the level of adjustment.

Most Canadian studies have focused on British Columbia's carbon tax, introduced in 2008. [Murray and Rivers \(2015\)](#) review British Columbia's carbon tax as one of the first broad carbon pricing policies in North America and summarize evidence suggesting that the policy reduced emissions while having limited effects on aggregate economic performance. Their discussion is useful because it shows both the promise and the difficulty of evaluating an actual carbon tax. The policy was comprehensive and relatively transparent, but its effects varied across outcomes and sectors.

Several studies have examined how British Columbia's carbon tax has affected fuel consumption. [Rivers and Schaufele \(2015\)](#) find that the policy reduced gasoline

demand more strongly than an equivalent change in the market price of gasoline, suggesting that the carbon tax may have been more important to consumers than ordinary price movements. [Lawley and Thivierge \(2018\)](#) also find reductions in household gasoline consumption, but show that the response differed across locations, with stronger effects in urban areas and weaker or no clear response in rural and northern areas. These findings suggest that carbon pricing can affect energy use, but the response depends on adjustment opportunities.

Other studies analyze residential energy use. [Xiang and Lawley \(2019\)](#) estimate the effect of British Columbia’s carbon tax on residential natural gas consumption using panel regression and synthetic control methods. They find evidence that the carbon tax reduced residential natural gas consumption, with estimates around 7 percent in their synthetic control approach and larger reductions in panel regressions. This work is relevant to this thesis because it focuses directly on energy consumption rather than only emissions, although the setting is residential rather than industrial.

The evidence on aggregate emissions is more mixed. [Pretis \(2022\)](#) studies British Columbia’s carbon tax using Difference-in-Differences, synthetic control, and a break-detection approach. The paper finds reductions in transportation emissions but does not find large statistically significant reductions in aggregate emissions. This distinction is important for this thesis because industrial energy consumption is also an aggregate outcome. Even if carbon pricing affects some sectors or behaviours, the effect may be harder to detect when the outcome combines many activities with different trends and adjustment margins.

Some other studies examine firm-level or sector-level responses, which are relevant for understanding industrial energy use. [Ahmadi et al. \(2022\)](#) use plant-level manufacturing data from British Columbia and find that the carbon tax reduced emissions by about 4 percent while also reducing emission intensity and increasing output. Their results suggest that some firms may respond to carbon pricing

by improving efficiency rather than simply reducing production. [Yamazaki \(2017\)](#) also shows that the employment effects of the British Columbia carbon tax were heterogeneous. Carbon-intensive and trade-sensitive industries experienced employment declines, while cleaner service industries saw employment gains. These studies highlight the importance of industrial composition.

More recent work has moved beyond the British Columbia case. [Maladkar \(2024\)](#) compares carbon taxes, cap-and-trade systems, and baseline-and-credit systems across Canadian provinces using a Difference-in-Differences framework. The study finds limited evidence that these policies significantly reduced overall emissions, but suggests that effects differ by policy type and sector. For example, cap-and-trade systems appear to be more effective in high-emitting sectors. This research is particularly relevant because it emphasizes policy heterogeneity across provinces, although its main focus is emissions and economic outcomes rather than industrial energy consumption.

The methodological literature is also important for this thesis. Traditional Difference-in-Differences methods are most straightforward when there are two groups and two periods. Canadian carbon pricing involves multiple provinces adopting different policies at different times. [Callaway and Sant’Anna \(2021\)](#) emphasize that staggered Difference-in-Differences designs require careful attention to the parallel trends assumption, comparison groups, and the causal parameter being estimated. This matters for the present thesis because a standard two-way fixed effects model may not be sufficient when treatment effects vary across provinces or over time.

[Xu \(2024\)](#) applies a staggered Difference-in-Differences approach to Canadian provincial carbon pricing in the context of food inflation. Although the outcome variable differs from this thesis, the paper is useful because it shows how staggered carbon pricing across provinces can be used to estimate dynamic effects in a Canadian setting. This supports the use of a modern staggered Difference-in-Differences

framework for studying other outcomes, including industrial energy consumption.

Overall, based on existing literature, carbon pricing can affect emissions, fuel demand, firm behaviour, and economic outcomes, but the results are not uniform. Many studies focus on British Columbia, household fuel consumption, transportation, or emissions. Fewer studies examine aggregate industrial energy consumption across all provinces. This thesis addresses this gap using variation in carbon pricing adoption to estimate whether these policies caused measurable changes in industrial energy use. In doing so, it connects the energy-consumption literature with recent developments in staggered Difference-in-Differences methods and contributes to the empirical evaluation of Canadian climate policy.

Chapter 4

Data and Variable Construction

4.1 Data Sources and Sample

This thesis uses a province-year panel dataset for the ten Canadian provinces over the period 1995 to 2024. The territories are excluded because their data reliability, population size, industrial structure, energy markets, and weather-station coverage are completely different from those of the provinces. Limiting the sample to the ten provinces keeps the units of analysis more consistent.

The main outcome variable is industrial energy consumption, measured annually at the provincial level in terajoules. The energy data are taken from Statistics Canada’s annual energy balance table. Population and GDP are also obtained from Statistics Canada and are used as economic controls. Population captures differences in provincial scale, while GDP controls for changes in the level of economic activity across provinces and years ([Statistics Canada, 2026a,b,c](#)).

Weather variables are constructed separately from station-level daily temperature data obtained from Environment and Climate Change Canada’s Historical Climate Data. The data are extracted using the `weathercan` package in R ([Environment and Climate Change Canada, 2026](#); [LaZerte and Albers, 2026](#)), which provides tools for

accessing Canadian historical climate station data. These station-level observations are used to construct annual heating degree days and cooling degree days, which are then merged with the main province-year panel.

Table 4.1 summarizes the main variables used in the analysis.

Table 4.1: Main Variables and Data Sources

Variable	Construction	Source
Industrial energy consumption	Annual provincial industrial energy consumption, measured in terajoules and used in logarithmic form in the empirical analysis.	Statistics Canada, Table 25-10-0029-01
GDP	Annual provincial gross domestic product, used in logarithmic form to control for economic activity.	Statistics Canada, Table 36-10-0222-01
Population	Annual provincial population estimate, used in logarithmic form to control for provincial scale.	Statistics Canada, Table 17-10-0005-01
Carbon pricing treatment	Province-year treatment indicator based on the first year in which a province is coded as subject to carbon pricing.	Provincial and federal policy timelines
HDD18	Annual heating degree days constructed from daily mean temperature using an 18°C base.	ECCC Historical Climate Data via weathercan
CDD18	Annual cooling degree days constructed from daily mean temperature using an 18°C base.	ECCC Historical Climate Data via weathercan

Notes: ECCC refers to Environment and Climate Change Canada. The treatment variable is described in more detail in the treatment-timing section. HDD18 and CDD18 are constructed from daily station-level temperature data and aggregated to the province-year level.

