Three Essays in Macroeconomics

Three Essays in Macroeconomics

By Oliver Loertscher B.Sc. M.A.

A Thesis Submitted to the School of Graduate Studies in the Partial Fulfillment of the Requirements for the Degree Doctor of Philosophy

McMaster University
Doctor of Philosophy (2025)
Hamilton, Ontario (Economics)

TITLE:
Three Essays in Macroeconomics

AUTHOR:
Oliver Loertscher (McMaster University)

SUPERVISORS:
Dr. Pau S. Pujolas
Associate Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

Dr. Gajendran Raveendranathan
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

SUPERVISORY COMMITTEE MEMBERS:
Dr. Bettina Brüggemann
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

Dr. Zachary L. Mahone
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

NUMBER OF PAGES: viii, 88

Abstract

This thesis consists of three papers in macroeconomics that investigate the following questions: (1) How do changes in global demand for fossil fuels affect welfare across households in small resource-rich economies? (2) How did the expansion of the oil sector in Canada affect measured aggregate productivity? (3) How sensitive are cross-country comparisons to measurement errors introduced by nominal-to-real conversions? Chapter 1 develops a quantitative model of a fossil fuel exporting economy to show that the oil price boom between 1997 and 2020 increased welfare among young, low-income households between 11% and 16%. I then simulate the transition to a Net Zero world between 2020 and 2050 and show that while the fall in global demand for fossil fuels reduces lifetime consumption by 0.56% (between 0.49% and 0.77% for the youngest low-income households), the growth of the clean energy sector can dampen these losses by 15% to 54% depending on the speed of the expansion. Chapter 2, co-authored with Pau Pujolas, studies the observed stagnation of Canadian Total Factor Productivity (TFP) between 2000 and 2018. We find that the entirety of the slowdown can be accounted for by the expansion of the oil sector, due to the massive capital investments that occurred. Comparing TFP growth in the rest of the economy to the United States, we find that Canadian TFP grew at comparable rates over the same period. Chapter 3, also co-authored with Pau Pujolas, explores how conclusions drawn from comparing GDP per capita of developed economies relative to the United States differ significantly depending on if current- or constant-Purchasing power parity (PPP) metrics are used. Using data from the Organization for Economic Co-operation and Development (OECD), we first document the differences in the evolution of GDP per capita relative to the US in current-PPPs and constant-PPPs before demonstrating in a numerical example how the choices made in constructing real metrics of GDP to make cross-country comparison can lead to contradictory interpretations.

Acknowledgements

Completing the Ph.D. has been a long journey with many unexpected twists and turns. I would like to express my sincerest gratitude to everyone that played a part in supporting, pushing, and lifting me up along the way.

First and foremost, I am incredibly grateful to my co-supervisors, Pau Pujolas and Gajendran Raveendranathan and my supervisory committee members, Bettina Brüggemann and Zachary Mahone, for their tireless support and guidance over the years. Each of them went above and beyond their roles as mentors. To Pau, I am grateful for his gamble in taking me on as his student. I am extremely fortunate to have found an advisor that gave me the opportunity to prove myself as a co-author, welcomed me into his home, and spent countless hours counselling me on matters both professional and personal. I do not think I could have seen this dissertation through to the end without his support and belief in me. The same can be said for Gajen, who helped me through the challenges of conducting computationally intensive research and never ceased to show enthusiasm for what I was working on, even when I was losing mine. I am also incredibly grateful to Bettina and Zach for their patience as I struggled to narrow my focus, for their thoughtful and constructive questions, and for their tireless work to help me improve in all aspects of conducting and communicating my research.

I am also thankful for the many incredible professors at McMaster beyond those on my supervisory committee. Many of them took the time to attend my presentations and contribute to my growth. At the top of this list I must highlight the incredible contributions of Professor Jevan Cherniwchan, who truly went above and beyond in helping me get across the finish line. I am also indebted to Professors Alok Johri, Arthur Sweetman, Adam Lavecchia, Angela Zheng, Jeffrey Racine, Marc-André Letendre, and the many other professors who were so generous with their time and advice. I also extend my thanks to Cynthia Zhao, Emma Beamson, Lihua Qian, and Mena Petta-Lorenzini for their assistance with the under appreciated complexities of navigating administrative tasks, as well as Peter Kitchen and Li Wang for all their assistance in guiding me through working in the Statistics Canada Research Data Centre.

What I will remember most fondly of all are the friends I met and often commiserated with. Having found myself in the odd position of only spending one full semester getting to know my peers in-person before Covid struck, I am incredibly grateful for the bond I share with my colleagues. Ryan Bacic, Thomas Palmer, Andrew Leal, and Saleha Jafri were each instrumental in bringing some light and camraderie to this challenging endeavour, and I feel blessed to have crossed paths with each of them.

Finally, I want to express my heartfelt gratitude to my wife Kei and my son Felix. They were my biggest cheerleaders when I needed it most. To Kei, thank you for being by my side every step of the way, for picking me up at each stumble, and for the light that you are.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Oliver Loertscher, declare that this thesis titled, *Three Essays in Macroeconomics* and the work presented in it are my own. Chapter 1 is a solo-authored paper for which I independently collected the data, performed the analysis, and wrote the manuscript. Chapter 2 and Chapter 3 are co-authored with Pau Pujolas. For each of these papers I collected the data, and conducted all of the empirical analysis. We contributed equally to the writing of the manuscripts. Chapter 2 has been published in the Canadian Journal of Economics. The paper can be found online at `https://onlinelibrary.wiley.com/doi/abs/10.1111/caje.12707`

Introduction

The transition to a clean economy, productivity growth, and the evolution of national income are all issues at the forefront of the macroeconomics literature with important policy implications. The apparent slowdown in productivity and economic growth among developed economies is at the heart of the policy debate in many developed economies. These issues are compounded by the uncertainty surrounding a global transition to clean energy in resource-rich economies such as Canada.

This thesis contributes to each of these areas through three independent yet related papers that study welfare and economic well-being through three lenses. Each chapter of this thesis builds on our understanding of modern macroeconomics, combining cutting edge quantitative tools with robust and consistent practices of understanding the National Accounts. I study how changes in the global demand for fossil fuels and the expansion of clean energy production affects lifetime consumption across the age, income, and regional distributions in resource-rich economies such as Canada, how the expansion of oil production in Canada has affected measured productivity growth since 2000, and how the standard practice of making cross-country comparisons by converting economic data to a "constant-price" common currency yields a wide range of conclusions to be drawn.

Chapter 1, titled "Small Resource-Rich Economies and the Green Transition", studies the welfare impacts across the age and income distributions and across regions in a small open resource-rich economy of a global adoption of clean energy and declining demand for fossil fuels. This contributes to a rich, expanding literature using quantitative macroeconomic models to study a transition to clean energy and away from fossil fuels. I build a model of a small open economy, with an embedded overlapping generations model with heterogeneous households, multiple regions and multiple production sectors. The price of one of the sectors, the fossil fuel sector, is treated as exogenous to the model. I quantify the distributional impacts of the commodities boom in the late 1990's and early 2000's before measuring the welfare changes of a global decline in the demand for fossil fuels and growth in clean energy production. Calibrating the model to the Canadian economy, I find that the observed rise in oil prices can account for differences in the growth of wages, consumption and GDP between regions. This boom was highly beneficial, particularly for young, low-income households in the fossil extracting region whose lifetime consumption increased by 15 percent. The primary channel that explains the distributional aspect of my results is through wages and profits from the fossil fuel sector. As the price of oil increases, wages in the fossil producing region (Region F) increase due to the expansion of the fossil industry. As a result, young households experience higher lifetime labour earnings. Similarly, the lowest income households benefit the most from the higher transfers of profits from the fossil sector, which act as a crucial stipend for their incomes.

In a forward-looking exercise, I measure the impact of a transition to a Net Zero world over a 30-year period. Through the lens of the model, this is captured by a 50 percent fall in the global demand for fossil fuels, captured by a fall in the price, and a rapid expansion in clean energy productivity by 2050, captured by an exogenous growth in the productivity of the clean energy sector. The decline in demand for fossil fuels produces welfare losses across all households. Younger households at the bottom of the income distribution and residing in the fossil producing region experience the largest losses with a decrease in lifetime consumption of 0.73 percent. However, an accelerated expansion in clean energy productivity can dampen these losses by as much as 54 percent, depending on the speed of the expansion.

The mechanism driving the welfare losses are understood as the same ones that produced the gains in the first quantitative experiment: as the global demand for fossil fuels falls, the industry contracts. This results in lower profits transferred back to the households and weaker wage growth during the transition. However, the expansion of the clean energy sector produces a spillover effect across the economy and across regions. First, the productivity growth induces a fall in the relative price of clean energy, leading to a substitution effect in aggregate energy production as the demand for clean inputs increases. Driven by this fall in the relative price, the price of the aggregate energy good in the economy falls, inducing an expansion in the final good sector. As a result, there is increased demand for the non-energy goods produced in each region. The gains are unequally distributed, however. The lowest-income households experience the lowest benefit from the clean sector expansion and still experience welfare losses over the transition, while the households at the top of the income distribution are the primary beneficiaries.

Chapter 2, titled "Canadian Productivity Growth: Stuck in the Oil Sands", coauthored with Pau Pujolas and published in the Canadian Journal of Economics, explores the observed stagnation in Canadian total factor productivity (TFP) from 2000 onward. While there is a rich literature exploring various explanations for this stagnation, we find that the entirety of the slowdown can be explained by the expansion of the oil sector. From 1997 onward, there was a massive shift of capital into oil extraction, largely due to technological progress and the rise in international oil prices which made industrial scale production of the Alberta oil sands profitable. We show that when we measure TFP in the rest of the economy outside the oil sector, the observed stagnation disappears. To reinforce our results, we first show how this result is supported at the provincial level by comparing TFP growth in Alberta and Ontario. We show that Alberta's productivity grew between 1997 and 2005 before plummeting through 2018, eliminating any productivity gains. However, these trends are eliminated when focusing on the Net-of-Oil economy.

Next, we compare aggregate and Net-of-Oil TFP growth in Canada, Norway, and the

United States. We found that in the US, there was very little difference in aggregate and Net-of-Oil TFP growth, and between 2001 and 2018 Net-of-Oil productivity in both Canada and the US were nearly identical. When comparing Canada with Norway, a comparable small and resource-rich economy, we find the same results that Net-of-Oil productivity growth was stronger than what the aggregate measurement suggests. In both countries, the growth of the oil sector results in a massive growth in the percentage of the aggregate capital stock allocated to oil production. Canada proves a unique example where the growth far outpaced the overall percentage of GDP accounted for by oil production. Oil production between 1961 and 2018 accounts for roughly 5 percent of GDP. On the other hand, over this period the capital stock in the oil sector account grew from 6 percent to approximately 31 percent.

Chapter 3, titled "What are we deflating for?" and co-authored with Pau Pujolas, studies how conclusions drawn from comparing the evolution of GDP per capita among developed economies relative to the United States differ markedly depending on whether researchers analyze the time series measured in nominal PPPs or chain volume indexed PPPs. When considered separately, chain volumes or real measures are useful for comparing the evolution of economic output within a country over time, while PPP measures are useful to compare two countries' output in a particular year. However, real-PPP metrics introduce a bias that appears to favour the United States.

To motivate our findings, we present the evolution of GDP per capita for each G7 country relative to the United States measured in nominal and chain volume indexed PPPs. We show that for each country, the conclusions drawn about relative economic performance are very different depending on which time series we observe. Except for Japan, the difference between the nominal and real values of the observed decline in relative GDP per capita is as high as 10 percentage points, a stark difference. We interpret this observation as suggesting that due to the dual approximations of converting price bases into a real measure in a common currency, the real valued PPP series tends to overstate the relative declines. To further strengthen the observation, we also demonstrate the identical trend in Net National Income (NNI). Focusing on NNI emphasizes the main point that the conclusions drawn are entirely dependent on which series we observe: in the case of Germany the evolution of relative NNI in nominal terms appears roughly constant between 2010 and 2020, while in real terms relative NNI declined by nearly 22 percentage points.

Through a series of numerical experiments, we demonstrate how the process of converting measured GDP in two economies into a common set of base prices determines the qualitative interpretation of the relative performance of two countries. Using a simple two period model of two economies, we show that measuring economic performance in real PPPs can show that an economy contracted, while real GDP (not in PPPs) grows. Across a series of examples meant to demonstrate the robustness of our results, we show how the choice of base year and the process of constructing the Chain Volume Index can lead to contradictory conclusions being drawn about the performance of the economy.

Across each robustness check, we find that the only qualitatively consistent story is the comparison of relative GDP in nominal PPPs.

These three papers advance our understanding of how changes in the global economy affect welfare across households and highlight the challenges in measuring these impacts. They emphasize the channels through which the contraction and expansion of particular industries can interact and spillover across the economy, and how the choices we make as practitioners to measure and compare these patterns influence the conclusions we draw. While each paper addresses a distinct question, they form a unified, cohesive body of work focused on three distinct but related channels of economic well-being. The first paper focuses broadly on welfare as understood through changes in lifetime consumption, focusing on the impact of the global move to a Net Zero world in small, resource-rich economies. This paper shines a light on how changes in global demand for the output in important sectors and the productivity growth in the production of close substitutes is transmitted through the economy and across different households, producing important policy implications. The second paper continues to pull on the thread of productivity and the special case of oil production in Canada. This paper sheds light on the challenges of drawing conclusions from aggregate data and sheds some insight on making cross-country comparisons. The third paper further expands on this by digging deeper into the problems that arise when comparing relative economic performance rather than growth rates. Each paper focuses on a different measure of economic well-being (consumption, productivity, income) to develop a unified body of knowledge and significant contribution to the literature.

This thesis offers a thorough study of vital issues at the forefront of the modern macroeconomic literature related to changing global demand during energy transitions, capital formation, productivity growth, and measures of national income for cross-country comparisons. The result is a collection of evidence-based insights using cutting edge methods to assist policy makers in understanding and navigating the changing economic landscape.

Chapter 1

Small Resource-Rich Economies and the Green Transition

Small Resource-Rich Economies and the Green Transition

Abstract

The global transition to clean energy carries significant implications for countries with large resource extraction sectors. I quantify the welfare impact of this transition for a small, resource-rich country. I develop a quantitative, overlapping generations model of a multi-region, multi-sector, small open economy with heterogeneous households featuring both fossil and clean energy production and calibrate it to the Canadian economy. In a backward-looking exercise, I find that the global rise in oil prices from 1997 to 2020 can quantitatively account for wage, consumption, and GDP growth differences across Canadian regions with and without fossil fuel extraction. In forward-looking counterfactual experiments, I quantitatively evaluate the effects of a global transition to clean energy as laid out in the International Energy Agency's path to a Net Zero world by 2050. Through the lens of the model, this scenario calls for a 50% reduction in global fossil fuel demand and an expansion of the clean energy sector to account for 90% of domestic energy production. I find that the reduction in international demand for fossil fuels decreases lifetime consumption by 0.56% overall. These losses are most pronounced among young, low-income households in the fossil extracting region, whose lifetime consumption decreases by 0.77%. The expansion of the clean energy sector dampens the aggregate losses by 15%. An immediate expansion in the clean sector reduces the losses by 54%.

1.1 Introduction

There is strong evidence to suggest that we are witnessing a transition towards clean energy. The cost per kilowatt-hour of clean energy sources has plummeted, while global adoption of clean and renewable energy sources is increasing. Almost every major country has introduced policies and signed international agreements to reduce dependence on fossil fuels. Concerns over carbon emissions and fossil fuel dependency have led to the implementation of carbon taxes and other efforts to curtail demand. This poses a unique set of challenges for fossil-exporters as they face diminishing demand, both internationally and domestically, in a key industry. A boom in commodity prices at the turn of the 2000s spurred massive investment in the resource extraction sectors of several economies. Nations such as Australia, Canada, Nigeria, and Norway (among others) derive a large fraction of their economic production and export revenues from fossil fuel extraction.

This paper asks how changes in the global demand for fossil fuels affect welfare across the age and income distributions in small, open, resource-rich economies. While the rise in oil prices between 1997 and 2020 particularly benefited young households at the bottom of the income distribution, a global fall in demand for fossil fuels produces an inverted outcome. Using the International Energy Agency's recommendations for attaining a Net Zero world by 2050, I find that the fall in demand for fossil fuels decreases welfare across the age, income, and regional distributions. Overall, lifetime consumption decreases by 0.56%. The largest losses are experienced by young households at the bottom of the income distribution in fossil extracting regions of the economy, primarily due to dampened wage growth, lower returns on savings and the declining profits from the fossil sector. The expansion of the clean energy sector dampens these losses by 15%. In a counterfactual exercise, I find that an immediate expansion of the clean sector dampens the losses due to fossil fuel demand by 54%. Among young, low income households, rapid expansion in the production of clean energy reduces losses by 20% in the fossil producing region and by 85% in the clean producing region.

I develop an overlapping generations model featuring two domestic regions with multiple production sectors to quantitatively assess the welfare effects of the global transition to clean energy. The model features household heterogeneity through age, assets, and productivity differences in the style of Bewley (1977), Imrohoroğlu (1989), Huggett (1993), and Aiyagari (1994). The production structure adapts Fried, Novan, and Peterman (2022), where the final consumption and investment good is produced using energy and non-energy inputs. Regions in the model each produce a unique energy and non-energy intermediate good. The key distinction here is that one region specializes in the production of a fossil input into energy, and the other produces a clean input. Each type of energy intermediate is highly substitutable in energy production, and non-energy intermediates are highly substitutable in non-energy production. Energy and non-energy inputs are modelled as perfect complements to capture the difficulties in substituting between the two over short time horizons. International demand for the fossil good is captured purely through the price, which is exogenous.

I calibrate the parameters to Canadian data and evaluate the welfare effects of the oil price boom between 1997 and 2020. To validate the model, I compare the model output against a number of untargeted empirical moments. The model can replicate several trends in aggregates, in particular the observed fact that wages, consumption, and GDP grew faster in the fossil-producing region than the clean-producing region between 1997 and 2020. Reinforcing the key mechanism of the model, I show that these trends are inverted if the price of the fossil good is held constant, highlighting the importance of the resource boom as a driving factor in this economy. I find that the boom in oil prices produced positive but heterogeneous welfare gains along the transition path. The gains specifically from the rise in the price of the fossil good are more pronounced in the fossil-producing region, particularly among young households and low-income households, whose lifetime consumption increased by 16%. These gains are monotonically decreasing with both age and income. This result is driven by the stronger wage growth, increased return to savings and higher profits from the fossil sector, which boost lifetime income.

Next, I quantify the welfare impacts of a "Green Transition." For the purpose of this paper, I adopt the International Energy Agency's "Net Zero by 2050" recommendations (IEA, 2021). This scenario calls for (1) 90% of the domestic electricity production to be derived from clean sources and (2) for demand for coal and oil falling by 50% from 2020 levels. To mimic this scenario in my model, I recalibrate the productivity in the clean sector and the price of the fossil good so that along the transition path, clean energy grows to account for 90% of inputs into energy production, and fossil production falls to 50% of 2020 levels. I allow TFP in the non-energy sectors to continue to grow along the observed trend between 2010 and 2020. Through the lens of the model, I find that the fall in demand for fossil fuels results in an aggregate decrease of lifetime consumption by 0.56%. These losses are most pronounced in the fossil-producing region and decrease with age and income. This scenario is largely the inverse story of the backward looking exercise: falling demand for fossil fuels (which is reflected in a decrease in the price) dampens growth in wages and the interest rate, and amplifies the fall in fossil sector profits along the transition path. Households at the bottom of the age and income distribution experience the largest losses in this scenario. However, the losses due to falling international demand for fossil fuels are partially offset by the gains from expansion of the clean sector. This Green Transition dampens aggregate losses by 15%. I investigate whether faster growth in the clean sector can further mitigate these losses. I find that an immediate expansion in clean energy sector dampens the aggregate welfare losses from declining demand for fossil fuels by 54%. Among young low-income households, this expansion dampens the losses by 19% in the fossil producing region and by 85% in the clean producing region.

This paper contributes to and expands upon three distinct areas of the literature. First, this paper contributes to the existing literature on clean energy transitions. Many papers focus on the roles of taxes and subsidies in driving a change from a dirty to a clean technology (see Besley and Persson, 2023; Helm and Mier, 2021; Lennox and Witajewski-Baltvilks, 2017; Acemoglu et al., 2016). Recent additions to the literature, such as Arkolakis and Walsh (2023), highlight the welfare losses fossil fuel exporters face in a world moving towards clean energy due to the loss of export revenues. My paper expands on this work by modelling and analyzing the effects of both a fall in global demand for fossil fuels and a rapid increase in clean technology across the age and income distributions, and is the first paper to quantitatively study the macroeconomic impacts of a global transition to a Net Zero world by 2050. This paper also contributes to the literature arguing that the Green Transition will produce heterogeneous outcomes, such as Baldwin et al. (2020), which argues that long-run outcomes of a clean energy transition depend on how easily the dirty capital stock can be converted to clean, and Borenstein and Davis (2016), which finds that U.S. income tax credits aimed at adopting clean technology have mostly benefited high-income Americans. Consistent with this last paper, I find that the top of the income distribution benefits the most over the course of the transition.

Second, this paper contributes to the expanding literature applying quantitative macroeconomic models to environmental and resource economics. A common feature of these models, as highlighted in Hassler, Krusell, and Olovsson (2021, 2022) and Casey (2023), is the inability to substitute between natural resources (or energy inputs) with other productive resources in the short-run. Given this paper's focus on short-run dynamics, I adhere to this approach by modelling energy and non-energy inputs as perfect complements in production. The literature typically employs a representative, infinitely lived agent framework and abstracts the production technology used in the fossil sector, omitting the role of capital and labour. My paper offers a novel contribution in two ways: first, I add an overlapping generations framework with heterogeneous agents. This allows me to produce a richer understanding of how changes in global demand and production of fossil fuels are transmitted across households by age and income. Second, I calibrate the production parameters to capture the fact that fossil extraction is extremely capital intensive, as evidenced in Loertscher and Pujolas (2024), which shows the disproportionate flow of capital into the oil and gas sector in Canada in the 2000s.

Finally, this paper contributes to the large literature concerned with the economic impacts of commodity price booms. Thus far, the existing literature suggests positive income and welfare effects for regions that benefit from expanding fossil fuel extraction. This is true of all fossil production booms, be it coal (see Black et al., 2005), natural gas (see Bartik et al., 2019), and oil (see Michaels, 2011). This paper is consistent with the existing literature, emphasizing that the largest beneficiaries of the fossil price boom were younger, low income households, driven primarily by the impacts of profits from the fossil sector on per capita income. My paper highlights how these benefits are distributed across the age distribution and is able to decompose the welfare gains derived from the rise in commodity prices and the welfare gains derived from productivity gains in other sectors.

The rest of the paper is organized as follows: Section 1.2 summarizes the relevant evidence indicative of a transition towards growing adoption of clean energy. Section 1.3 presents the model, which features an overlapping generations framework with heterogeneous agents, two regions, an energy sector utilizing both clean and dirty inputs, a fossil extraction sector, a clean sector, and non-energy intermediate production. Section 1.4 details the estimation and calibration of the model. Section 1.5 presents the model performance with respect to non-targeted moments and the results of the benchmark transition exercise to quantify the welfare impact of the observed oil price trends from 1997 to 2020. Section 1.6 discusses the model predictions regarding the welfare effects of a clean energy transition, characterized by a drop in global demand for fossil inputs and a decrease in the relative price of clean inputs. Through the lens of the model, this occurs via a fall in the exogenous price of the fossil good, and rapid growth in the productivity of the clean sector Finally, Section 3.4 concludes the paper and highlights areas for future research.

1.2 Empirical Evidence of the Green Transition

The primary motivation of this paper is outlined in the facts presented in this section. Nearly every nation is signalling an intent to dramatically reduce their emissions by 2050. The price of the technology to substitute away from fossil fuels has plummeted in the last decade, and there are signs that key players in this transition are beginning this substitution towards clean energy. For small, resource-rich economies, this foreshadows a significant contraction in a key sector.

There are currently 120 nations that have some commitment to Net Zero status in writing, whether in an official policy document or actively signed into law. An additional 70 countries have proposed or pledged (in some non-binding manner) to attain Net Zero status. While pledges and proposals are not enforceable or binding agreements, there is a consensus among almost every country that reducing or removing emissions is an urgent priority. Table 1.1 summarizes the state of Net Zero commitments internationally. The numbers are sourced from data taken from "The Net Zero Tracker," a collaborative project between the Energy & Climate Intelligence Unit, the Data-Driven EnviroLab, the NewClimate Institute, and the Oxford Net Zero project, which collects data on targets and progress towards emissions reduction and net zero status among all nations, as well as cities and companies. The concept of Net Zero refers to the reduction or removal of greenhouse gas emissions from the earth's atmosphere.

Table 1.1: Net Zero Progress among all nations

	Proposed	In policy document	Pledged	In law
Number of countries	53	87	17	33

Most of these countries have also highlighted 2050 as an important benchmark year. 50% of countries that have an official policy document or enacted a law and 84% of those who have pledged or proposed a Net Zero goal have signalled 2050 as the target date. Meeting this deadline requires rapid improvements in the costs and efficiency of non-fossil fuel energy sources, and for these emerging technologies to be adopted.

Figure 1.1 tracks the fall in the levelized cost of energy (LCOE) measured in 2022 US dollars per kilowatt-hour from a number of renewable energy sources. LCOE is a measure used to assess the minimum price at which the energy generated can be sold to offset the production costs. The data presented here is taken from the International Renewable Energy Agency (IRENA).

Other than hydroelectricity, the costs of using renewables have been declining steadily. The cost of onshore wind power has been trending downwards for the last 40 years. In the last decade, the LCOE of offshore wind has become comparable to hydro. The most dramatic improvements are in solar power. Photovoltaic solar power, which absorbs sunlight and converts it directly into electricity, has seen a particularly stark improvement in recent years. The price per kilowatt-hour has fallen from 0.44$ to 0.05$. This also coincides with a dramatic fall in the price of a single solar PV panel. This is indicative of a trend in which clean energy is becoming increasingly cost-effective and cheaper to adopt as well.

Concurrent with this fall in costs, many large countries have been increasing their use of clean energy in primary energy consumption. Figure 1.2 tracks the growth in clean sources as a percentage of all terawatt-hours of primary energy consumed. For

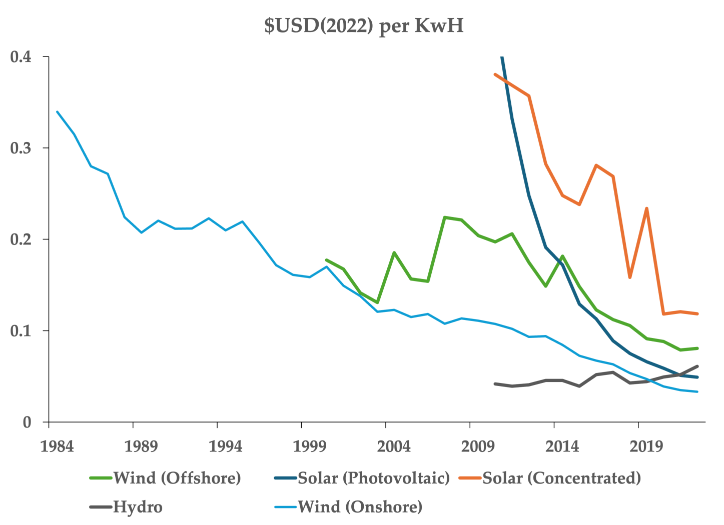

FIGURE 1.1: Levelized costs of clean energy
***Source**: International Renewable Energy Agency*

the purpose of this paper, any reference to clean energy implies renewables (solar, wind, hydro) and nuclear energy. Any reference to fossil fuels or oil implies oil, gas, and coal. The data is taken from Our World in Data, who processed the data from the Energy Institute's Statistical Review of World Energy (2024). The period from 2012 to 2022 shows an upward trend (to varying degrees) among a number of major countries. Japan had a sharp decline at the time of the 2011 Fukushima nuclear accident but has since steadily increased the share of energy derived from clean sources. China has accelerated clean adoption over this period, nearly doubling from 9% to 18%. India and the United States have also increased their shares, albeit at a slower pace. In 2012, clean sources accounted for 8% of India's primary energy consumption and 14% in the United States. In 2023, those numbers had grown to 10% and 17%, respectively. Germany shows the most consistent growth in clean energy adoption since 1997, growing from 14% to 18% by 2012, and to 24% by 2023.

These five countries account for a significant fraction of global emissions (in particular the United States, China and India), and have the largest ability to affect global demand for fossil fuels if they maintain or accelerate these trends. They are also responsible for the lion's share of the international demand among small resource-rich countries. "Small" in this context refers to the country's ability to influence the world price of an export good. Australia, Canada, Nigeria, and Norway (among others) are significant exporters of fossil fuels. Figure 1.3 plots fossil fuel exports as a percentage of the value of all exports for each country. Among this set of nations, fossil fuels account for at least 20% of all export revenue. If global demand for fossil fuels falls, these countries and other small fossil exporters are significantly exposed to potential losses.

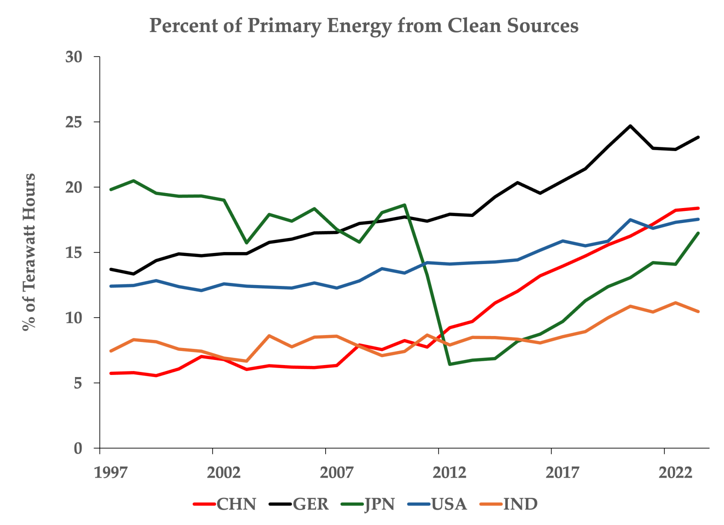

FIGURE 1.2: Adoption of clean energy
Source: *Our World in Data and author's own calculations*

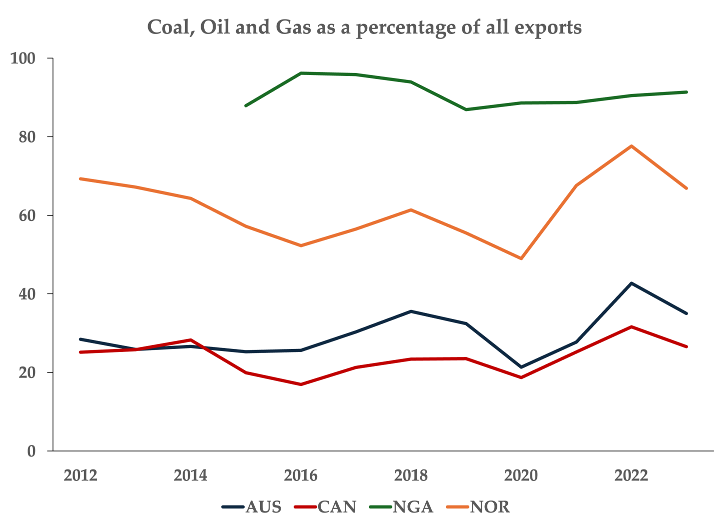

FIGURE 1.3: Coal, Oil, and Gas as a share of exports
Source: *UN Comtrade and author's own calculations*

1.3 Model

In this section, I build a Small Open Economy model with two regions and multiple sectors. Households follow an overlapping generations structure with heterogeneity due to both age-specific productivity and survival probability differences, as well as idiosyncratic productivity differences. On the production side, the economy is organized into a final good producer, which uses non-energy and energy intermediate inputs, an energy sector which uses clean and fossil intermediate inputs, a clean sector which produces the clean energy input using a capital-labour composite, a fossil sector which produces the fossil intermediate using a capital-labour composite and a fossil intermediate, and region-specific non-energy intermediate producers who use capital and labour. The price of the fossil good is assumed to be exogenous. Both the fossil good and the final consumption good can be bought and sold on the international market, and trade is balanced every period. The structure of the model adds to the use of overlapping generations models in the climate literature, as in Laurent-Lucchetti and Leach (2011) and Jaimes (2023), both of which study climate policy in an OLG setting. On the production side, this paper builds on the structures of Fried et al. (2022) and Casey (2023), by explicitly modelling the clean and fossil sectors, and assuming the extreme case that final good production employs a Leontief structure where energy and non-energy inputs are complements.