4.2 Energy Consumption, GDP, and Population

Since industrial energy consumption, GDP, and population differ considerably in scale across provinces, the analysis uses their logarithmic forms. Using logarithms reduces the influence of large cross-provincial differences and allows estimated effects to be interpreted approximately as percentage changes. GDP and population are

included as controls for provincial economic activity, industrial energy demand, and broader differences in provincial size.

Figure 4.1 shows the evolution of log industrial energy consumption by province from 1995 to 2024. The figure illustrates that consumption follows different paths. Alberta and Saskatchewan show strong upward trends over much of the sample period, while several Atlantic provinces and Ontario display more visible declines or fluctuations. These differences support the use of additional controls in the empirical analysis.




Figure 4.1: Log Industrial Energy Consumption by Province, 1995–2024
Notes: The figure plots the logarithm of annual industrial energy consumption for each province. The vertical scale is allowed to vary across provinces to highlight within-province changes over time.

4.3 Carbon Pricing Treatment Variable

The policy background and provincial timing of carbon pricing are discussed in chapter 2, with the main timing summarized in Table 2.1. The treatment variable is defined at the province-year level. Let G_p denote the first year in which province p is coded as treated. The treatment indicator is defined as:

$$D_{pt} = \begin{cases} 1 & \text{if } t \geq G_p, \\ 0 & \text{if } t < G_p. \end{cases}$$

This definition is based on the staggered structure. Provinces enter treatment in different years, and after the first treatment year they remain coded as treated for the purpose of the empirical model. The treatment variable captures the timing of policy adoption rather than differences in the detailed design of provincial systems.

Table 4.2 reports the treatment cohorts used in the empirical analysis. The earlier-treated cohorts are 2007, 2008, 2013, and 2017, while six provinces enter treatment in 2019. Since all provinces are coded as treated from 2019 onward, the 2019 cohort has limited post-treatment identification under the not-yet-treated comparison design.

4.4 Weather Data and Climate Variables

Weather variables are included to control for temperature-related variation in industrial energy consumption, consistent with industrial energy-intensity models that include climate controls for seasonal weather variation in heating and cooling requirements (Bernstein et al., 2003). For this purpose, heating degree days and cooling degree days are constructed as annual province-level climate controls. These variables follow the standard degree-day approach, in which daily temperature deviations from a base temperature are used to measure heating and cooling requirements (U.S.

Table 4.2: Treatment Cohorts Used in the Empirical Analysis

Treatment cohort G_p	Province(s)	Number of provinces	Empirical role
2007	Alberta	1	Early-treated cohort
2008	British Columbia	1	Early-treated cohort
2013	Quebec	1	Mid-period treated cohort
2017	Ontario	1	Late pre-2019 treated cohort
2019	Saskatchewan, Manitoba, New Brunswick, Nova Scotia, Newfoundland and Labrador, Prince Edward Island	6	Later-treated cohort

Notes: G_p denotes the first year in which province p is coded as treated. Policy details are discussed in Chapter 2.

[Energy Information Administration](#), nd; [Price et al., 2020](#)).

To construct these variables, daily station-level temperature data are obtained from Environment and Climate Change Canada’s Historical Climate Data using the `weathercan` package in R ([Environment and Climate Change Canada, 2026](#); [LaZerte and Albers, 2026](#)). The construction starts from a station inventory. Since station coverage differs across provinces, a 20 percent random sample of stations is drawn within each province, with a minimum of two stations per province.

For each sampled station, daily mean temperature is extracted for the desired period, and missing data rates are also calculated. Stations with less than 5 percent missing daily mean temperature observations and sufficient coverage are preferred. For provinces without a preferred station, the most suitable fallback station is selected, as shown in Figure 4.2. The final climate panel uses 38 stations across the ten provinces, including 35 preferred stations and 3 fallback stations.

For selected stations, a complete daily date sequence is created, and remaining missing temperatures are filled using nearby observations from the previous and following two days when available. Then, heating degree days and cooling degree

days are constructed using a base temperature of 18°C:

$$HDD18_d = \max(18 - T_d, 0), \quad (4.1)$$

$$CDD18_d = \max(T_d - 18, 0), \quad (4.2)$$

where T_d is daily mean temperature. These daily measures are summed to the station-year level and then averaged across selected stations within each province-year. This gives annual province-level measures of $HDD18$ and $CDD18$, which are used as climate controls summarized with the other variables in Table 4.3.

After constructing the province-year climate panel, a limited outlier adjustment is applied. The adjustment identifies large deviations from each province’s historical pattern relative to the yearly trend in other provinces, and replaces flagged values using that trend-based relationship. This mainly affects only a small number of observations in New Brunswick and Prince Edward Island and is likely a consequence of missing data.

4.5 Final Panel and Descriptive Statistics

This analysis uses a balanced province-year panel for the ten provinces over the 1995–2024 period. The final panel places the outcome, treatment variable, economic controls, and climate controls at the same province-year frequency.

An important feature of the final panel is that not all variables enter the empirical specification in the same form. In particular, the model includes a semi-log structure that should be considered when interpreting the estimated effects and covariate adjustments. The functional form of these variables and their role in the baseline specification are discussed in Chapter 5, while summary statistics for the




Figure 4.2: Selected Weather Stations After Quality Screening
Notes: Preferred stations have less than 5 percent missing daily mean temperature observations. Fallback stations are used where no preferred station is available.

Table 4.3: Summary Statistics

Variable	N	Mean	SD	Min	Max
CDD18	300	142.139	85.191	5.800	448.942
GDP, billion dollars	300	185.057	222.255	4.313	956.994
HDD18	300	4631.243	709.565	3182.050	6469.634
Industrial energy (petajoules)	300	249.602	300.576	1.564	1270.452

final variables are reported in [Table 4.3](#).

Chapter 5

Empirical Strategy

5.1 Overview of the Empirical Strategy

This thesis uses the staggered Difference-in-Differences approach of [Callaway and Sant’Anna \(2021\)](#) to estimate the effect of carbon pricing on industrial energy consumption, accounting for variation in treatment timing and potential differences in effects across provinces, policy designs, industrial structures, and years since implementation.

Therefore, the analysis does not rely on a conventional two-way fixed effects coefficient as the main causal parameter. Instead, fixed provincial differences and common year shocks are addressed through the Difference-in-Differences comparison structure, while treatment effects are first estimated by treatment cohort and calendar year and then aggregated into summary measures. This approach is better suited to the Canadian carbon pricing setting, where treatment timing and policy design vary across provinces.