1.3.1 Regions

There are two regions in the model denoted by $s \in \{F, C\}$. The regions each produce a specialized energy and non-energy intermediate good. The key distinction is region F produces the fossil energy intermediate (denoted by $\mathcal{F}$) and region C produces the clean energy intermediate (denoted by $\mathcal{C}$). The fossil good, whose price $p_{\mathcal{F}}$ is determined exogenously, is sold both domestically and internationally. The clean good and the two non-energy goods are both produced and sold domestically.

1.3.2 Households

Households in each region supply labour inelastically and maximize expected lifetime utility given their age j, current level of assets a, and idiosyncratic productivity ϵ by choosing consumption c and next period assets a'. I assume that there is no labour mobility between regions, but capital is perfectly mobile. Hence, wages are region specific but the interest rate is common across regions. The problem of the household is given by:

$$V(j, a, \epsilon, s) = \max_{a', c} \frac{c^{1-\sigma}}{1-\sigma} + \psi_j \beta \mathbb{E}_{\epsilon'|\epsilon} V(j+1, a', \epsilon', s) \tag{1.3.1}$$

subject to:

$$(1+\tau_{c,s})c + a' = \begin{cases} w_s \epsilon \theta_j + (1+r)(a+b) + \pi_s - T(y) & if\ j < j_{ret} \\ ss + (1+r)(a+b) + \pi_s - T(y) & if\ j \geq j_{ret} \end{cases}$$

$$c, a' \geq 0$$

where time subscripts t are suppressed for ease of notation.

Profits π_s for $s \in \{F, C\}$ correspond to the profits from the energy intermediate producers specific to each region. A household living in region F receives the profits from the fossil producers $\pi_{\mathcal{F}}$, while a household living in region C receives the profits from the clean producers $\pi_{\mathcal{C}}$.

Agents enter at $j = 1$ and live until a maximum age of J years. Agents all retire exogenously at age j_{ret}. From ages $j = 1$ to $j_{ret} - 1$, agents supply one unit of labour inelastically at the market wage rate for their region w_s. Retired households receive no labour income and collect social security benefits ss. At the beginning of each period, agents are transferred accidental bequests b, which consist of the assets of agents who did not survive from period $t - 1$ to t. These bequests are uniformly distributed among the remaining agents in both regions and added to their existing asset holdings a. Agents aged $j = 1$ enter with the bequest transfers as their initial endowment of assets.

Household idiosyncratic productivity ϵ is assumed to follow an AR(1) process so that $\log \epsilon' = \rho_\epsilon \log \epsilon + e$ where $e \sim N(0, \sigma_\epsilon^2)$. There is also an age-specific productivity component θ_j common to all agents with age j. An individual aged j survives to $j + 1$ with probability ψ_j. Agents who do not survive to the next period have their assets uniformly redistributed in the form of accidental bequests b.

Consumption choices c are taxed at a rate $\tau_{c,s}$ depending on the region. Following Heathcote et al. (2017), market income is taxed progressively according to the tax function:

$$T(y) = y - \lambda y^{1-\tau_p} \tag{1.3.2}$$

where:

$$y = \begin{cases} w_s \epsilon \theta_j + (1+r)(a+b) + \pi_s & if\ j < j_{ret} \\ (1+r)(a+b) + \pi_s & if\ j \geq j_{ret} \end{cases} \tag{1.3.3}$$

The parameter τ_p governs the tax progressivity, and λ is determined in equilibrium to balance the government's budget each period. The income tax function is assumed to be the same across each region.

1.3.3 Government

The government consumes the final good and pays social security transfers. These expenditures are financed via consumption taxes and income taxes. The government budget is balanced each period so that:

$$SS + G = \tau_{c,F} C_F + \tau_{c,C} C_C + \int T(y)\ d\Omega(j, a, \epsilon, s) \tag{1.3.4}$$

where $\Omega(j, a, \epsilon, s)$ is the distribution across all household types in the economy, and C_F and C_C are the aggregated consumption decisions in each region.

Government consumption G is assumed to be a fraction of the total output of the final consumption good, so that $G = gY$, $g \in (0, 1)$. The government issues no bonds and incurs no debts.

1.3.4 Production

Intermediate production in this economy is region-specific. Each region specializes in producing a non-energy intermediate good and an energy intermediate good. Non-energy inputs are denoted with subscripts $\{F, C\}$, while energy intermediate inputs are denoted with calligraphic subscripts $\{\mathcal{F}, \mathcal{C}\}$. Intermediates are then aggregated into the energy good and the final good. The energy good (denoted with $\mathcal{E}$) uses region specific energy intermediates $x_{j,\mathcal{E}}$ for $j \in \{\mathcal{F}, \mathcal{C}\}$. Final good production uses non-energy intermediates $x_{s,Y}$ for $s \in \{F, C\}$ and the aggregated energy good $x_{\mathcal{E},Y}$.

Fossil Producers

The fossil energy good $Y_{\mathcal{F}}$ is produced in region F using capital $K_{\mathcal{F}}$, labour $L_{\mathcal{F}}$, and fossil intermediate $x_{\mathcal{F},\mathcal{F}}$ with a decreasing returns to scale technology governed by $\nu_{\mathcal{F}} \in (0, 1)$. Fossil producers take the world price $p_{\mathcal{F}}$ as given each period and make their production decisions accordingly. The problem of a fossil producer is given by:

$$\max_{K_{\mathcal{F}}, L_{\mathcal{F}}, x_{\mathcal{F},\mathcal{F}}} p_{\mathcal{F}} Y_{\mathcal{F}} - (r + \delta) K_{\mathcal{F}} - w_F L_{\mathcal{F}} - p_{\mathcal{F}} x_{\mathcal{F},\mathcal{F}} \tag{1.3.5}$$

subject to:

$$Y_{\mathcal{F}} = \min\{Z_{\mathcal{F}} (K_{\mathcal{F}}^{\gamma} L_{\mathcal{F}}^{1-\gamma})^{\nu_{\mathcal{F}}}, \mu_{\mathcal{F}} x_{\mathcal{F},\mathcal{F}}\}$$
$$\nu_{\mathcal{F}} \in (0, 1)$$

Profits $\pi_{\mathcal{F}} = p_{\mathcal{F}} Y_{\mathcal{F}} - (r + \delta) K_{\mathcal{F}} - w_F L_{\mathcal{F}} - p_{\mathcal{F}} x_{\mathcal{F},\mathcal{F}}$ are distributed back to the households in region F uniformly. Here the Leontief production structure captures the costs of extraction, and the decreasing returns to scale capture the notion that scaling production is costly. As the world price $p_{\mathcal{F}}$ rises, fossil producers will want to increase production to take advantage. However, they are constrained in that they need to use some fraction of their own output, $x_{\mathcal{F},\mathcal{F}}$ to produce the desired output. The decreasing returns to scale parameter $\nu_{\mathcal{F}}$ captures the notion that each marginal unit of fossil production requires more inputs to produce.

Clean Energy Producers

The clean energy good $Y_{\mathcal{C}}$ is produced in region C using capital $K_{\mathcal{C}}$ and labour $L_{\mathcal{C}}$. Clean producers make production choices to maximize profits according to:

$$\max_{K_{\mathcal{C}},L_{\mathcal{C}}} p_{\mathcal{C}}Y_{\mathcal{C}} - (r+\delta)K_{\mathcal{C}} - w_C L_{\mathcal{C}} \tag{1.3.6}$$

subject to:

$$Y_{\mathcal{C}} = Z_{\mathcal{C}}(K_{\mathcal{C}}^{\eta} L_{\mathcal{C}}^{1-\eta})^{\nu_{\mathcal{C}}}$$
$$\nu_{\mathcal{C}} \in (0,1)$$

Clean profits $\pi_{\mathcal{C}} = p_{\mathcal{C}}Y_{\mathcal{C}} - (r+\delta)K_{\mathcal{C}} - w_C L_{\mathcal{C}}$ are redistributed uniformly to households in region C.

Non-Energy Producers

The non-energy good Y_s produced in each region $s \in \{F, C\}$ uses capital K_s and labour L_s. The problem of a non-energy producer is given by:

$$\max_{K_s,L_s} p_s Y_s - (r+\delta)K_s - w_s L_s \tag{1.3.7}$$

subject to:

$$Y_s = Z_s K_s^{\alpha} L_s^{1-\alpha}$$

Energy Producers

The final energy good $Y_{\mathcal{E}}$ is produced by combining the Clean energy intermediate $x_{\mathcal{C},\mathcal{E}}$ and the Fossil energy intermediate $x_{\mathcal{F},\mathcal{E}}$, both of which are region-specific. The problem of the energy producer is:

$$\max_{x_{\mathcal{C},\mathcal{E}} x_{\mathcal{F},\mathcal{E}}} p_{\mathcal{E}}Y_{\mathcal{E}} - p_{\mathcal{C}}x_{\mathcal{C},\mathcal{E}} - p_{\mathcal{F}}x_{\mathcal{F},\mathcal{E}} \tag{1.3.8}$$

subject to:

$$Y_{\mathcal{E}} = \left(x_{\mathcal{C},\mathcal{E}}^{\rho_{\mathcal{E}}} + x_{\mathcal{F},\mathcal{E}}^{\rho_{\mathcal{E}}}\right)^{1/\rho_{\mathcal{E}}}$$

The CES production function is chosen to capture the imperfect substitutability of clean and fossil inputs in the production of energy. This structure allows the model to produce a mix of aggregate energy that mirrors what is observed in the data, and follows the structure in Fried et al. (2022). The final energy mix depends on the relative prices between fossil and clean inputs.

Final Good Producer

The final good Y is produced using region-specific non-energy intermediates $x_{F,Y}$, $x_{C,Y}$, and the energy good $x_{\mathcal{E},Y}$. The final good producer takes prices for each intermediate as given and makes input decisions according to:

$$\max_{x_{F,Y},x_{C,Y},x_{\mathcal{E},Y}} p_Y Y - p_F x_{F,Y} - p_C x_{C,Y} - p_{\mathcal{E}} x_{\mathcal{E},Y} \tag{1.3.9}$$

subject to:

$$Y = \min\left\{x_{N,Y}, \mu_{\mathcal{E}} x_{\mathcal{E},Y}\right\}$$
$$x_{N,Y} = \left(x_{F,Y}^{\rho_Y} + x_{C,Y}^{\rho_Y}\right)^{1/\rho_Y}$$

Here, the Leontief production function stands in for the high degree of complementarity between energy and non-energy inputs, treating them as perfect complements. This choice is in line with the insights in Casey (2023), who points out the extreme complementarity in the short run between energy and non-energy inputs. Other papers, such as Fried et al. (2022), model final good production as a CES production function with an elasticity of substitution that implies a high degree of complementarity. Given the short time horizon I am considering for my quantitative analysis, I favour the extreme case of treating the inputs in final good production as perfect complements.

1.3.5 Equilibrium

The equilibrium of the model is defined as follows. A stationary, recursive equilibrium consists of a value function $V(j,a,\epsilon,s)$, decision rules $c(j,a,\epsilon,s)$ and $a'(j,a,\epsilon,s)$ for all j,a,ϵ,s, prices $\{w_F, w_C, r, p_F, p_C, p_{\mathcal{C}}, p_{\mathcal{E}}, p_{\mathcal{F}}\}$, a stationary distribution $\Omega(j,a,\epsilon,s)$, factor demands K_F, K_C, $K_{\mathcal{F}}$, $K_{\mathcal{C}}$, L_F, L_C, $L_{\mathcal{F}}$, $L_{\mathcal{C}}$, $x_{F,Y}$, $x_{C,Y}$, $x_{\mathcal{E},Y}$, $x_{\mathcal{C},\mathcal{E}}$, $x_{\mathcal{F},\mathcal{E}}$, $x_{\mathcal{F},\mathcal{F}}$ and aggregate quantities $Y, Y_{\mathcal{E}}, Y_F, Y_C, Y_{\mathcal{C}}, Y_{\mathcal{F}}$ such that:

1. Household decision rules $c(j, a, \epsilon, s)$ and $a'(j, a, \epsilon, s)$ solve the household problem (1.3.1) given prices $\{w_F, w_C, r, p_F, p_C, p_{\mathcal{C}}, p_{\mathcal{E}}, p_{\mathcal{F}}\}$.
2. The government budget constraint (1.3.4) holds each period.
3. The fossil, clean, non-energy, energy, and final good producers maximize their respective profits given prices.
4. The labour market clears in each region, so that total labour demand equals total labour supply, ie.

$$\int_1^{j_{ret}} \quad d\Omega(j, a, \epsilon, F) = L_F + L_{\mathcal{F}}$$
$$\int_1^{j_{ret}} \quad d\Omega(j, a, \epsilon, C) = L_C + L_{\mathcal{C}}$$

5. The capital market clears such that total capital demand equals total savings, ie.

$$\int a'(j, a, \epsilon, s) \quad d\Omega(j, a, \epsilon, s) = K'_F + K'_C + K'_{\mathcal{F}} + K'_{\mathcal{C}}$$

6. The markets for clean intermediates, non-energy intermediates and energy intermediates clears, ie.

$$Y_{\mathcal{C}} = x_{\mathcal{C},\mathcal{E}}$$
$$Y_F = x_{F,Y}$$
$$Y_C = x_{C,Y}$$
$$Y_{\mathcal{E}} = x_{\mathcal{E},Y}$$

7. Trade is balanced in each period, ie.

$$p_{\mathcal{F}}(Y_{\mathcal{F}} - x_{\mathcal{F},\mathcal{F}} - x_{\mathcal{F},\mathcal{E}}) = Y - C_F - C_C - I - G \tag{1.3.10}$$

where

$$C_i = \int_{s=i} c(j, a, \epsilon, i) \quad d\Omega(j, a, \epsilon, i)$$

for $i \in \{F, C\}$

1.4 Calibration

A period in the model corresponds to 4 years. I calibrate my model in two steps. First, I externally calibrate several parameters. I set $j = 1$ to correspond to an agent who is

20 years old. Retirement age $j_{ret} = 13$ so agents work until their age corresponds to 64 and are retired at 68. Agents do not live past 100. I fix the parameters σ, β, δ, ρ_E and ρ_Y using values from the literature (summarized in Table 1.2). I also externally calibrate and estimate a number of production, government and household age specific parameters directly from the data. These parameters are summarized in Tables 1.3, 1.4, and 1.5. Second, I internally calibrate the idiosyncratic productivity distribution parameters ρ_ϵ and σ_ϵ and the productivity in the fossil and clean sectors $Z_{\mathcal{F}}$ and $Z_{\mathcal{C}}$. I calibrate the parameters governing the idiosyncratic productivity process directly outside of equilibrium, and the results are summarized in Table 1.6. TFP in the fossil and clean sectors are calibrated in equilibrium, summarized in Table 1.7.

1.4.1 External Calibration

In the first step of parametrizing my model, I assume values for 6 parameters, presented in Table 1.2. I assume that the household risk aversion coefficient σ is set to 2, a standard value in the literature. Similarly, I assume an annualized household discount factor of 0.96 so that $\beta = 0.85$ and capital depreciates at an annual rate 0.05 so that $\delta = 0.18$. Following Fried, Novan, and Peterman (2022) and Papageorgiou, Saam, and Schulte (2017), I assume that the elasticity of substitution between clean and fossil inputs in energy production is $\rho_{\mathcal{E}} = 0.66$. Papageorgiou, Saam, and Schulte (2017) report empirical estimates of the elasticity of substitution between clean and dirty inputs in energy production that fall between 0.23 an 0.66. Given the similarities between my model and Fried, Novan, and Peterman (2022), I adopt the higher value. In line with Albrecht and Tombe (2016), who estimate the interprovincial trade elasticity in Canada, I set the elasticity of substitution in the non-energy composite to $\rho_Y = 0.80$.

TABLE 1.2: Assumed parameters

Parameter		Description	Value
σ	Risk Aversion	Standard	2
β	Discount factor	Annualized rate of 0.96	0.85
δ	Capital depreciation	Annualized rate of 0.05	0.18
$\rho_{\mathcal{E}}$	Elasticity of substitution in Energy production	Papageorgiou et al. (2017)	0.66
ρ_Y	Elasticity of substitution in Non-Energy composite	Albrecht and Tombe (2016)	0.80

In the next step, I estimate several parameters directly from the data. The production parameters are presented in Table 1.3. I treat the production technology in this sector as identical between the two regions. The capital share in non-energy production is estimated using data on labour compensation (Table 36-10-0489-01, Statistics Canada) and value added (Table 36-10-0402-01, Statistics Canada) in the Canadian provinces of Alberta, British Columbia, Ontario and Québec net of mining (NAICS code 21) and utilities (NAICS code 22). I take the average across the years 1997-2020. I get a value of $\alpha = 0.4$. To compute Non-Energy TFP Z_s I compute the capital-labour composite

$K^\alpha L^{1-\alpha}$ using the calibrated value of α, data on the stock of fixed non-residential capital by industry (Table 36-10-0096-01, Statistics Canada) and hours worked by industry (Table 36-10-0489-01, Statistics Canada). I use the values for "Geometric end-year net stock" reported in current prices for the years 1996 to 2019 and deflate them by the annual GDP deflator for Canada taken from the Federal Reserve Economic Data (FRED). Both hours worked and capital stock once again correspond to the values for all industries net of mining and utilities, in the provinces of Alberta, British Columbia, Ontario and Québec.

The decreasing returns to scale parameters $\nu_{\mathcal{F}}$ and $\nu_{\mathcal{C}}$ are set using data on profits and revenues (Table 33-10-0500-01, Statistics Canada) for "Oil and gas extraction and support services" and "Utilities" respectively. I impute profits in each industry by taking the difference between "Sales of goods and services" and "Cost of goods sold." I then compute the profit-to-revenue ratio where profit corresponds to the imputed profits and revenue corresponds to "Sales of goods and services." The parameter ν_j for $j \in \{\mathcal{F}, \mathcal{C}\}$ is then given by

$$\nu_j = 1 - \frac{\pi_j}{Y_j}$$

where π_j is profits and Y_j is revenues. This returns a value of $\nu_{\mathcal{F}} = 0.73$ and $\nu_{\mathcal{C}} = 0.72$. The capital share in Fossil sector γ is computed using data on labour compensation (Table 36-10-0489-01, Statistics Canada) and value added (Table 36-10-0402-01, Statistics Canada) in Oil and Gas extraction (NAICS 211) and Support activities for oil and gas extraction (NAICS 21311A) in the province of Alberta in the years 1997 to 2020. Data for value added in Support activities in oil and gas is missing between 1997-2006. I impute the value of these years by computing

$$\kappa = \frac{1}{9} \sum_{t=1997}^{2006} \frac{Y^S_{\mathcal{F},t}}{Y^S_{\mathcal{M},t}}$$

where $Y^S_{\mathcal{F}}$ corresponds to value added in "Support activities for oil and gas extraction", and $Y^S_{\mathcal{M}}$ corresponds to "Support activities for mining, and oil and gas extraction". I then multiply this fraction by the value added in Support activities for mining, and oil and gas extraction between 1997 and 2006. For Alberta this is a good approximation, as Support activities in oil and gas as a fraction of all support activities in mining is fairly constant, accounting for approximately 96% value added between 2007 and 2020. I then construct a series for value added in the fossil sector as

$$\hat{Y}_{\mathcal{F},t} = \begin{cases} \kappa Y^S_{\mathcal{M},t} + Y_{\mathcal{F},t} & if \quad t \in \{1997, \dots, 2006\} \\ Y^S_{\mathcal{F},t} + Y_{\mathcal{F},t} & if \quad t \in \{2007, \dots, 2020\} \end{cases}$$

where $Y_{\mathcal{F},t}$ corresponds to the value added in Oil and Gas extraction (excluding support activities). Using the value $\nu_{\mathcal{F}}$, I use data on compensation (denoted $w_t L_{\mathcal{F},t}$) and value added (denoted $\hat{Y}_{\mathcal{F},t}$) compute

$$\gamma = 1 - \frac{1}{23}\sum_{t=1997}^{2020} \frac{w_t L_{\mathcal{F},t}}{\nu_{\mathcal{F}} \hat{Y}_{\mathcal{F},t}}$$

to arrive at $\gamma = 0.66$. Here, $w_t L_{\mathcal{F},t}$ corresponds to the sum of employee compensation in NAICS 211 and NAICS21311A. Similarly, the capital share in the clean sector is estimated using data on labour compensation and value added in Electric power generation, transmission and distribution (NAICS code 2211) in British Columbia, Ontario and Québec, taking the average across 1997-2020. This is a good approximation for clean energy production in Canada. Electricity production from hydro, nuclear, solar and wind account for 94% of total terawatt hours produced in these provinces, and 81% of total clean production across all of Canada. Using the same process as computing γ, I compute

$$\eta = 1 - \frac{1}{23}\sum_{t=1997}^{2020} \frac{w_t L_t}{\nu_{\mathcal{C}}(p_t Y_t)}$$

to arrive at $\eta = 0.60$

Table 1.3: Estimated Production Parameters

Parameter		Source	Value
α	Capital share, non-energy	StatsCan Tables	0.4
$Z_{F,1997}, Z_{C,1997}$	TFP in non-energy	StatsCan Tables	0.11
$Z_{F,2020}, Z_{C,2020}$	TFP in non-energy	StatsCan Tables	0.22
$\nu_{\mathcal{F}}$	DRS parameter	StatsCan Tables	0.73
$\nu_{\mathcal{C}}$	DRS parameter	StatsCan Tables	0.72
γ	Capital share, Fossil	StatsCan Tables	0.66
η	Capital share, Clean	StatsCan Tables	0.60
$\mu_{\mathcal{E}}$	Intermediate use of energy	I/O Table	38.43
$\mu_{\mathcal{F}}$	Intermediate use of Fossil	I/O Table	18.19
$p_{\mathcal{F},1997}$	Price of fossil good, 1997	WTI average	0.19
$p_{\mathcal{F},2020}$	Price of fossil good, 2020	WTI average	0.53

The Leontief intermediate parameters $\mu_{\mathcal{E}}$ and $\mu_{\mathcal{F}}$ are computed from the Canadian Input-Output tables by taking the ratio of value added over intermediate used. Given the Leontief structure of production, we can compute the parameter governing energy used in final good production, $\mu_E = Y/x_{\mathcal{E},Y}$ where Y maps to total value-added net of mining and utilities and $x_{\mathcal{E},Y}$ is energy intermediates used in this sector. "Energy" intermediates map to the sum of "Electric power generation, transmission and distribution" and "Natural gas distribution." I compute these ratios using the Input-Output tables for the years 2013-2020 and then take the average across all years to arrive at $\mu_{\mathcal{E}} = 38.43$. In the Fossil sector production function, $\mu_{\mathcal{F}} = Y_{\mathcal{F}}/x_{\mathcal{F},\mathcal{F}}$ where $Y_{\mathcal{F}}$ maps to the value added in the fossil sector and $x_{\mathcal{F},\mathcal{F}}$ is fossil intermediates used in this sector. The "fossil sector" corresponds to the sum of "Conventional oil and gas extraction", "Non-conventional oil extraction" and "Support activities for oil and gas extraction" in 2013, and "Oil and gas extraction (except oil sands)", "Oil sands extraction" and "Support activities for oil and gas extraction" for the years 2014-2020. I take the average across 2013-2020 and arrive at $\mu_{\mathcal{F},\mathcal{F}} = 18.19$.

The price of the fossil good $p_{\mathcal{F}}$ in the initial and final steady states corresponds to the annualized average of the West Texas Intermediate price. I convert the values to Canadian dollars using the FRED exchange rate and deflate them with the Canadian GDP deflator. They are then re-scaled to ensure the model produces an internal solution.

Next, I estimated the consumption tax rates $\tau_{c,s}$, the income tax progressivity parameter τ_p and the government share of final good consumption g. These values are reported in Table 1.4. The consumption tax rates in each region $\tau_{c,s}$ is computed using the method employed in Mendoza et al. (1994), Krueger and Ludwig (2016), and Moschini et al. (2024). I sum Taxes on products with Taxes on production/(Household final consumption expenditure + Non-profit institutions serving households' final consumption expenditure + General governments final consumption expenditure - Numerator) for Alberta to correspond to region $s = F$ and for the composite Canada that consists of British Columbia, Ontario and Québec to map into region $s = C$. Values are annualized by summing the quarterly values. I then take the average for the series. Data on taxes is taken from Statistics Canada Table 36-10-0221-01 while data on consumption is taken from Statistics Canada Table 36-10-0222-01.

Table 1.4: Estimated Government Parameters

Parameter		Source	Value
$\tau_{c,F}$	Consumption tax rate, region A	StatsCan tables	0.12
$\tau_{c,C}$	Consumption tax rate, region B	StatsCan tables	0.18
τ_p	Income Tax progressivity	StatsCan tables	0.1232
g	Government consumption	I/O Table	0.2

I estimate the income tax progressivity parameter τ_p by regressing log average pre-tax income on log average post-tax by income decile for the years 1976-2021. I use share of post-tax income as regression weights. So, if Y_{AT} is after tax income, we get

$$Y_{AT} = \lambda Y^{1-\tau_p}$$

Taking logs, we get

$$\log Y_{AT} = \log \lambda + (1 - \tau_p) \log Y$$

or

$$\log Y_{AT} = \beta_0 + \beta_1 \log Y$$

I run the regression separately for each year, compute $\tau_p = 1 - \beta_1$ for each year, and then take the average across all years. All data is taken from Statistics Canada Table 11-10-0193-01.

Finally, the government consumption parameter g is taken from quarterly expenditure side national accounts (Statistics Canada Table 36-10-0104-01). Values are annualized by summing the quarterly values. I compute G by summing Household final consumption expenditure, Non-profit institutions serving households' final consumption expenditure, General governments final consumption expenditure, Gross fixed capital formation, Investment in inventories, and Exports of goods and services then netting out imports of goods and services. For each year I compute $g = G/Y$ as General governments final consumption expenditure/Y where Y is GDP. I then take the average across years.

The age specific parameters governing productivity θ_j and survival probability ψ_j are summarized in Table 1.5. Age specific productivity θ_j is estimated using 4th degree polynomial on age in the following way: Using data on mean log income and mean log residual income for ages 25 to 55 from the Global Repository of Income Dynamics (GRID), I construct $y_j = \log Inc_j - \log Res\ Inc_j$. I run the regression

$$y_j = \beta_0 + \beta_1 year + \mu_j \tag{1.4.1}$$

Here, μ_j captures $\log \theta_j$. I save the estimated $\hat{\mu}_j$ from the previous regression, and I run the regression

$$\hat{\mu}_j = \delta_0 + \delta_1 age + \delta_2 age^2 + \delta_3 age^3 + \delta_4 age^4 + \zeta \tag{1.4.2}$$

I save the δ_i's from this regression and interpolate

$$\hat{\theta}_j = \exp\{\hat{\delta}_0 + \hat{\delta}_1 j + \hat{\delta}_2 j^2 + \hat{\delta}_3 j^3 + \hat{\delta}_4 j^4\} \tag{1.4.3}$$

for $j \in \{20, \ldots, j_{ret}\}$. I then normalize the values so that $\theta_1 = 1$. The age productivity profile can be seen in Figure 1.4a. It follows the expected hump shape over an individuals working life: earnings are lowest as an agent starts working and grow before gradually declining again in the later career years.

TABLE 1.5: Estimated Age specific Parameters

	Parameter	**Source**
θ_j	Age spec. productivity	GRID
ψ_j	Age spec. survival probabilities	StatsCan tables

Age specific survival probabilities ψ_j are taken directly from Statistics Canada Table 13-10-0114-01. Given that a period in my quantitative model corresponds to four years, then the probability of survival conditional on age j, ψ_j is computed as the product of 1 year survival probabilities between years t and $t+1$. For example, to compute the probability that an individual aged 20 (corresponding to $j = 1$) survives to the next period, $\psi_j = p(21|20) \times p(22|21) \times p(23|22) \times p(24|23)$ where $p(n+1|n)$ is the probability that an individual with age n survives to age $n+1$. Statistics Canada reports these estimates annually for sample cohorts lasting two years. The first sample was measured between 1980 and 1982, and the most recent sample was measured between 2020 and 2022. I compute the probabilities for each sample period and then take the average across periods. Given that agents in my model all die at $J = 100$, I fix the probability of survival at age 96 to $\psi_{J-1} = 0$

FIGURE 1.4: Age specific parameters

(A) Age Specific Productivity (B) Survival Probabilities

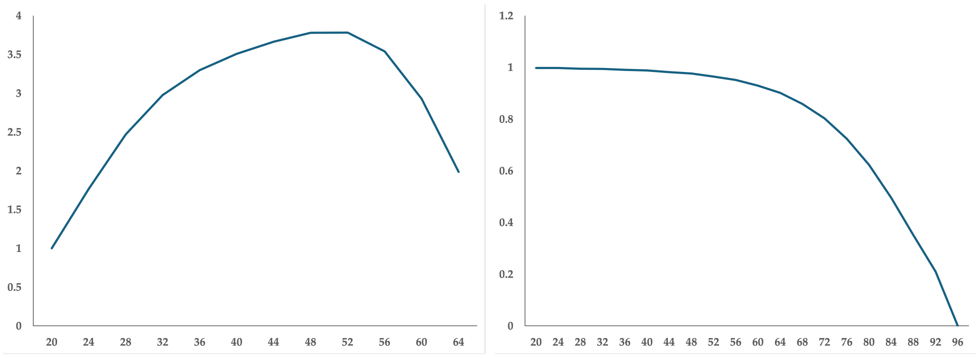

1.4.2 Internal Calibration

Moments related to the household productivity distribution are calibrated outside of equilibrium. These parameter values are presented in Table 1.6.

TABLE 1.6: Parameters determined outside of equilibrium

Parameter	Target	Source	Data	Model	Value
ρ_ϵ	ACF 1 period log res. earnings	GRID	0.74	0.74	0.75
σ^2_ϵ	SD 1 period change in log res. earnings	GRID	0.53	0.53	0.21

The data and target moments are taken from the GRID for Canada. The persistence of the AR(1) process ρ_ϵ is calibrated to match the autocorrelation function of the 1 period log residual earnings, and the variance σ^2_ϵ is calibrated to match the standard deviation of 1 period change in log residual earnings.