The outcome variable is the logarithm of industrial energy consumption in province p and year t , denoted by $\ln(Energy_{pt})$. The treatment variable is defined by the first year in which each province is coded as being subject to intervention. The main ques-

tion is whether industrial energy consumption changes after implementation relative to the counterfactual path that would have occurred in the absence of the policies.

5.2 Treatment Timing and Target Parameter

Let G_p denote the first year in which province p is treated. Provinces with the same first treatment year belong to the same treatment cohort. The main causal parameter is the group-time average treatment effect:

$$ATT(g, t) = E[Y_t(g) - Y_t(0) \mid G = g], \quad (5.1)$$

where g denotes the treatment cohort, t denotes the calendar year, $Y_t(g)$ is the potential outcome in year t for a province first treated in year g , and $Y_t(0)$ is the untreated potential outcome. In this thesis, Y_t refers to log industrial energy consumption.

The parameter $ATT(g, t)$ measures the average effect of carbon pricing in year t for provinces first treated in year g . This parameter is useful because it does not impose a common treatment effect across all provinces or all years. Instead, it allows the estimated effect to vary across treatment cohorts and over time.

5.3 Identification Assumptions

The first identifying assumption is staggered adoption. Once a province is coded as treated, it remains treated in subsequent years for the purpose of the empirical treatment definition. This assumption follows the coding of carbon pricing as a continuing policy exposure after the first treatment year.

The second assumption is conditional parallel trends. In the baseline specification, treated provinces are compared with provinces that are not yet treated in a given year. The identifying assumption is that, in the absence of carbon pricing,

treated provinces and not-yet-treated provinces would have followed parallel trends in industrial energy consumption after conditioning on observed covariates.

A simplified representation of this assumption is:

$$E[Y_t(0) - Y_{t-1}(0) \mid X_{pt}, G_p = g] = E[Y_t(0) - Y_{t-1}(0) \mid X_{pt}, G_p > t], \quad (5.2)$$

where X_{pt} is the vector of observed covariates and $G_p > t$ denotes provinces that are not yet treated in year t . In other words, after controlling for observed factors, not-yet-treated provinces are assumed to provide the counterfactual trend for provinces already exposed to carbon pricing.

The third assumption is no anticipation in the baseline model. The main specification sets the anticipation period equal to zero. This means that provinces are assumed not to change industrial energy consumption before the official treatment year solely because of future carbon pricing adoption. Since policy announcements may occur before implementation, a one-year anticipation specification is examined later as a robustness check.

The fourth assumption is overlap. Treated provinces must have comparable not-yet-treated provinces available in the relevant years. This condition is important because all provinces are treated from 2019 onward. Therefore, the not-yet-treated comparison group is available before 2019 but is exhausted once all provinces become treated.

5.4 Covariates and Baseline Implementation

The baseline model is estimated using the `att_gt()` function from the `did` package in R. The outcome variable is log industrial energy consumption, and the covariate vector includes log GDP, log population, *HDD18*, and *CDD18*:

$$Y_{pt} = \ln(Energy_{pt}), \quad (5.3)$$

$$X_{pt} = [\ln(GDP_{pt}), \ln(Population_{pt}), HDD18_{pt}, CDD18_{pt}]. \quad (5.4)$$

This specification separates economic scale controls from climate-exposure ones. GDP and population are entered in logarithmic form because they are positive scale variables and differ substantially across provinces. In contrast, *HDD18* and *CDD18* are kept in levels because degree-day variables measure accumulated heating and cooling exposure rather than economic size. This treatment is consistent with applied energy specifications that include degree-day variables as level climate controls. For example, [Bernstein et al. \(2003\)](#) include climate in an industrial energy-intensity regression as the sum of heating-degree and cooling-degree days, while the remaining variables are entered in logarithmic form. Similarly, [Paul et al. \(2009\)](#) estimate log electricity demand by sector, including the industrial class, and include heating and cooling degree days directly in the demand equation.

This specification uses not-yet-treated provinces as the comparison group, a universal base period, no anticipation period, and the outcome-regression estimator. The estimated objects are group-time average treatment effects, which are then aggregated in the next section. As a robustness check, the analysis also estimates an alternative specification in which the climate controls are entered in logarithmic form to assess whether the results are sensitive to this functional-form choice.

5.5 Aggregation of Treatment Effects

After estimating $ATT(g, t)$, the group-time effects are aggregated into summary measures. The first aggregation is the simple overall average treatment effect. Following

Callaway and Sant’Anna (2021), this can be written in two steps. First, the average post-treatment effect for each treatment cohort is:

$$\theta^S(g) = \frac{1}{T - g + 1} \sum_{t=2}^T \mathbf{1}\{t \geq g\} ATT(g, t). \quad (5.5)$$

The overall simple aggregate effect is then obtained by averaging these cohort-specific effects across treated cohorts:

$$\theta_O^S = \sum_{g \in \mathcal{G}} \theta^S(g) P(G = g), \quad (5.6)$$

where $P(G = g)$ is the share of treated provinces that first receive treatment in cohort g . In the empirical application, this share is determined by the number of provinces in each treatment cohort. This aggregation provides a single summary measure of the average effect of carbon pricing across treated provinces and identified post-treatment periods.

The second aggregation examines dynamic treatment effects by event time. Event time is defined as:

$$e = t - g, \quad (5.7)$$

where $e = 0$ is the first treatment year, positive values represent years after treatment, and negative values represent pre-treatment years. The event-study aggregation can be written generally as:

$$\theta_D(e) = \sum_{g \in \mathcal{G}_e} w(g, e) ATT(g, g + e), \quad (5.8)$$

where $\mathcal{G}_e$ is the set of treatment cohorts for which the group-time effect $ATT(g, g + e)$ is observed and identified at event time e . The weights are based on the relative size of each treatment cohort among the cohorts contributing to that event-time estimate:

$$w(g, e) = \frac{N_g}{\sum_{g' \in \mathcal{G}_e} N_{g'}}, \quad (5.9)$$

where N_g is the number of provinces in treatment cohort g . Thus, each dynamic estimate is a weighted average of the identified group-time treatment effects at the same length of exposure. This aggregation is used to examine whether the estimated effect changes as provinces remain exposed to carbon pricing for longer periods.

The third aggregation summarizes effects by treatment cohort. This is useful because carbon pricing policies differ across provinces in timing, design, and industrial coverage. Cohort-specific estimates therefore help show whether the effect is similar across treated provinces or whether it differs by treatment timing.