Table 1.7 presents parameter values calibrated in equilibrium. I match the productivity in the fossil sector $Z_{\mathcal{F}}$ in the initial steady state (1997) and the final steady state (2020) to match the percentage of capital used by the oil sector using data from Statistics Canada Table 36-10-0096-01. To match productivity in the clean sector $Z_{\mathcal{C}}$, I use data on primary energy consumption by source from Our World in Data (OWID). OWID sources the values from the Energy Institute's "Statistical Review of World Energy" and reports units of primary energy in terawatt hours broken down by source. To arrive at my target moments, I sum up terawatt hours of clean sources (hydro, solar, wind, other renewables and nuclear) and divide by the total number of terawatt hours across all sources.

TABLE 1.7: Parameters determined in equilibrium

Parameter	Target	Source	Data	Model	Value
$Z_{\mathcal{F}}^{1997}$	$K_{\mathcal{F}}/\bar{K}$ in 1997	StatsCan	0.10	0.10	1.66
$Z_{\mathcal{C}}^{1997}$	Share clean in 1997	OWID	0.37	0.37	0.32
$Z_{\mathcal{F}}^{2020}$	$K_{\mathcal{F}}/\bar{K}$ in 2020	StatsCan	0.22	0.22	1.46
$Z_{\mathcal{C}}^{2020}$	Share clean in 2020	OWID	0.36	0.36	0.29

1.5 Model Validation

In this section I present the model performance relative to what is observed in the data. A key feature of the data is that growth in wages, consumption and GDP in the fossil producing region (Region F), corresponding to the province of Alberta, outperformed the clean producing region (Region C), corresponding to a composite of B.C., Ontario, and Québec. Table 1.8 presents the results. The model is able to qualitatively replicate all three of these patterns that were not explicitly targeted in the calibration.

TABLE 1.8: Growth in Aggregates between Steady States

	Region F		**Region C**	
	Data	Model	Data	Model
Wage Growth	370.11	313.60	320.66	290.15
Consumption Growth	447.29	358.61	368.28	312.64
GDP Growth	431.81	384.16	371.53	290.63

The model reasonably captures the extent to which the economy of Region F benefited from the boom in fossil prices. Wages, consumption and GDP all grew by more between the two steady states in Region F than they did in Region C. To highlight the contribution specifically due to changes in $p_{\mathcal{F}}$, I evaluate the ratio of wages, consumption and GDP in Region F relative to Region C along the transition path in the benchmark transition and in a counterfactual world where $p_{\mathcal{F}}$ stays at the 1997 level. In Figure 1.5a, we can see that in the benchmark transition, the rise in $p_{\mathcal{F}}$ increases the relative wage between Regions F and C by approximately 12 percentage points. If the rise in $p_{\mathcal{F}}$ never happened, wages in Region C would have grown by more along the transition path than they did in Region F, decreasing the gap by about 8 percentage points.

Similarly, Figures 1.5b and 1.5c show the same inverse trend happening if the fossil boom never happens. In the benchmark transition, consumption in Region F relative to Region C grows by approximately 7 percentage points, while in the counterfactual transition relative consumption falls by about 4 percentage points. Relative GDP between the two regions grows by nearly 10 percentage points in the benchmark example, and in the absence of growth in $p_{\mathcal{F}}$, falls by 7 percentage points. Note that the fossil price is not the sole driver of growth in the two regions. The continued TFP growth in the non-energy sector produces wage growth in either scenario. However, the growth in $p_{\mathcal{F}}$ is the mechanism that drives the differences between the two regions. This result is

FIGURE 1.5: Counterfactual path of untargeted moments

(A) Relative Wages (B) Relative Consumption (C) Relative GDP

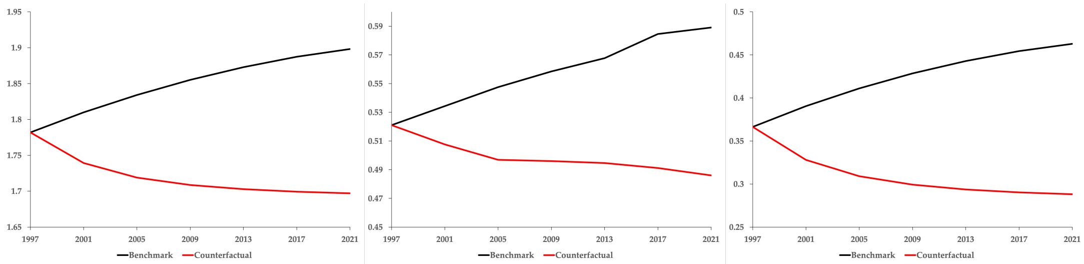

consistent with the literature on resource booms in other countries. For instance, Allcott and Keniston (2018) finds that US counties with oil and gas endowments experienced modestly higher real wages between 1969 and 2014.

1.5.1 Oil boom and welfare

This section evaluates the welfare impact of the rise in oil prices between 1997 and 2020 along the transition path. There are five objects in the model that change between 1997 and 2020 and then remain at a constant level until the economy reaches the new steady state: the fossil price ($p_{\mathcal{F}}$), and non-energy productivity (Z_F and Z_C) grow, while fossil productivity ($Z_{\mathcal{F}}$) and clean productivity ($Z_{\mathcal{C}}$) contract. All objects grow (or contract) linearly between the initial and final values over 24 years before remaining constant for another 400 years. The economy reaches the final steady state well before the final period, after approximately 60 years.

Welfare is reported as a weighted average of consumption equivalent variation (CEV). CEV is a measure of the constant percentage change in consumption where an individual is indifferent between two states of the world. That is, the consumption equivalent variation for an agent with a period t state vector (j, a, ϵ, s) would be

$$g = [\mathbb{E}_0 V_{bench}(t, j, a, \epsilon, s)/\mathbb{E}_0 V_{counter}(t, j, a, \epsilon, s)]^{1/(1-\sigma)} - 1$$

where V_{bench} corresponds to the value function of the benchmark transition and $V_{counter}$ is the value function from the counterfactual transition. A positive value for g implies the benchmark transition improves welfare, and a negative value implies it reduces welfare. I compute the expected value function for agents prior to the start of the transition in period $t = 2$. Conceptually, consider a household with current state vector (j, a, ϵ, s) that observes the paths of all key aggregates, but has not resolved the uncertainty about their idiosyncratic productivity ϵ. I then compare this agent's expected value function with the equivalent agent in a counterfactual world where the path of those aggregates is the same except for the price of the fossil good ($p_{\mathcal{F}}$), which is held constant. Since the two worlds are starting from the same initial steady state, the initial distribution over household types is the same, and the initial income distribution is the same.

Table 1.9 reports the weighted average CEV for households in specific age and income groupings between the two regions for this scenario. The rise in fossil prices benefits households in both regions. Lifetime consumption for households in Region F increases across all age and income groupings, with the youngest and poorest households increasing their consumption by about 16% along the transition path. The reason for this is that the wage growth over the course of their working lives in the benchmark transition is higher than in the counterfactual world, and profits from the fossil sector increase their income by more during the boom. Among these young, low income households, the gain in lifetime consumption is nearly 4 percentage points higher than the equivalent household in the clean region. For the wealthier households, the oil price boom benefitted lower income households by 0.12 percentage points more in the fossil region than in the clean region. These results are consistent with the literature on resource booms, which find that booms in resource prices benefit households at the bottom of the income distribution in the regions most exposed to the boom (for examples, see Jacobsen, 2019; Fortin and Lemieux, 2015).

TABLE 1.9: Welfare change due to p_F increase

			Region F		
	First Quintile	Second Quintile	Third Quintile	Fourth Quintile	Fifth Quintile
20-34	15.76%	12.01%	10.74%	10.21%	9.98%
35-49	15.63%	12.14%	10.27%	9.46%	9.01%
50-64	12.73%	9.88%	8.39%	7.94%	7.87%
65-99	NA	7.75%	6.79%	6.46%	6.63%
			Region C		
20-34	11.06%	9.65%	9.39%	9.77%	9.86%
35-49	10.06%	8.80%	8.83%	8.86%	9.00%
50-64	6.73%	6.82%	7.08%	7.37%	7.77%
65-99	4.61%	5.67%	6.18%	6.55%	6.86%

These welfare outcomes are largely driven by wages, profits and the interest rate. Region F wages are higher in the benchmark world, while Region C wages are higher in the counterfactual world. This boost in Region F's wage growth is a contributing factor to the larger welfare gains that the lowest income households experience. Bequests and social security transfers are nearly identical in the two scenarios. The return on savings also grow significantly more in the benchmark transition than in the counterfactual. Similarly, profits in Region F are lower in the counterfactual world. These trends are presented in Figure 1.6

Figure 1.6a highlights how much households in Region F gain in the from the rise in $p_{\mathcal{F}}$. The benchmark transition results in massive growth in profits from the fossil sector, which are redistributed uniformly across households in the region, raising lifetime earnings for all households. These transfers are particularly beneficial to the lower income households

FIGURE 1.6: Profits and Interest Rate

(A) Fossil Profits (B) Interest Rate

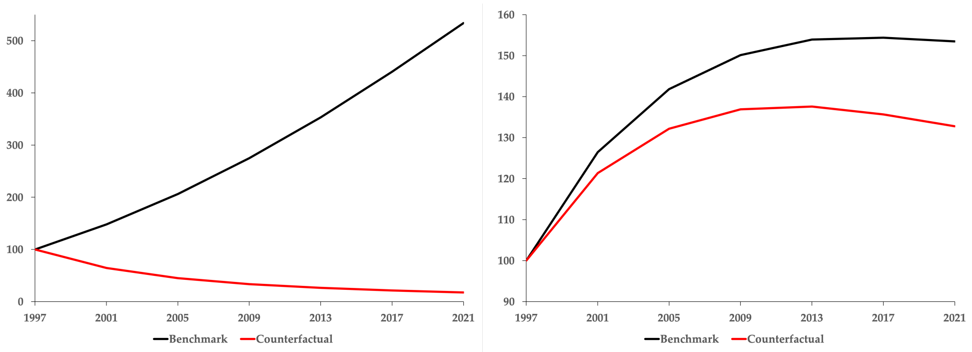

as they act as an additional supplement on top of the higher lifetime labour earnings. In the counterfactual experiment, profits in the fossil sector are actually decreasing. Since $p_{\mathcal{F}}$ is constant while both w_F and r are still increasing over time, the sectors costs relative to revenue increase, even as they demand less capital and labour. Profits from the clean sector actually grow in both the benchmark and the counterfactual exercises. However, in either scenario these profits account for less than 1% of a households income in Region C across the transition. The interest rate (Figure 1.6b) grows at a faster rate in the benchmark than it does in the counterfactual. Along the remainder of the transition path, as the economy converges to the new steady state, the interest rate is roughly 9 percentage points higher than it would have been if $p_{\mathcal{F}}$ had remained at the 1997 level. This is particularly beneficial to the youngest households in the economy as they earn significantly higher returns on their savings over their lifetime. These households labour earnings are at their lowest at the start of the transition, given the path of the age productivity profile: at age 20, their age specific component of labour earnings is at the lowest level. During these early years, their incomes are largely supplemented by the profits. As their labour earnings grow with their age, the profit transfers from the fossil sector account for a decreasing but important share of their total income. During retirement years, the percentage of household income derived from profits increases once social security replaces labour income.

1.6 Green Transition

This section evaluates the welfare impact of the Green Transition. For the purpose of this paper, I refer to the "Green Transition" as the transition path to a Net Zero world. This is difficult to define precisely, as there is no universal agreement on what the path to Net Zero looks like or how it ought to be achieved. For this reason, I adopt a simplified version of the International Energy Agency's recommendations in their "Net Zero by 2050" document (IEA, 2021). To reach Net Zero by 2050, the IEA estimates that fossil fuel demand needs to fall by half, and clean energy production needs to account for 90%

of electricity production. Through the lens of the model, this occurs via two mechanisms: a fall in $p_{\mathcal{F}}$ and an increase in $Z_{\mathcal{C}}$. This path serves as a useful, though conservative, benchmark. Larger reduction targets for fossil demand serve to quantitatively amplify the welfare changes presented here, but do not qualitatively change the results.

In the benchmark experiment, I start from 2020 as an initial steady state. I recalibrate a subset of parameters so that in the final steady state, fossil production $Y_{\mathcal{F}}$ falls to 50% of the initial level, and clean intermediates used in domestic energy production, $x_{\mathcal{C},\mathcal{E}}$ account for 90% of the intermediates used (that is $x_{\mathcal{C},\mathcal{E}}/(x_{\mathcal{C},\mathcal{E}} + x_{\mathcal{F},\mathcal{E}}) = 0.9$). I assume that $Z_{\mathcal{F}}$ remains constant at the 2020 level throughout the transition. I also assume that TFP in the non-energy sectors continues to grow along the same trend line as what is observed in the data from 2010 to 2020. In the transition exercise, I construct a linear path for $p_{\mathcal{F}}$, $Z_{\mathcal{C}}$ and Z_F, Z_C and solve the model along these trends. Table 1.10 presents the calibration results.

Table 1.10: Benchmark Green Transition calibration

Parameter	Target	Data	Model	Value
$p_{\mathcal{F}}^{2050}$	$0.5 \times Y_{\mathcal{F}}^{2020}$	0.61	0.61	0.49
$Z_{\mathcal{C}}^{2050}$	90% of energy intermediates are clean	0.9	0.9	1.29

Using these calibrated values, I establish my benchmark “Green Transition” (GT) scenario. Fossil demand $p_{\mathcal{F}}$ decreases from 2020 to 2050, and then remains constant, while productivity in the clean sector ($Z_{\mathcal{C}}$) grows over the same period and then remains constant. Non-energy productivity (Z_F, Z_C) grow over the same period, but by less than $Z_{\mathcal{C}}$.

1.6.1 Decrease in demand for fossil fuels

I begin by evaluating the welfare impacts strictly from a decrease in fossil fuel demand with no concurrent rapid expansion in clean energy production relative to a “No Green Transition” (NGT) scenario. The NGT transition is a world where (a) fossil demand remains at the 2020 levels (so $p_{\mathcal{F}}$ is constant) and (b) $Z_{\mathcal{C}}$, Z_F and Z_C all grow at the same rate, meaning $Z_{\mathcal{C}}$ grows slower than it does in the GT benchmark. To isolate the costs from a fall in fossil demand, I compare the NGT scenario to an identical world, except that $p_{\mathcal{F}}$ declines until 2050 along the GT path.

The decline in demand for fossil fuels produces aggregate welfare losses of 0.55% across the entire economy. The detailed breakdown is presented in Table 1.11. The distribution of losses is largely the mirror image of the outcomes from the backward-looking exercise in Section 1.5. The largest decreases in lifetime consumption are felt by households in Region F, who are the most exposed to the contraction of the fossil industry. Young, low-income households’ consumption decreases by 0.77% along the transition.

The decrease in $p_{\mathcal{F}}$ implies a fall in overall fossil production $Y_{\mathcal{F}}$. Fossil production in the NGT scenario also contracts due to movement in other aggregates. The growth in both the clean and non-energy sectors leads to rising wages and interest rate, raising production costs across the economy. As both productivity $Z_{\mathcal{F}}$ and the price $p_{\mathcal{F}}$ stay constant, this results in lower demand for both labour and capital in the sector, and lower output overall. However, the contraction of the fossil sector is more pronounced with a fall in $p_{\mathcal{F}}$.

While productivity growth in clean energy $Z_{\mathcal{C}}$ lowers the relative price between the fossil good and the clean good in both scenarios, the price does not fall enough to match the decline in $p_{\mathcal{F}}$, so fossil fuels account for an increasing share of energy inputs in this scenario. However, the combination of falling $p_{\mathcal{F}}$ alongside increasing wages w_F and interest rate r also mean profits fall. Relative to the NGT scenario, fossil profits fall by an average of 17% more over the lifetime of an individual aged 20 at the beginning of the transition, significantly reducing lifetime income. The difference in wage growth in Region F is negligible between the two scenarios, only about 0.05% which does not serve to offset the fall in profits. The interest rate is also on average 1.19% lower than it would have been if $p_{\mathcal{F}}$ stayed constant, so the return to lifetime savings is dampened as well.

Wages in Region C are about 0.6% higher on average during the working life of a 20-year-old, while profits from the clean sector are lower by about 7%. This is due to the combination of increased production costs but also due to the substitution towards using more fossil fuels in domestic energy production. In both regions, profits are lower than in the NGT scenario. However, this fall in profits is 10 percentage points more pronounced in Region F. The contraction in fossil production and increase in domestic fossil use also implies a decline in exports. Due to trade being balanced in this economy, this also implies a reduction in imports. For an initial 20-year-old, imports are on average 20% lower, while aggregate production of the final good only increases by about 1%. As a result, there is less of the consumption good available overall.

TABLE 1.11: Impact of $p_{\mathcal{F}}$

	Region F				
	First Quintile	Second Quintile	Third Quintile	Fourth Quintile	Fifth Quintile
20-34	-0.77%	-0.59%	-0.52%	-0.49%	-0.47%
35-49	-0.73%	-0.58%	-0.50%	-0.46%	-0.44%
50-64	-0.61%	-0.51%	-0.44%	-0.41%	-0.39%
65-99	NA	-0.47%	-0.38%	-0.35%	-0.34%
	Region C				
20-34	-0.49%	-0.42%	-0.40%	-0.40%	-0.41%
35-49	-0.45%	-0.40%	-0.38%	-0.38%	-0.38%
50-64	-0.35%	-0.33%	-0.33%	-0.33%	-0.34%
65-99	-0.26%	-0.29%	-0.30%	-0.30%	-0.31%

1.6.2 Clean energy expansion

To assess the overall benefits to the economy from faster growth in the clean energy sector, I compare the welfare impact of the benchmark Green Transition (GT) (relative to NGT) with the impact due solely to falling $p_{\mathcal{F}}$ from the previous section. Table 1.12 summarizes the welfare results by age and income groups across the two regions.

Table 1.12: Impact of the Green Transition

	Region F				
	First Quintile	Second Quintile	Third Quintile	Fourth Quintile	Fifth Quintile
20-34	-0.73%	-0.51%	-0.42%	-0.37%	-0.34%
35-49	-0.68%	-0.50%	-0.41%	-0.35%	-0.32%
50-64	-0.57%	-0.45%	-0.37%	-0.32%	-0.29%
65-99	NA	-0.43%	-0.31%	-0.27%	-0.25%
	Region C				
20-34	-0.38%	-0.28%	-0.24%	-0.23%	-0.23%
35-49	-0.34%	-0.26%	-0.24%	-0.22%	-0.22%
50-64	-0.25%	-0.22%	-0.21%	-0.20%	-0.19%
65-99	-0.16%	-0.21%	-0.20%	-0.19%	-0.19%

To isolate the effect of expansion in clean energy production (i.e. growth in $Z_{\mathcal{C}}$), I take the ratio of the CEV's. Then the impact on overall welfare from growth in the clean sector is given by

$$1 - \frac{CEV_{GT}}{CEV_{p_{\mathcal{F}}}}$$

where the numerator is the measure of consumption equivalence from the benchmark Green Transition (relative to NGT) and the denominator is the measure of consumption equivalence from the fall in fossil demand from the previous section. The domestic growth in clean energy production produces modest welfare improvements, dampening the aggregate losses by 15% in aggregate. This expansion is not sufficient to completely negate the welfare losses from falling demand for fossil fuels, but it does produce some substantial improvement for some households. The youngest, poorest households in Region F benefit the least and only see an improvement in lifetime consumption of 5.68%. Moving up the income distribution, the wealthiest households in the region see their overall welfare improve by between 27 and 28% thanks to the expansion in the clean sector. In Region C, the bottom of the age and income distribution see their welfare improve by roughly 22%. The losses for the wealthiest households decrease between 38% and 44%.

Fossil production contracts by more in this scenario than in the previous section, due to the falling domestic demand for fossil fuels. The relative price between clean and fossil inputs falls by enough so that energy production substitutes away from fossil fuels and increases the use of clean inputs. This also means the price of the energy good is significantly lower during the Green Transition. On average, over the course of the transition the energy price in the GT scenario is 38% lower than it was in the previous exercise. This results in final good production expanding by slightly more (about 0.31% higher per year). Expanding production of the final good produces an increase in demand for non-energy inputs from both regions, which can account for the stronger wage growth that the Green Transition produces. A 20-year-old in Region F will have labour income that is on average 0.61% higher over the course of the Green Transition than they would have if there was no concurrent expansion in the clean sector. In Region C, wages are 0.66% higher over this period.

Due to the fall in domestic demand for fossil energy, exports as a fraction of output in the fossil sector increase. While the fossil sector contracts slightly more in the Green Transition than it would absent the expansion of the clean sector, and even though in both scenarios exports decline over time, the economy still sells slightly more of the fossil good to the rest of the world in the GT scenario than in the previous section. As a result, due to the balanced trade condition the economy imports slightly more of the final consumption and investment good. Since the aggregate economy can both produce more of the final consumption good and import more of it, households in both regions are able to consume and invest slightly more thanks to the expansion of the clean sector.

The interest rate is on average 0.15% lower over the lifetime of an initial 20-year-old than the interest rate in the previous section was, but aggregate savings are higher, due to the increase in the total quantity of the final good available for consumption and investment. As a result, the annual income from savings over the lifetime of an initial 20-year-old is 0.68% higher thanks to the expansion of the clean sector. This increase in savings income and labour income in both regions work to dampen the welfare losses from falling fossil demand, particularly at the top of the income distribution among young households. These households benefit from higher labour income during their working lives, and can save more, thus reaping additional gains from the higher savings.

This increase in labour income and return to savings is higher than the decrease in profits from the fossil sector. Over the lifetime of the youngest households, profits from the fossil sector are 0.32% lower than in the previous section. This decline is less than the growth in labour income and savings, which explains why households in Region F gain. In Region C, profits from the clean sector are 24% higher, which helps to explain the significant dampening of the losses from the previous section.

1.6.3 Alternative Scenario: Faster growth in clean productivity

I now explore how much welfare improves if the domestic production of clean energy happens more rapidly. I compare the "No Green Transition" world to a new benchmark where $Z_{\mathcal{C}}$ arrives immediately at the final steady state level in the first period of the transition. This is a world where the demand for the fossil good falls between 2020 and 2050 so that the economy produces 50% of 2020 levels, non-energy productivity grows along the 2010-2020 trend, and clean technology is instantaneously productive enough to account for 90% of energy inputs. I then treat this transition as the alternative benchmark and derive the welfare results relative to the “No Green Transition” scenario from the previous section. For the entire economy, this immediate expansion of the clean energy sector reduces the losses from falling fossil fuel demand by 54%. Table 1.13 summarizes the welfare outcomes from the alternative benchmark relative to the “No Green Transition” scenario.

TABLE 1.13: Alternative Benchmark: Instantaneous growth in $Z_{\mathcal{C}}$

	First Quintile	Second Quintile	Third Quintile	Fourth Quintile	Fifth Quintile
			Region F		
20-34	-0.62%	-0.23%	-0.06%	0.05%	0.10%
35-49	-0.59%	-0.28%	-0.11%	0.00%	0.08%
50-64	-0.52%	-0.32%	-0.16%	-0.06%	0.02%
65-99	NA	-0.35%	-0.16%	-0.07%	-0.01%
			Region C		
20-34	-0.08%	0.17%	0.27%	0.32%	0.35%
35-49	-0.04%	0.13%	0.21%	0.26%	0.30%
50-64	-0.05%	-0.06%	0.13%	0.19%	0.23%
65-99	-0.04%	-0.02%	0.05%	0.08%	0.11%

An instantaneous growth in $Z_{\mathcal{C}}$ dampens the welfare losses by significantly more than in the previous exercise, particularly among households in Region C. Faster growth in the clean sector makes the path nearly welfare neutral among the youngest and oldest households in Region C, and welfare enhancing for middle- and late-career households.

For households in Region F, the dampening in the costs of the fall in oil demand is largely driven by the expansion of both final good production and imports and the growth in wages. Production of the fossil good is on average 0.22% lower than it was in the benchmark Green Transition. Due to the accelerated substitution towards clean energy inputs, from 2020 to 2052 fossil exports account for a slightly higher share of total fossil production (between 1 and 3 percentage points) before converging to approximately the same levels. As a result, imports are slightly higher than in the benchmark GT. Final good production is also on average 0.13 percentage points higher than in the benchmark transition, for the same reasons as discussed in the previous section: lower energy costs due to the extreme growth in $Z_{\mathcal{C}}$ allows for an expansion in final good production. This generates an increase in the demand for non-energy inputs in both regions; demand for Region F's non-energy intermediate is 0.18% higher, and 0.13% higher for Region C's non-energy intermediate. For the working life of the 20-year-olds at the start of the transition, wages in Region F are 0.27% higher than they are in in the benchmark, and 0.30% higher than the benchmark in Region C. The profits from the fossil sector are 0.22% lower than they are in the benchmark case, and the interest rate is 0.06% lower.

In Region C, profits from the clean sector increase during the transition but are about 7% lower on average than they are in the GT benchmark. This is because the price $p_{\mathcal{C}}$ falls much more rapidly and labour costs increase by more. Across the economy, the return on savings is slightly lower than the benchmark but the actual capital stock is about 0.20% higher. Between 2020 and 2028 the total return on savings in the economy is slightly higher in the alternative benchmark (with faster growth in Z_C), which largely benefits the households at the upper end of the income distribution. In Region F, we can see this in the last column Table 1.13 where the younger, wealthier households gain from the Green Transition. The youngest households in Region F benefit from the higher lifetime labour earnings, and the slightly higher return on savings early in the transition. Similarly in Region C, the gains from the transition are most pronounced among the youngest and wealthiest households for the same reasons. Their total savings are higher at the start of the transition, and they benefit from the combination of those initial higher returns to savings, growing profits and higher wages during their working years.

1.7 Concluding Remarks

This paper evaluates the welfare impact of a Green Transition to a Net Zero world in small, resource-rich economies. While the fossil boom in the 2000's produced large welfare gains across the age and income distribution, the model predicts decreases in lifetime consumption of 0.56% along the transition path in a world with declining demand for fossil fuels. Using the IEA's recommendations for reaching a Net Zero world, I find that these losses are most pronounced among the youngest and poorest households residing in the fossil producing region of the economy, for whom lifetime consumption decreases by 0.77%. This is largely driven by falling profits from the fossil sector, slower wage growth and lower returns to savings.

These losses are dampened by the expansion of the clean energy sector. In the benchmark Green Transition scenario, the expansion of clean energy to account for 90% of domestic energy inputs by 2050 reduces the aggregate welfare losses from falling demand for fossil fuels by 15%. Accelerating this expansion of the clean energy sector to has an even stronger impact, reducing the losses by 54%. While the decrease in consumption due to a fall in fossil fuel demand are largest among young, low-income households, the expansion of the clean sector dampens their losses 19% and 85% in the fossil producing and clean producing regions respectively. The largest beneficiaries of an accelerated expansion of the clean energy sector are the wealthiest young households.

This paper leaves some of the finer details of how local labour markets respond to these changes to future work. The model abstracts from potential challenges in switching jobs that workers in the fossil producing region experience, or for possible changes to inter-regional migration patterns that may arise. Using microdata on employment dynamics within the fossil fuel industry to inform a similar model that includes job search would be particularly insightful.

Extensions of this paper may also shed light on the potential costs stemming from capital adjustment costs and stranded assets. This paper assumes capital is perfectly mobile across regions and between industries. There is a rich literature on the economics of climate change and the potential for capital investments to become stranded and worthless. For a summary of the relevant literature, see von Dulong et al. (2023). The results in this model are partially driven by changes in the interest rate and profits. However, this occurs within a static production structure. Adjusting the model to incorporate dynamic, industry specific investment decisions will be beneficial.

This paper also highlights the channels through which the costs of the Green Transition can be negated, namely via growth in other sectors of the economy. Extending the model to allow for the emergence of a clean energy sector in the fossil producing region would likely dampen some of the negative impacts in the fossil region. Future work focusing on policy changes that boost productivity to offset the costs of the transition is needed. Alternatively, further work studying redistributive policies to compensate the low-income households who are adversely impacted may yield further insights.

References

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). 'Transition to Clean Technology'. Journal of Political Economy, volume 124, no. 1, 52–104.

Aiyagari, S. R. (1994). 'Uninsured Idiosyncratic Risk and Aggregate Saving'. The Quarterly Journal of Economics, volume 109, no. 3, 659–684.

Albrecht, L. and Tombe, T. (2016). 'Internal trade, productivity and interconnected industries: A quantitative analysis'. Canadian Journal of Economics, volume 49, no. 1, 237–263.

Allcott, H. and Keniston, D. (2018). 'Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America'. The Review of Economic Studies, volume 85, no. 2, 695–731.

Arkolakis, C. and Walsh, C. (2023). 'Clean Growth'.

Baldwin, E., Cai, Y., and Kuralbayeva, K. (2020). 'To build or not to build? Capital stocks and climate policy'. Journal of Environmental Economics and Management, volume 100, 102235.

Bartik, A. W., Currie, J., Greenstone, M., and Knittel, C. R. (2019). 'The Local Economic and Welfare Consequences of Hydraulic Fracturing'. American Economic Journal: Applied Economics, volume 11, no. 4, 105–155.

Besley, T. and Persson, T. (2023). 'The Political Economics of Green Transitions'. The Quarterly Journal of Economics, volume 138, no. 3, 1863–1906.

Bewley, T. (1977). 'The permanent income hypothesis: A theoretical formulation'. Journal of Economic Theory, volume 16, no. 2, 252–292.

Black, D., McKinnish, T., and Sanders, S. (2005). 'The Economic Impact Of The Coal Boom And Bust*'. The Economic Journal, volume 115, no. 503, 449–476.

Borenstein, S. and Davis, L. W. (2016). 'The Distributional Effects of US Clean Energy Tax Credits'. Tax Policy and the Economy, volume 30, no. 1, 191–234.

Casey, G. (2023). 'Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation'. The Review of Economic Studies, pages 192–228.

Fortin, N. M. and Lemieux, T. (2015). 'Changes in wage inequality in Canada: An interprovincial perspective'. Canadian Journal of Economics, volume 48, no. 2, 682–713.

Fried, S., Novan, K., and Peterman, W. B. (2022). 'Climate policy transition risk and the macroeconomy'. European Economic Review, volume 147, 104174.

Hassler, J., Krusell, P., and Olovsson, C. (2021). 'Directed Technical Change as a Response to Natural Resource Scarcity'. Journal of Political Economy, volume 129, no. 11, 3039–3072.

Hassler, J., Krusell, P., and Olovsson, C. (2022). 'Finite resources and the world economy'. Journal of International Economics, volume 136, 103592.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). 'Optimal Tax Progressivity: An Analytical Framework*'. The Quarterly Journal of Economics, volume 132, no. 4, 1693–1754.

Helm, C. and Mier, M. (2021). 'Steering the energy transition in a world of intermittent electricity supply: Optimal subsidies and taxes for renewables and storage'. Journal of Environmental Economics and Management, volume 109, 102497.

Huggett, M. (1993). 'The risk-free rate in heterogeneous-agent incomplete-insurance economies'. Journal of Economic Dynamics and Control, volume 17, no. 5, 953–969.

IEA (2021). 'Net Zero by 2050'. Technical report, IEA, Paris.

Imrohoroğlu, A. (1989). 'Cost of Business Cycles with Indivisibilities and Liquidity Constraints'. Journal of Political Economy, volume 97, no. 6, 1364–1383.

Jacobsen, G. D. (2019). 'Who Wins in an Energy Boom? Evidence from Wage Rates and Housing'. Economic Inquiry, volume 57, no. 1, 9–32.

Jaimes, R. (2023). 'Optimal climate and fiscal policy in an OLG economy'. Journal of Public Economic Theory, volume 25, no. 4, 727–752.

Krueger, D. and Ludwig, A. (2016). 'On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium'. Journal of Monetary Economics, volume 77, 72–98.