5.6 Assessment of Pre-Treatment Patterns

The identifying assumptions cannot be proven directly. However, the empirical analysis examines pre-treatment estimates as indirect evidence on the plausibility of the conditional parallel trends assumption. In the event-study results, negative event times correspond to years before carbon pricing adoption.

If conditional parallel trends is plausible, pre-treatment estimates should generally be close to zero and should not display a systematic pattern. These estimates are therefore used as a diagnostic tool rather than as a definitive test. If pre-treatment estimates are large or systematic, the estimated treatment effects must be interpreted more cautiously.

This diagnostic is especially important in this thesis because the number of provinces is small and treatment cohorts are unevenly sized. The pre-treatment evidence is therefore considered together with the magnitude, direction, and stability of the estimated treatment effects.

5.7 Inference

Inference is based on bootstrap standard errors and confidence bands produced by the `did` package. The baseline model uses 999 bootstrap iterations. Standard errors are clustered at the province level because observations within a province may be correlated over time.

The small number of provinces is an important limitation for inference. Since the analysis includes ten provinces and some treatment cohorts contain only one province, statistical precision should be interpreted cautiously. For this reason, the results chapter emphasizes not only statistical significance but also the direction, magnitude, and consistency of the estimates across event time and treatment cohorts.

Therefore, the baseline specification combines a staggered treatment design, not-yet-treated comparison provinces, no anticipation, economic controls, climate controls, and bootstrap-based inference.

Chapter 6

Results

6.1 Overview of the Baseline Estimates

This chapter discusses the results from the baseline staggered Difference-in-Differences model described in Chapter 5. The model estimates the effect of carbon pricing on log industrial energy consumption for the ten Canadian provinces over 1995–2024, using log GDP, log population, *HDD18*, and *CDD18* as controls. Treatment cohorts are defined by the first year of carbon pricing adoption, with Alberta, British Columbia, Quebec, and Ontario treated before 2019 and the remaining provinces treated in 2019. Since all provinces are treated from 2019 onward, the not-yet-treated comparison group is only available before that point, which limits identification for later-treated provinces. The results are reported as aggregate, dynamic, and cohort-specific treatment effects.

6.2 Identification Diagnostics and Pre-Treatment Patterns

Based on the pre-treatment diagnostic described in Section 5.6, Figure 6.1 provides mixed evidence on the identifying assumption. The pre-treatment estimates are not fully centered around zero, and some of the early estimates are imprecise or visibly different from zero. At the same time, the figure does not show a clear systematic pattern before treatment. Overall, this suggests that the baseline design does not deliver an ideal pre-treatment pattern, but the evidence also does not point to a strong pre-existing trend.

These results should not be interpreted as an automatic rejection of the design. Pre-treatment estimates provide diagnostic evidence on the plausibility of the conditional parallel trends assumption, but they do not directly test the unobserved counterfactual trend (Callaway and Sant’Anna, 2021; Marcus and Sant’Anna, 2021; Roth, 2022). Given the small number of provinces and the limited variation across treatment cohorts, the results are interpreted cautiously, with attention to their direction, precision, and heterogeneity rather than as evidence of a single uniform causal effect.

6.3 Aggregate Treatment Effects

Table 6.1 reports the aggregate treatment effect estimates from the baseline specification. The simple aggregate ATT is 0.2575 log points, interpreted in levels, this is approximately 29.4 percent, with a standard error of 0.2414. The 95 percent confidence interval ranges from -0.2156 to 0.7305. Since the interval includes zero, this estimate is not statistically significant.

The dynamic aggregate ATT is 0.4208 log points and is statistically significant

at the 5 percent level. This result should be interpreted with care. As discussed above, the event-study pattern is not uniform, and the pre-treatment estimates do not provide a fully clean diagnostic pattern. The cohort aggregate ATT is -0.0576 and is not statistically significant.

Taken together, the aggregate estimates do not provide evidence of a clear reduction in industrial energy consumption after carbon pricing adoption. The simple aggregate estimate is positive but statistically insignificant, while the cohort aggregate estimate is close to zero and also statistically insignificant. Although the dynamic estimate is positive and statistically significant, it should not be interpreted in isolation from the event-study pattern, pre-treatment diagnostics, and cohort composition.

Table 6.1: Aggregate Treatment Effect Estimates

Aggregation type	Estimate	Std. error	95% CI lower	95% CI upper
Simple aggregate ATT	0.2575	0.2414	-0.2156	0.7305
Dynamic aggregate ATT	0.4208	0.1900	0.0484	0.7932
Cohort aggregate ATT	-0.0576	0.3361	-0.7163	0.6011

Notes: The outcome variable is log industrial energy consumption. The baseline model uses not-yet-treated provinces as the comparison group, assumes no anticipation, and includes log GDP, log population, HDD18, and CDD18 as controls. Confidence intervals are reported at the 95 percent level.

6.4 Dynamic Treatment Effects

Figure 6.1 illustrates the event-study estimates by length of exposure to carbon pricing. The estimates around the treatment year and in the first few post-treatment years are generally close to zero and imprecise. For longer exposure periods, some of the estimates become positive and statistically significant.

This pattern cannot support a clear or stable decline in industrial energy consumption after intervention. What it shows instead is that the estimated response changes across event time. The later positive estimates also need to be considered with caution, since longer exposure periods are mainly informed by earlier-treated

provinces, especially Alberta and British Columbia. As a result, the composition of provinces contributing to the estimates is not the same across all event-time periods.

Overall, the dynamic results do not provide strong evidence of a consistent reduction in industrial energy consumption. They are more aligned with a response that varies by treatment cohort and length of exposure, which is expected given the institutional and policy differences discussed earlier.




Figure 6.1: Dynamic Effects of Carbon Pricing on Log Industrial Energy Consumption

Notes: The figure reports event-study estimates from the baseline staggered Difference-in-Differences specification. Negative event times refer to pre-treatment years, while non-negative event times refer to the treatment year and post-treatment years. The comparison group consists of not-yet-treated provinces.

6.5 Treatment-Cohort Heterogeneity

Based on Table 6.2, the estimates across treatment cohorts indicate substantial variation in the effect of carbon pricing. The differences are not only in the size of the estimates, but also in their direction. The 2007 and 2008 cohorts have positive esti-

mated effects, while the 2013 and 2017 cohorts have negative estimated effects. For all four reported cohorts, the corresponding confidence intervals do not include zero, which suggests that these estimates are statistically different from zero.

The estimated effect for the 2007 cohort is 0.6101 log points, and for the 2008 cohort it is 0.4021 log points. In contrast, the estimated effects for the 2013 and 2017 cohorts are -0.3193 and -0.9234 log points, respectively. This result also indicates that the effect of the policy on each cohort is different. Moreover, it suggests that the timing of treatment may matter, possibly because provinces adopted different types of regulatory policies and also had different industrial structures. For Alberta, the positive estimate may also reflect the province’s intensity-based large-emitter system and expanding industrial production.

The 2019 cohort is not reported in this table since all provinces are treated from 2019 onward. With the not-yet-treated comparison group, there are no untreated or not-yet-treated provinces left after 2019 to be used as a comparison group for this cohort. Therefore, the baseline design cannot separately estimate post-treatment effects for the 2019 cohort under this setting.

Table 6.2: Treatment-Cohort Effects

Treatment cohort	Estimate	Std. error	95% CI lower	95% CI upper
2007	0.6101	0.1959	0.2261	0.9941
2008	0.4021	0.1311	0.1452	0.6590
2013	-0.3193	0.0401	-0.3979	-0.2406
2017	-0.9234	0.1668	-1.2503	-0.5965

Notes: Cohort effects summarize average post-treatment effects for provinces first treated in each treatment year. The 2019 cohort is not reported because the not-yet-treated comparison group is exhausted once all provinces become treated.

Figure 6.2 presents the cohort-specific estimates in graphical form. The figure helps to show the differences across cohorts more clearly. The estimates for the 2007 and 2008 cohorts are located to the right of zero, while the estimates for the 2013 and 2017 cohorts are located to the left of zero. This again shows that the estimated

effects are not similar across treatment cohorts.




Figure 6.2: Average Treatment Effects by Treatment Cohort
Notes: The figure reports cohort-specific average treatment effects from the baseline staggered Difference-in-Differences model. The horizontal axis reports the estimated effect on log industrial energy consumption.

6.6 Summary of Main Results

Overall, the baseline results do not provide strong evidence of a clear reduction in provincial industrial energy consumption after carbon pricing adoption. The aggregate estimates are either statistically insignificant or difficult to interpret in isolation, especially given the imperfect pre-treatment pattern and the changing cohort composition across event time.

The results also show meaningful variation across treatment cohorts. The estimated effects differ in both magnitude and direction, which suggests that the policy effect was not uniform across provinces. This variation may reflect differences in treatment timing, policy design, industrial structure, and other province-specific conditions.

Therefore, the baseline evidence should be interpreted cautiously. Rather than showing a consistent reduction in industrial energy use, the results point to a mixed and uneven pattern. The next chapter examines whether this conclusion is robust to alternative specifications and additional checks.

Chapter 7

Robustness Checks

7.1 Purpose of the Robustness Analysis

This chapter examines whether the baseline conclusion is sensitive to alternative modelling choices. The baseline model in Chapter 6 uses a staggered Difference-in-Differences specification with not-yet-treated provinces as the comparison group, no anticipation period, the outcome-regression estimator, and weather controls entered in levels.

The robustness analysis considers five checks. First, it allows for a one-year anticipation period to account for possible adjustment before the official start of carbon pricing. Second, it replaces the outcome-regression estimator with the doubly robust estimator while keeping the same sample, controls, treatment timing, and comparison group. Third, it uses logged weather controls to check whether the results depend on the functional form of the climate variables. Fourth, it restricts the sample to 1995–2018 to check whether the baseline results are affected by observations after all provinces become treated. Finally, it estimates an alternative ETWFE specification as an additional check on the baseline pattern.

These checks are not used to search for a more convenient result. The aim is

simply to see whether the main interpretation remains reasonable when key modelling choices are changed. Since the baseline results do not show a clear and consistent reduction in industrial energy consumption, the robustness analysis asks whether that conclusion still holds under these alternative specifications.

7.2 One-Year Anticipation Specification

The first robustness check allows for a one-year anticipation period. The baseline model assumes that provinces do not adjust industrial energy consumption before the official treatment year. This may be restrictive if firms start preparing for carbon pricing before formal implementation, or if policy announcements affect behaviour before the policy becomes active.

To check this possibility, the anticipation period is set equal to one while the rest of the baseline structure is kept unchanged. Based on the results reported in the table 7.1, the estimates do not provide evidence of a reduction in industrial energy consumption. Instead, the aggregate effects are positive, and some are statistically significant. This suggests that the main conclusion is not driven by the strict no-anticipation assumption.

Table 7.1: Robustness Check: Aggregate Effects with One-Year Anticipation

Aggregation type	Estimate	Std. error	95% CI lower	95% CI upper
Simple aggregate ATT	0.5887	0.1679	0.2596	0.9178
Dynamic aggregate ATT	0.7032	0.1836	0.3434	1.0631
Cohort aggregate ATT	0.3375	0.2537	-0.1596	0.8347

Notes: The outcome variable is log industrial energy consumption. This specification allows for a one-year anticipation period. The model uses not-yet-treated provinces as the comparison group and controls for log GDP, log population, HDD18, and CDD18. Confidence intervals are reported at the 95 percent level.

Figure 7.1 shows the dynamic estimates under the one-year anticipation specification. The pattern is still uneven. The early post-treatment estimates are mostly close to zero or imprecise, while some later estimates are positive. This supports the

same cautious interpretation as the baseline results.




Figure 7.1: Dynamic Effects with One-Year Anticipation

Notes: The figure reports event-study estimates from the robustness specification with one anticipation period. Negative event times refer to pre-treatment years, while non-negative event times refer to the treatment year and post-treatment years under the anticipation-adjusted specification.

The cohort-specific results in Appendix Table A.1 and Figure A.1 also show that some cohort-level estimates are sensitive to the treatment-timing assumption, but they do not change the overall interpretation of this robustness check.

Overall, the one-year anticipation specification does not overturn the baseline conclusion. It does not show a clear reduction in industrial energy consumption after carbon pricing adoption. Instead, it confirms that the estimates are somewhat sensitive to assumptions about treatment timing, which somewhat supports real-world evidence that firms usually adjust in advance to expected policy changes.

7.3 Estimator Sensitivity

Another concern is whether the baseline findings depend on the specific estimator used in the [Callaway and Sant’Anna \(2021\)](#) framework. The baseline results are estimated using the outcome-regression approach. As a check, the model is re-estimated using the doubly robust estimator. This estimator combines outcome-regression adjustment with inverse-probability weighting, so it is less dependent on one specific modelling component than the outcome-regression estimator alone. The sample, treatment timing, comparison group, anticipation period, and covariates are kept unchanged.

This comparison is useful because it changes the estimation method without changing the main empirical design. If the results changed sharply, it would suggest that the findings are sensitive to the estimator rather than the underlying treatment pattern. As shown in [Table 7.2](#), this is not the case. The doubly robust estimates are very close to the baseline estimates.