Laurent-Lucchetti, J. and Leach, A. J. (2011). 'Generational Welfare under a Climate-Change Policy with Induced Innovation*'. The Scandinavian Journal of Economics, volume 113, no. 4, 904–936.

Lennox, J. A. and Witajewski-Baltvilks, J. (2017). 'Directed technical change with capitalembodied technologies: Implications for climate policy'. Energy Economics, volume 67, 400–409.

Loertscher, O. and Pujolas, P. S. (2024). 'Canadian productivity growth: Stuck in the oil sands'. Canadian Journal of Economics, volume 57, no. 2, 478–501.

Mendoza, E. G., Razin, A., and Tesar, L. L. (1994). 'Effective tax rates in macroeconomics: Cross-country estimates of tax rates on factor incomes and consumption'. Journal of Monetary Economics, volume 34, no. 3, 297–323.

Michaels, G. (2011). 'The Long Term Consequences of Resource-Based Specialisation'. The Economic Journal, volume 121, no. 551, 31–57.

Moschini, E. G., Raveendranathan, G., and Xu, M. (2024). 'Optimism About Graduation and College Financial Aid'.

Papageorgiou, C., Saam, M., and Schulte, P. (2017). 'Substitution between Clean and Dirty Energy Inputs: A Macroeconomic Perspective'. The Review of Economics and Statistics, volume 99, no. 2, 281–290.

von Dulong, A., Gard-Murray, A., Hagen, A., Jaakkola, N., and Sen, S. (2023). 'Stranded Assets: Research Gaps and Implications for Climate Policy'. Review of Environmental Economics and Policy, volume 17, no. 1, 161–169.

Chapter 2

Canadian Productivity Growth: Stuck in the Oil Sands

Canadian Productivity Growth: Stuck in the Oil Sands

Abstract

We study the behaviour of Canadian total factor productivity growth over the past 60 years. We find that the observed stagnation during the last 20 years is entirely accounted for by the Oil sector. Higher oil prices made capital-intensive sources of oil like the oil sands viable to extract on a commercial scale. However, the greater input required per barrel of oil slowed productivity growth. Comparing Canadian TFP growth to that of the United States reinforces these results. However, our result should not be interpreted to carry any welfare implications.

2.1 Introduction

Economists and policy-makers have expressed ongoing concern about the lack of productivity growth in Canada.[1] Conesa and Pujolas (2019) identifies the "Canadian Productivity Stagnation" during the period from 2002 to 2014, where Canadian Total Factor Productivity (TFP) growth was negligible, lagging behind both previous decades in Canada and contemporaneous US TFP growth. Our paper extends the period of Canadian Productivity Stagnation up to at least 2018 (the last year for which data is available, as detailed in Appendix A), and demonstrates that the absence of aggregate TFP growth during this period, and the divergence of TFP growth with respect to the United States, can be attributed entirely to the Oil sector.[2] Increases in oil prices made capital-intensive, lower TFP oil production (such as the oil sands) viable, which in turn reduced aggregate Canadian TFP growth.

In essence, the lack of TFP growth is entirely accounted for by excluding the Oil sector from the Canadian national accounts and recalculating TFP accordingly (we refer to this measure as "Net-of-Oil TFP"). From 2001 to 2018, Canada's TFP grows at 0.06 % per year. By contrast, Canada's Net-of-Oil TFP grows at 0.60 % per year, a similar rate to that of the United States (0.47 % if oil is included, 0.49 % if oil is excluded).[3] It is

[1]See, for instance, the Fraser Institute's monograph on Improving Productivity Growth in Canada (Douglas et al., 2021); Op-Ed by William Robson, CEO of the C.D. Howe Institute, in the Financial Post titled "Faster Productivity Growth Would Solve Many Problems" (Robson, 2022); OECD's Canada Economic Snapshot (OECD, 2023); Deloitte's Future of Productivity volume (Currie, Scott, and Dunn, 2021); or the Charter of Professional Accountants in Canada's "Solution to Canada's plummeting productivity" (Fong, 2019) to name a few.

[2]Throughout the paper we refer to the industry named "Oil and Gas Extraction," NAICS code 211, as the "Oil sector." Similarly, we refer to industry "Non-conventional Oil and Gas extraction," NAICS code 211114, as the "Oil sands." Note that this sector was relabelled in the 2017 revision as "Oil sands extraction."

[3]As shown in Fernald (2015), Cette et al. (2016), there is a decrease in the TFP growth of many rich economies, including Canada and the United States. To be precise, our work demonstrates that the Oil sector can account for the different evolution of TFP between Canada and the United States, but not the overall productivity decrease.

important to note that TFP does not account for the positive effects of rising oil prices.

To measure TFP, we need data on output (real GDP), inputs (a measure of the capital stock in the economy, as well as total hours worked), and an assumption on how inputs combine to generate output (which we employ through a Cobb-Douglas production function). Then, TFP is calculated as changes in output that cannot be attributed to changes in inputs, as first proposed by Solow (1957). Our methodology is based on the approach in Kehoe and Prescott (2007) and is similar to that in Conesa, Kehoe, and Ruhl (2007).

We continue our analysis by looking at different Canadian provinces. Namely, we compare Alberta (a large, oil-producing province) to Ontario (a large, non-oil-producing province).[4] First, we find that Ontario's TFP evolves almost identically regardless of whether oil is included in the calculations (over the period of analysis, TFP and Net-of-Oil TFP grow at 0.82 % per year). Second, we find that Alberta's Net-of-Oil TFP growth is higher than Ontario's (0.97 % per year). However, when oil is included, Alberta's TFP grows at a negative rate during this time period, at -0.20 %.

Next, we compare Canadian TFP growth to that of the United States and Norway. We use US data from the Bureau of Economic Analysis and Norwegian data from the Statistisk sentralbyrå. Canada, the United States, and Norway are the three largest oil producers among developed, western economies. US TFP growth is not significantly affected by excluding the Oil sector during the period from 2001 to 2018. We find that the key difference between Canadian and US TFP growth lies in the differential evolution of capital and output by the Oil sector. While the Canadian Oil sector's share of installed capital almost doubles during the period of analysis (from using 17.1 % of all the installed capital in 2001 to 31.4 % in 2018), the Oil sector's share of Canadian GDP remains fairly constant, and even decreases (5.7 % in 2001 to 4.7 % in 2018). By contrast, the US Oil sector's capital and GDP stay relatively constant and at lower rates throughout (from 2.3 % in 2001 to 3.6 % in 2018 for capital, and from 0.8 % in 2001 to 1.1 % in 2018 for value added).

Norwegian TFP growth is affected by excluding the Oil sector during the period from 2001 to 2018, aligning with Canada. The difference between Net-of-Oil TFP and TFP growth rates is 0.37 % in Norway, and 0.54 % in Canada (both growth rates are small in Norway, at -0.32 % for TFP and 0.05 % for Net-of-Oil TFP). Norwegian results are also driven by the evolution of capital and output by the Oil sector. In 2001, 13.9 % of capital in Norway is in the Oil sector (17.1 % in Canada), increasing to 14.2 % by 2018 (31.4 % in Canada). At the same time, Norway's Oil sector accounts for 22.4 % of its GDP in 2001 (5.7 % in Canada), and decreases to 19.0 % in 2018 (4.7 % in Canada).

The surge in capital used by the Canadian Oil sector coincides with the early 2000s oil price boom and the commencement of commercial oil sands extraction in 2001.[5] It is

[4] Due to data availability reasons, our provincial-level analysis spans from 1997 to 2018.

[5] The offshore oil project Hibernia in Newfoundland and Labrador, which started producing in the late 1990's, also fits this analysis. However, its production is much lower than the oil sands.

noteworthy that the proportion of capital in the oil sands as a percentage of the overall capital in the Oil sector has seen a remarkable increase, rising from an average of 5.77 % between 1961 and 2000 to an astonishing 30.35 % as of 2019.

Lastly, it is important not to interpret our results as a critique of the Oil sector. While the industry confronts a host of challenges, such as carbon emissions, our findings do not necessarily imply any additional negative aspects. Rather, the drop in TFP may be attributed to a combination of increased oil prices and a technology that exhibits decreasing returns to scale. Higher oil prices might encourage the extraction of costlier barrels of oil, which would lead to a lower TFP due to a composition effect. Hence, it is plausible that the Canadian economy is responding optimally by exploiting a resource when its value is high.

The paper is organized as follows: Section 2.2 contextualizes our contribution in light of the literature. Section 2.3 presents the methodology used to measure TFP and highlights that TFP in Canada has remained stagnant since the early 2000s, but increases when the Oil sector is excluded from the calculations. In Section 2.4, we perform a comparative analysis of TFP between Canada, the United States, and Norway. Section 2.5 offers additional information regarding the Canadian Oil sector that is pertinent to understanding the productivity of the sector. Finally, in Section 3.4, we provide further context and propose potential avenues for future research.

2.2 Literature Review

The lack of productivity growth in Canada post-2000 is an ongoing topic of discussion among scholars and policy-makers. This paper builds on Conesa and Pujolas (2019) and other studies that have investigated the sluggish productivity in Canada. For example, Boothe and Roy (2008) reviews labour and multi-factor productivity (MFP) in the Canadian business sector and links weak MFP growth to Canadian firms' lackluster innovation performance. They also note a sharp decline in productivity in the oil and gas sector from 2000 to 2006. Similarly, Alexopoulos and Cohen (2018) finds that the slowdown in productivity growth in the Canadian business sector since 2000 was due to a decrease in the rate at which Canadian firms adopt new technologies and a lack of innovative activity. Our paper complements these studies by focusing on the connection between TFP and the Oil sector, rather than firm level innovative activity.

Shao and Tang (2021) examines the role of allocative efficiency in driving aggregate labour productivity growth and explores the reasons behind the labour productivity gap between Canada and the United States. The paper identifies capital allocation as the primary factor responsible for the decline in Canadian allocative efficiency. Similarly, Gu (2018) investigates the impact of various measurements of capital on slow productivity growth in Canada, and finds that a quarter of Canada's productivity slowdown between 2000 and 2015 is due to the use of capital in the Oil and gas sector. Our analysis in

Section 2.5 is consistent with these findings, as we demonstrate that Canada's capital has been heavily utilized by the Oil sector since the late 1990s and that this trend can explain the lack of TFP growth during the same period.

Sharpe (2010) focuses on 12 different industries in Canada and argues that the decline in labour productivity growth in the manufacturing sector is responsible for the entire slowdown in business sector productivity growth between 2000 and 2007. Similarly, Baldwin and Willox (2016) suggests that the low productivity growth in three different industries (including oil extraction) explains the entirety of the slowdown in business sector labour productivity growth from 2000 to 2014. While we recognize the validity of these analyses, in Appendix B, we show that even if one excludes the manufacturing sector (or the agricultural, services, or "rest of mining" sectors) the lack of TFP growth persists. Therefore, while all these findings invite further investigations into areas where Canadian productivity growth may be improved, it is striking that the stagnation in aggregate TFP growth can be so singularly attributed to the Oil sector.

According to Keay (2009), the resource extraction sector had a positive impact on per capita economic performance in the Canadian economy from 1970 to 2005. Although our paper attributes the recent lack of TFP growth to the Oil sector, our growth accounting decomposition with and without the Oil sector also reveals that TFP growth was higher during the 1970s and the 1990s thanks to the Oil sector (see Table 2.1 in Section 2.3 for more information). Similarly, Olewiler (2017) suggests that Canada has benefited from exporting natural resources, but notes that failing to account for the environmental externalities of resource extraction raises concerns about the long-term economic benefits.

The notion that the oil sands are expensive and capital intensive has already been documented in the literature. Heyes, Leach, and Mason (2018) documents that a barrel of crude bitumen trades at $12.77 below the WTI price,[6] and the world oil price for new oil sands projects to be profitable is $9 higher than other oil extraction projects would be. Leach (2022) shows that, at a minimum, a new oil sands project requires an initial investment of $1 billion and takes up to 5 years for production to reach full capacity.

One way to evaluate the productivity for oil and gas extraction sectors is to look at the Energy Return on Investment (EROI), a measure that reports the ratio between energy produced per unit of energy used. Gagnon, Hall, and Brinker (2009), Hall, Lambert, and Balogh (2014) and others povide evidence that the EROI for oil and gas extraction at large has been declining. Poisson and Hall (2013) finds that in Canada, conventional oil and gas extraction EROI fell from 20-to-1 to 12-to-1 from the mid-1990's to 2008, while the oil sands has fluctuated around a significantly lower average EROI of 4-to-1. Brandt, Englander, and Bharadwaj (2013) suggests that although the EROI of oil sands extraction has improved over time, it remains less efficient than conventional oil production.

Finally, our paper is also related to a long-standing discussion about the Dutch Disease (DD) in Canada. The DD phenomenon arises when a sector (in this case, oil) captures

[6]This is significantly more discounted than Mexican Maya crude ($6.98 below WTI price).

factors of production to the extent that it ends up harming the rest of the economy. For instance, Beine, Bos, and Coulombe (2012) argues that over a third of Canada's manufacturing employment loss in the early 2000s is related to an appreciation of the exchange rate. Boadway, Coulombe, and Tremblay (2012) builds on those findings to analyse the policy challenges faced by the provincial and federal governments. Papyrakis and Raveh (2014) finds that Dutch Disease mechanisms are relevant at the regional level for Canada. On the other extreme, Carney (2012) writes about the DD in the following terms: *"[w]hile the tidiness of the argument is appealing and making commodities the scapegoat is tempting, the diagnosis is overly simplistic and, in the end, wrong."* Our paper's findings align with the DD story when it documents that the Oil sector has indeed captured an enormous fraction of the overall capital stock in Canada (from 6 per cent in 1961 to over 30 per cent in 2018). However, our findings do not align with the DD in that the rest of the economy has not experienced a systematic reduction in importance (approximately, non-oil accounts for 95 per cent of the economy throughout the period). Moreover, our results show that productivity for the Canadian economy without the Oil sector continues to grow at levels comparable to those of the United States and Norway, which is not necessarily compatible with what is usually thought of as the DD.

2.3 Canadian TFP, with and without the Oil sector

In this section, we describe our methodology for measuring TFP, which is then used to conduct a growth accounting decomposition of the Canadian economy. We also compare TFP growth rates with and without the Oil sector across different time periods. At the end, we also compare TFP growth rates in Alberta and Ontario.

2.3.1 How to measure TFP

To measure TFP, we assume a standard Cobb-Douglas production function, where GDP (Y_t) is a function of capital (K_t), labour (L_t), and a productivity factor (A_t):[7]

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha} \tag{2.3.1}$$

where α is the parameter that measures the capital intensity of the economy. We obtain data on GDP (Y_t), capital (K_t), labour (L_t), and compensation of employees from StatsCan.[8] Using the compensation of employees data, we can calculate the capital share of income (α) as:

$$\alpha = 1 - \frac{1}{T}\sum_t \frac{w_t \times L_t}{Y_t} \tag{2.3.2}$$

[7]Our results are obtained using series of GDP in constant prices. In Appendix C we re-compute the analysis using "Output-side real GDP at chained PPPs" from the Penn World Tables 10.01 (?, ?), which is useful "to compare relative productive capacity across countries and over time." Our results are the same regardless of which metric is used.

[8]See Appendix A for details.

where $w_t \times L_t$ is the series of compensation of employees.

With all this information, we can calculate TFP as a residual,

$$A_t = \frac{Y_t}{K_t^{\alpha} L_t^{1-\alpha}} \tag{2.3.3}$$

and decompose GDP per working-age population (N_t, henceforth WAP) as

$$\underbrace{\frac{Y_t}{N_t}}_{\text{GDP per WAP}} = \underbrace{A_t^{\frac{1}{1-\alpha}}}_{\text{TFP}} \times \underbrace{\left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}}}_{\text{Capital-Output ratio}} \times \underbrace{\frac{L_t}{N_t}}_{\text{Hours per WAP}} \tag{2.3.4}$$

2.3.2 Growth Accounting Decomposition with and without oil

The Growth Accounting Decomposition of the Canadian economy from 1961 to 2018 is shown in Figure 2.1, where the y-axis is presented in logarithmic scale due to the exponential growth observed in the series for GDP per WAP and TFP.