Table 7.2: Estimator Sensitivity: Baseline and Doubly Robust Estimates

Estimator	Aggregation type	Estimate	Std. error	95% CI lower	95% CI upper
OR	Simple aggregate ATT	0.2575	0.2414	-0.2156	0.7305
OR	Dynamic aggregate ATT	0.4208	0.1900	0.0484	0.7932
OR	Cohort aggregate ATT	-0.0576	0.3361	-0.7163	0.6011
DR	Simple aggregate ATT	0.2608	0.2481	-0.2256	0.7471
DR	Dynamic aggregate ATT	0.4249	0.1627	0.1060	0.7438
DR	Cohort aggregate ATT	-0.0610	0.3520	-0.7509	0.6289

Notes: OR refers to the outcome-regression estimator, and DR refers to the doubly robust estimator. Both specifications use not-yet-treated provinces as the comparison group, assume no anticipation, and control for log GDP, log population, HDD18, and CDD18. Confidence intervals are reported at the 95 percent level.

The main message from this check is straightforward. The simple aggregate effect remains statistically insignificant, the dynamic aggregate effect remains positive, and the cohort aggregate effect remains close to zero. Therefore, switching from the outcome-regression estimator to the doubly robust estimator does not materially

change the results.

The cohort-level pattern is also preserved. The 2007 and 2008 cohorts remain positive, while the 2013 and 2017 cohorts remain negative. This supports the interpretation that the results reflect heterogeneity across treatment cohorts rather than a single common effect.

The dynamic and cohort figures closely follow the baseline patterns reported in Appendix Figures A.2 and A.3. Overall, this robustness check supports the stability of the baseline findings with respect to the estimation method.

7.4 Log-Weather Specification

The baseline specification keeps *HDD18* and *CDD18* in levels. However, it is useful to check whether the results depend on this functional-form decision. For this reason, an additional robustness check replaces *HDD18* and *CDD18* with $\ln(HDD18)$ and $\ln(CDD18)$, while keeping the rest of the baseline model unchanged.

Before estimating this specification, the logged weather variables were checked to ensure that the transformation did not create missing or non-finite values. This check is important because logarithms are only defined for positive values. In the final panel, both weather variables are positive, so the log transformation can be implemented without losing observations.

Table 7.3 reports the aggregate estimates from the log-weather specification. The results are somewhat weaker than the baseline estimates. The simple aggregate effect remains statistically insignificant, the dynamic aggregate effect is no longer statistically significant, and the cohort aggregate effect remains close to zero. Overall, the logged-weather specification does not suggest a clear reduction in industrial energy consumption following carbon pricing adoption.

The main point from this check is that changing the weather controls to logarith-

Table 7.3: Robustness Check: Aggregate Effects with Logged Weather Controls

Aggregation type	Estimate	Std. error	95% CI lower	95% CI upper
Simple aggregate ATT	0.1677	0.1987	-0.2218	0.5571
Dynamic aggregate ATT	0.1940	0.1571	-0.1138	0.5018
Cohort aggregate ATT	-0.0245	0.2253	-0.4661	0.4170

Notes: The outcome variable is log industrial energy consumption. This specification replaces *HDD18* and *CDD18* with $\ln(HDD18)$ and $\ln(CDD18)$. The model uses not-yet-treated provinces as the comparison group, assumes no anticipation, and controls for log GDP and log population. Confidence intervals are reported at the 95 percent level.

mic form does not change the overall interpretation. The estimates remain mixed and do not provide evidence of a robust negative effect. This supports the view that the baseline conclusion is not mainly driven by the decision to include *HDD18* and *CDD18* in levels.

The corresponding dynamic and cohort figures are reported in Appendix Figures [A.4](#) and [A.5](#). These figures provide additional detail, but the main message is already captured by the aggregate estimates in Table [7.3](#).

7.5 Restricted Sample Check

As an additional check, the baseline [Callaway and Sant’Anna \(2021\)](#) specification is re-estimated after restricting the sample to 1995–2018. In this restricted sample, the 2019 cohorts are still coded with their actual treatment timing, but they are not treated within the estimation window and therefore serve as not-yet-treated comparison provinces. This check is useful because all provinces become treated from 2019 onward in the full sample.

Table [7.4](#) shows that the main aggregate estimates are unchanged after restricting the sample. The standard errors change only slightly. Therefore, the baseline results are not driven by the presence of post-2018 observations in the dataset.

Table 7.4: Restricted Sample Check: Baseline and 1995–2018 Estimates

Aggregation type	Baseline estimate	Baseline SE	Restricted estimate	Restricted SE
Simple aggregate ATT	0.2575	0.2414	0.2575	0.2335
Dynamic aggregate ATT	0.4208	0.1900	0.4208	0.1825
Cohort aggregate ATT	-0.0576	0.3361	-0.0576	0.3429

Notes: The restricted specification uses the 1995–2018 sample. The 2019 cohorts retain their actual treatment year but are not treated within this restricted window. Both specifications use not-yet-treated provinces as the comparison group and include the same controls.

7.6 Alternative ETWFE Specification

The final robustness check uses an extended two-way fixed effects specification, following [Wooldridge \(2025\)](#), as an additional comparison model. This specification is not used as the main estimator, since the baseline analysis relies on the [Callaway and Sant’Anna \(2021\)](#) framework. Instead, it is used to check whether the main interpretation changes under a different modelling approach. The model is estimated using the `etwfe` package in R ([McDermott et al., 2025](#)).

The ETWFE model keeps the same outcome variable and controls as the baseline specification. In compact form, the model can be written as:

$$\ln(Energy_{pt}) = X'_{pt}\beta + \mu_p + \lambda_t + \sum_{g \in \mathcal{G}} \sum_{s=g}^T \delta_{gs} \mathbf{1}\{G_p = g\} \mathbf{1}\{t = s\} + \varepsilon_{pt}. \quad (7.1)$$

where X_{pt} includes log GDP, log population, $HDD18$, and $CDD18$. Province fixed effects, μ_p , control for time-invariant provincial differences, while year fixed effects, λ_t , control for common shocks across provinces. The interaction terms allow the estimated treatment effects to vary by treatment cohort and year. Standard errors are clustered at the province level.

Table [7.5](#) reports the average treatment effect from the ETWFE specification. The estimate is close to zero and statistically insignificant. Therefore, this alternative model does not provide clear evidence that carbon pricing reduced industrial energy

consumption.

Specification	Estimate	Std. error	95% CI lower	95% CI upper
ETWFE average effect	0.0338	0.3242	-0.6016	0.6693

Notes: The outcome variable is log industrial energy consumption. The ETWFE specification includes log GDP, log population, HDD18, and CDD18 as controls. Province and year fixed effects are included, and standard errors are clustered at the province level. Confidence intervals are reported at the 95 percent level.

Figure 7.2 reports the event-study estimates from the ETWFE specification. The estimates do not show a stable post-treatment decline. Some event-time estimates are positive, while others are negative, and most confidence intervals include zero. This pattern suggests that the alternative ETWFE model does not identify a clear dynamic treatment effect.




Figure 7.2: Event-Study Estimates from the Alternative ETWFE Specification

Notes: The figure reports event-study estimates from the ETWFE specification. The model includes economic and climate controls, province and year fixed effects, and province-clustered standard errors.

Overall, the ETWFE results support the cautious interpretation of the base-

line analysis. The average estimate is small and statistically insignificant, and the event-study pattern does not show a consistent reduction in industrial energy consumption after carbon pricing adoption. The cohort and calendar-year ETWFE figures are reported in Appendix Figures [A.6](#) and [A.7](#). These supplementary figures show additional heterogeneity, but they do not change the main conclusion.

7.7 Summary and Comparison

Overall, the robustness checks do not change the main conclusion of the baseline analysis. None of the alternative specifications provides clear evidence that carbon pricing reduced aggregate provincial industrial energy consumption.

The one-year anticipation model shows that the estimates are sensitive to assumptions about treatment timing. This is not surprising, since firms may respond before a policy formally begins. The doubly robust specification, however, gives results very close to the baseline model, suggesting that the main findings are not mainly driven by the choice of estimator.

The log-weather specification also leads to the same general interpretation. Changing *HDD18* and *CDD18* from levels to logarithms does not reveal a clear negative effect. The ETWFE model gives an average estimate close to zero, and it is statistically insignificant.

Taken together, these checks support a cautious reading of the results. The estimates remain mixed and heterogeneous across specifications. Therefore, the evidence does not support a strong conclusion that carbon pricing consistently reduced industrial energy consumption across Canadian provinces.

Chapter 8

Conclusion

8.1 Summary of the Study

This thesis examined whether carbon pricing policies affected industrial energy consumption across Canadian provinces. The focus on the industrial sector is important because this sector uses large amounts of energy, but its response to carbon pricing may be less direct than in household or transportation settings. Firms may adjust through efficiency, fuel choice, production decisions, or technology rather than through an immediate reduction in total energy use.

The analysis used a balanced province-year panel for the ten Canadian provinces over the period 1995 to 2024. The outcome variable was log industrial energy consumption. The empirical model controlled for log GDP, log population, heating degree days, and cooling degree days. The main method was a staggered Difference-in-Differences approach based on the [Callaway and Sant'Anna \(2021\)](#) framework, which is suitable because carbon pricing was adopted at different times across provinces.

8.2 Main Findings

The main results do not provide clear evidence that carbon pricing reduced aggregate provincial industrial energy consumption. The simple aggregate treatment effect is positive but statistically insignificant, while the cohort aggregate effect is close to zero. The dynamic estimates are also not easy to interpret as a clear policy effect because the event-study pattern is uneven and the pre-treatment evidence is not fully clean.

The cohort results provide the clearest evidence of heterogeneity. The 2007 and 2008 cohorts show positive and statistically significant estimated effects, while the 2013 and 2017 cohorts show negative and statistically significant estimated effects. This suggests that carbon pricing was associated with measurable changes in industrial energy consumption for some treatment cohorts, but the direction and size of the response differed across provinces.

The robustness checks do not overturn this conclusion. Allowing for one year of anticipation does not produce evidence of a reduction in industrial energy consumption. The doubly robust estimator gives results close to the baseline estimates. The log-weather specification and the alternative ETWFE model also do not provide strong evidence of a stable negative effect. Overall, the findings point to heterogeneous and timing-sensitive estimates rather than one clear average response.

8.3 Policy Implications

The results should not be read as evidence that carbon pricing has no environmental value. Industrial energy consumption is a broad outcome, and carbon pricing may affect other margins that are not captured here. For example, firms may reduce emissions intensity, switch fuels, invest in cleaner equipment, or change production processes without reducing total energy consumption in the short run.

The findings suggest that policy evaluation should go beyond aggregate industrial energy use. More detailed outcomes, such as emissions, emissions intensity, fuel composition, and sector-level industrial activity, may better capture the channels through which carbon pricing works. The results also show why provincial context matters. A single average estimate can hide large differences across policy designs and industrial structures.

8.4 Limitations

This thesis has several limitations. The first is the small number of units. The analysis uses ten provinces, and several treatment cohorts contain only one province. This limits statistical precision and makes inference more sensitive to province-specific patterns.

A second limitation is the level of aggregation. The outcome variable measures total industrial energy consumption at the provincial level. More detailed industrial subsector data would allow a stronger analysis of whether carbon pricing affected energy-intensive industries differently from other industrial activities. However, in practice, more disaggregated industrial data often contain missing values or are suppressed for confidentiality reasons. Because of this, the analysis could not be reliably broken down into detailed industrial subsectors.

A third limitation is the treatment definition. The treatment variable captures carbon pricing adoption, but it does not fully measure policy intensity, coverage, exemptions, or differences between carbon taxes, cap-and-trade systems, and output-based pricing systems. These differences matter, but they are difficult to represent with a single treatment indicator.

The sample also includes the COVID-19 period, when industrial activity and energy use followed unusual patterns in many provinces. These years fall in the post-

treatment period for earlier-treated cohorts and may affect some longer-exposure estimates.

Another limitation is the absence of consistent disaggregated energy price data for the full period and desired level of detail. If province-sector energy prices by fuel type were available consistently, they could improve the model by helping distinguish the effect of carbon pricing from broader changes in energy markets.

Finally, the Difference-in-Differences assumptions require caution. The pre-treatment patterns are not perfect, and the not-yet-treated comparison group becomes unavailable after all provinces are treated. As a result, the main identified effects are driven mostly by the 2007, 2008, 2013, and 2017 cohorts, and the results say less about the six provinces that entered treatment in 2019. These issues do not remove the value of the analysis, but they limit how strongly the estimates can be interpreted as causal.

8.5 Future Research and Concluding Remarks

Future research could extend this work in several directions. A more detailed industrial dataset would make it possible to study subsectors separately and test whether energy-intensive industries respond differently. Better energy price data could also strengthen the empirical design by accounting for changes in fuel and electricity costs. Another useful extension would be to examine emissions and emissions intensity directly, since carbon pricing is designed primarily to reduce emissions rather than energy consumption itself.

Overall, this thesis contributes to the evaluation of Canadian carbon pricing by studying industrial energy consumption using a modern staggered Difference-in-Differences approach. The evidence does not support a strong conclusion that carbon pricing consistently reduced aggregate provincial industrial energy consump-

tion during the sample period, although the cohort-level results indicate measurable and heterogeneous responses. The more defensible conclusion is cautious: the estimated effects are mixed, heterogeneous, and sensitive to some modelling choices. This suggests that the relationship between carbon pricing and industrial energy use in Canada is more complex than a single average treatment effect can fully capture.

Bibliography

- Ahmadi, Y., Yamazaki, A., and Kabore, P. (2022). How do carbon taxes affect emissions? plant-level evidence from manufacturing. *Environmental and Resource Economics*, 82:285–325.
- Bernstein, M. A., Fonkych, K., Loeb, S., and Loughran, D. S. (2003). State-level changes in energy intensity and their national implications. Technical Report MR-1616-DOE, RAND Corporation, Santa Monica, CA.
- Callaway, B. and Sant’Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Environment and Climate Change Canada (2026). Historical climate data. Online database. Available at climate.weather.gc.ca. Accessed May 16, 2026.
- Lawley, C. and Thivierge, V. (2018). Refining the evidence: British columbia’s carbon tax and household gasoline consumption. *The Energy Journal*, 39(2):147–172.
- LaZerte, S. and Albers, S. (2026). *weathercan: Easy Downloading of Weather Data from Environment and Climate Change Canada*. R package documentation. Available at docs.ropensci.org/weathercan.
- Maladkar, P. (2024). Comparing provincial carbon emission regulation policies in

- canada: Impacts on greenhouse gas emissions, production, and labor measures. Honors Thesis, University of North Carolina at Chapel Hill.
- Marcus, M. and Sant’Anna, P. H. C. (2021). The role of parallel trends in event study settings: An application to environmental economics. *Journal of the Association of Environmental and Resource Economists*, 8(2):235–275.
- McDermott, G., Kluser, F., and Morawetz, U. (2025). *etwfe: Extended Two-Way Fixed Effects*. R package version 0.6.2.
- Murray, B. C. and Rivers, N. (2015). British columbia’s revenue-neutral carbon tax: A review of the latest “grand experiment” in environmental policy. *Energy Policy*, 86:674–683.
- Paul, A., Myers, E., and Palmer, K. (2009). A partial adjustment model of u.s. electricity demand by region, season, and sector. Technical Report RFF DP 08-50, Resources for the Future, Washington, DC.
- Pretis, F. (2022). Does a carbon tax reduce co2 emissions? evidence from british columbia. *Environmental and Resource Economics*, 83:115–144.
- Price, C. R., Nimbalkar, S. U., and Wenning, T. J. (2020). Energy intensity baselining and tracking guidance. Technical Report ORNL/SPR-2020/1566, Oak Ridge National Laboratory, Oak Ridge, TN.
- Rivers, N. and Schaufele, B. (2015). Salience of carbon taxes in the gasoline market. *Journal of Environmental Economics and Management*, 74:23–36.
- Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322.
- Statistics Canada (2026a). Table 17-10-0005-01: Population estimates on july 1, by

- age and gender. Data table. DOI: 10.25318/1710000501-eng. Accessed May 16, 2026.
- Statistics Canada (2026b). Table 25-10-0029-01: Supply and demand of primary and secondary energy in terajoules, annual. Data table. DOI: 10.25318/2510002901-eng. Accessed May 16, 2026.
- Statistics Canada (2026c). Table 36-10-0222-01: Gross domestic product, expenditure-based, provincial and territorial, annual. Data table. DOI: 10.25318/3610022201-eng. Accessed May 16, 2026.
- U.S. Energy Information Administration (n.d.). Degree days. Energy Explained: Units and Calculators. Accessed May 26, 2026.
- Wooldridge, J. M. (2025). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. *Empirical Economics*, 69(5):2545–2587.
- Xiang, D. and Lawley, C. (2019). The impact of british columbia’s carbon tax on residential natural gas consumption. *Energy Economics*, 80:206–218.
- Xu, J. (2024). The role of carbon pricing in food inflation: Evidence from canadian provinces. Working paper, University of Alberta.
- Yamazaki, A. (2017). Jobs and climate policy: Evidence from british columbia’s revenue-neutral carbon tax. *Journal of Environmental Economics and Management*, 83:197–216.

Appendix A

Additional Robustness Results

Unless otherwise stated, appendix figures report estimated effects on log industrial energy consumption. They are included as supplementary results for the robustness checks discussed in Chapter 7.

A.1 One-Year Anticipation Cohort Results

Table A.1 reports the cohort-specific estimates from the one-year anticipation specification. The main text focuses on the aggregate and dynamic results, while the cohort-level estimates are reported here as supplementary evidence.

Table A.1: Cohort Effects with One-Year Anticipation

Treatment cohort	Estimate	Std. error	95% CI lower	95% CI upper
2007	0.6637	0.2315	0.2101	1.1174
2008	0.7559	0.1912	0.3812	1.1306
2013	0.2763	0.0549	0.1687	0.3839
2017	-0.3458	0.1317	-0.6038	-0.0877

Notes: The 2019 cohort is not reported because the not-yet-treated comparison group is exhausted once all provinces become treated.

Figure A.1 shows the same cohort estimates graphically.




Figure A.1: Cohort Effects with One-Year Anticipation
Notes: Estimates are based on the one-year anticipation specification.

A.2 Doubly Robust Specification

Figure A.2 reports the dynamic treatment effects from the doubly robust specification.

Figure A.3 reports the corresponding cohort-specific effects.

A.3 Log-Weather Specification

Figure A.4 reports the dynamic treatment effects from the specification that uses logged weather controls.

Figure A.5 reports the corresponding cohort-specific effects.

A.4 Supplementary ETWFE Results

Figure A.6 reports the cohort-specific effects from the alternative ETWFE specification.




Figure A.2: Dynamic Effects from the Doubly Robust Specification
Notes: The specification uses the doubly robust estimator with the same sample, controls, and comparison group as the baseline model.




Figure A.3: Cohort Effects from the Doubly Robust Specification
Notes: Cohort-specific estimates are based on the doubly robust specification.




Figure A.4: Dynamic Effects with Logged Weather Controls
Notes: The specification replaces *HDD18* and *CDD18* with their logarithmic forms.




Figure A.5: Cohort Effects with Logged Weather Controls
Notes: Cohort-specific estimates are based on the log-weather specification.




Figure A.6: Cohort Effects from the Alternative ETWFE Specification
Notes: Standard errors are clustered at the province level.

Figure [A.7](#) reports the calendar-year effects from the same ETWFE specification.




Figure A.7: Calendar-Year Effects from the Alternative ETWFE Specification

Notes: Estimates are from the supplementary ETWFE specification.