It is worth noting that from 1961 to 2001, the Canadian economy followed the typical pattern of developed economies, with almost all growth in GDP per WAP attributed to improvements in TFP. During this period, the Capital-Output ratio and Hours per WAP series remained relatively stable, with minor fluctuations reflecting business cycle movements. Neither variable had a significant impact on the overall growth of GDP per WAP.

Figure 2.1: Growth Accounting Decomposition.

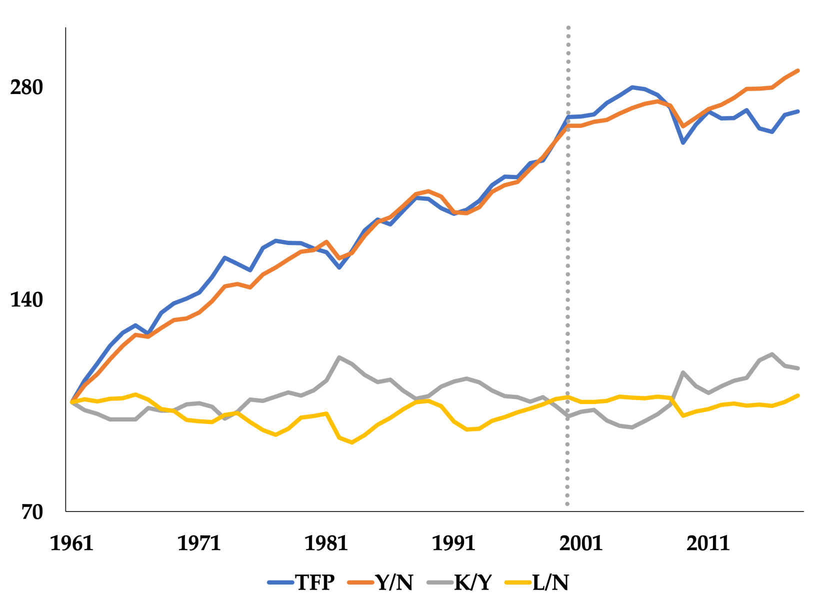

Beginning in 2001, the analysis reveals a marked shift in the Canadian economy's productivity trends, labelled as the "Canadian Productivity Stagnation" by Conesa and Pujolas (2019). Between 2001 and 2018, the TFP series exhibits little-to-no growth, fluctuating around a horizontal line and yielding an annualized growth rate of a mere 0.06 per cent.

In the next phase of our analysis, we eliminate all components of the Oil sector from our TFP calculation. We perform the same analysis but substitute $Y_t^{NO}, K_t^{NO}, L_t^{NO}$ in place of Y_t, K_t, L_t, respectively, where X_t^{NO} is defined as X_t minus X_t^{Oil} for X being Y, K, or L, and "NO" indicates "Net-of-Oil". To ensure consistency in our analysis, we revise α^{NO} to exclude Oil sector labour payments from the economy.[9] We present the results of the Net-of-Oil TFP Growth Accounting Decomposition exercise in Figure 2.2.

Figure 2.2: Net-of-Oil Growth Accounting Decomposition

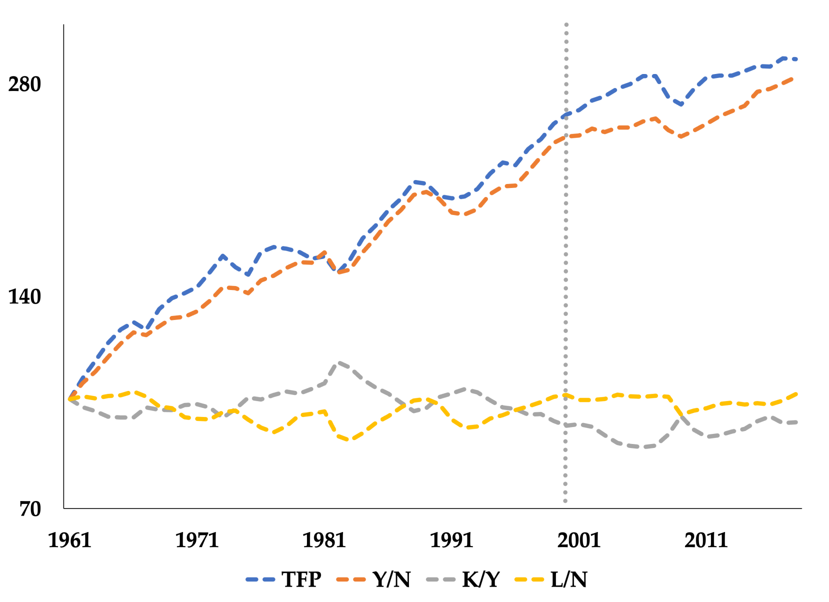

After removing the Oil sector components from the analysis, our results indicate that the economy did not display stagnant TFP post-2001. On the contrary, TFP increased at an annualized rate of 0.60 % between 2001 and 2018. These findings align with Canada's historical TFP growth rate.[10]

[9]We provide an exact description of how the Net-of-Oil series are calculated at the end of Appendix A. Moreover, we find that using the same α value as in the previous Net-of-Oil TFP analysis (Figure 2.1) has no impact on the results, as shown in Figures 2.20 in Appendix D. Likewise, we find that using the exact evolution in the labour share (using a time-varying α) does not affect our results either (Figures 2.21b and 2.21a also in Appendix D).

[10]We find that removing any sector other than oil from the Canadian economy is inconsequential for the evolution of TFP. The result of these exercises can be seen in Figures 2.14, 2.15, 2.16, and 2.17 in Appendix B.

TABLE 2.1: TFP and Net-of-Oil TFP growth rates

Period	**TFP growth**	**Net-of-Oil TFP growth**
1961-2018	0.99%	1.21%
1971-1981	0.78%	0.62%
1981-1991	0.75%	1.17%
1991-2001	1.90%	1.79%
2001-2018	0.06%	0.60%

Table 2.1 presents the annualized growth rates of TFP for different periods, comparing the economy with and without the Oil sector. TFP has grown at 0.99 % per year for the period considered, from 1961 to 2018. When the Oil sector is excluded, this figure increases to 1.21 %. This difference in growth rates, however, is not consistent throughout the period. Net-of-Oil TFP rises faster than overall TFP during the 1980s but not during the 1970s or the 1990s. Most importantly, while the annual TFP growth rate when the Oil sector is included is a meagre 0.06 % since 2001, the growth rate increases to 0.60 % during that time period when we exclude the Oil sector from the analysis. This figure is more in line with other growth rates presented in the table, like the Net-of-Oil TFP growth rate in the 1970s and the TFP growth in the 1970s and 1980s.

2.3.3 Alberta vs. Ontario

The impact of the oil extraction sector varies across different Canadian provinces. Alberta, which accounts for 17.26 % of Canadian GDP,[11] stands out for having a significant portion of its GDP derived from the Oil sector (24.27 %). In contrast, Ontario, which accounts for 38.94 % of Canadian GDP, only has a minimal fraction of its GDP attributed to the Oil sector (0.02 %). The significance of these two provinces in the overall economy, along with the substantial contrast in the Oil sector's importance within their respective economies, make them excellent candidates for studying the Oil sector's impact on TFP.

Figure 2.3a illustrates the growth accounting decomposition for Alberta when the Oil sector is included, while Figure 2.3b displays the growth accounting decomposition for Alberta when the Oil sector is excluded. Correspondingly, Figures 2.3c and 2.3d depict the equivalent growth accounting decompositions for Ontario.

Analysing the figures indicates a shared behaviour between Alberta's growth accounting decomposition figures without the Oil sector and Ontario's figures (whether considering the Oil sector or not, as the two are almost indistinguishable). Despite encountering several ups-and-downs, particularly notable during the Great Recession period, the figures present variations of the canonical growth accounting exercise figure of developed

[11] All the percentages in this paragraph are calculated based on the average from 1997 to 2018, which represents the available provincial-sectoral data time-frame. Refer to Appendix A for details regarding the datasets used in this analysis.

Figure 2.3: Provincial Growth Accounting Decompositions

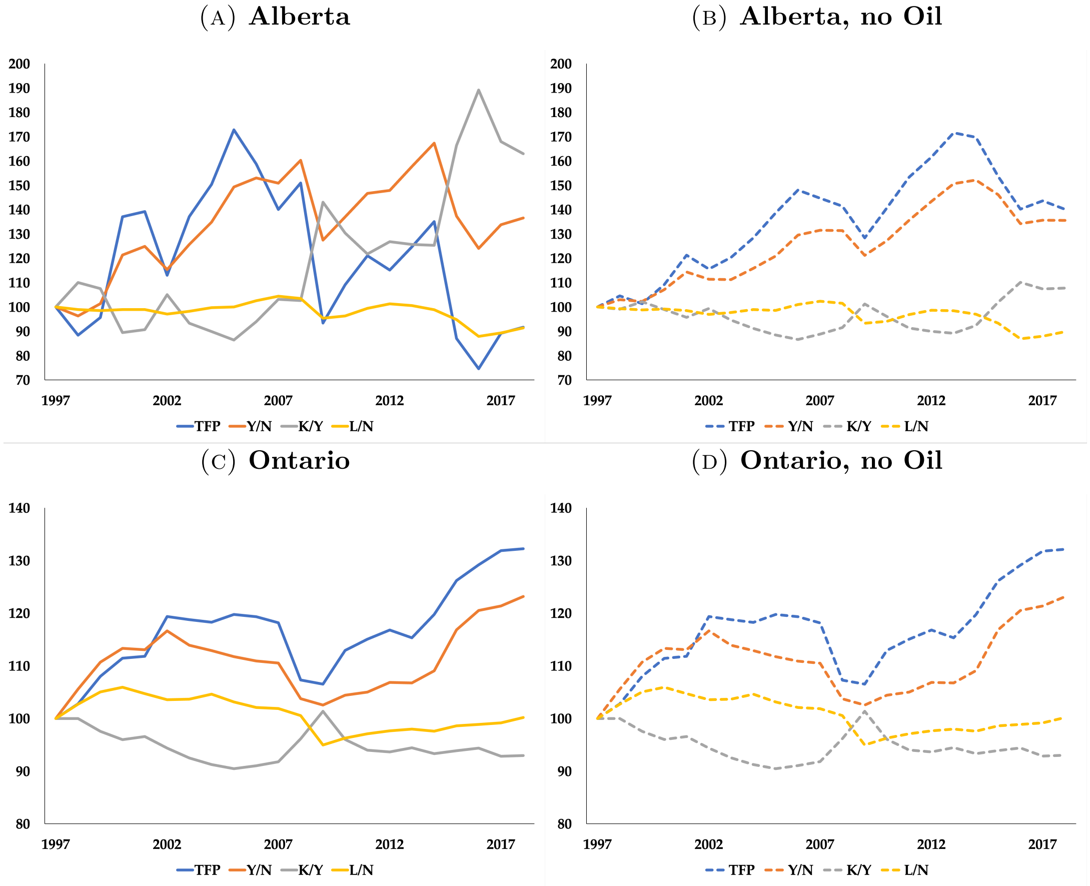

economies: two relatively flat lines (the Capital-Output ratio and the Hours per WAP) and two lines that move in parallel with an upward trend (TFP and GDP per WAP).

The figure that notably stands out is Alberta's growth accounting decomposition with the Oil sector. Initially, there is a remarkable surge in TFP from 1997 to the early 2000s, followed by a sharp decline that erases all the gained productivity by 2018. Interestingly, GDP per WAP decouples from the TFP trajectory, continues to rise (albeit with some fluctuations along the way), and the increase aligns with the significant upswing in the Capita-to-Output ratio starting in the mid-2000s.

Overall, the four pictures reinforce the story that Canadian productivity sluggishness since the turn of the century is largely driven by the Oil sector. Next, we show that a comparison of Canada to the United States and Norway reinforces this result.

2.4 International Comparisons

In this section, we compare Canada to the United States, a standard benchmark, and to Norway, a small, open economy with a large Oil sector. These three countries are the three largest oil producers among rich, western economies.

2.4.1 Canada and the United States

We explore whether the lack of TFP growth driven by the Oil sector, which was observed in Canada post-2000, also occurs in the United States. Specifically, we investigate the evolution of TFP and Net-of-Oil TFP in the United States and compare them to their Canadian counterparts. We find that, unlike Canada, excluding the Oil sector has no significant effect on measured aggregate TFP growth in the United States.

We use the same approach as in our analysis of Canada, utilizing data on capital, value added, hours worked, and labour compensation from the Bureau of Economic Analysis in the United States.[12]

Figure 2.4a displays the TFP and Net-of-Oil TFP evolution for Canada and the United States from 1961 to 2018, normalized to 100 in 1961.[13] As expected, all series increase with some fluctuations, and the US curve is smoother. Nevertheless, until 2000, both Canadian TFP series followed similar trends. The primary contrast between the two Canadian series emerges after this year.

FIGURE 2.4: TFP and Net-of-Oil TFP, Canada and United States

(A) **Since 1961** (B) **Since 2001**

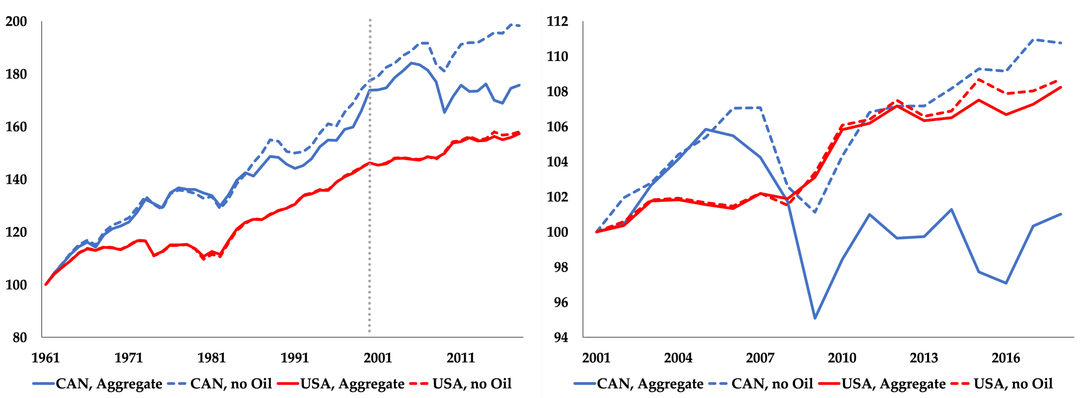

To gain insights into the evolution of TFP post-2000, we narrow our focus in Figure 2.4b, normalizing the series to 100 in 2001. The figure highlights a clear trend: while all the TFP series for Canada and the United States have business cycles, only Canada's TFP series displays stagnant growth post-2001. In contrast, the TFP series for the United States and the Net-of-Oil TFP series for both countries exhibit strikingly similar growth patterns. They all experience fluctuations but consistently grow throughout the period, indicating that the lack of TFP growth in Canada is a unique phenomenon, linked to the Oil sector.

[12] See Appendix A for details.

[13] While Canadian TFP has grown faster than the US TFP during this period, it is still the case that Canadian output per person is lower.

We consider Figure 2.4b to be a clear representation of the key message conveyed in this paper. If measured without the Oil sector, Canada's TFP growth would have been similar to that of the United States, and there would not have been a "Canadian Productivity Stagnation." Except for the Oil sector, Canada has been on par with the United states in terms of TFP growth.

Figure 2.5 illustrates the trend in the ratio of capital invested in the Oil sector compared to the total capital in the Canadian economy (dashed blue line) and the ratio of value added contributed by the Oil sector as a fraction of the total value added in Canada (solid blue line). Respectively, the lines for the United States are plotted in red.

FIGURE 2.5: Importance of oil, Canada and United States

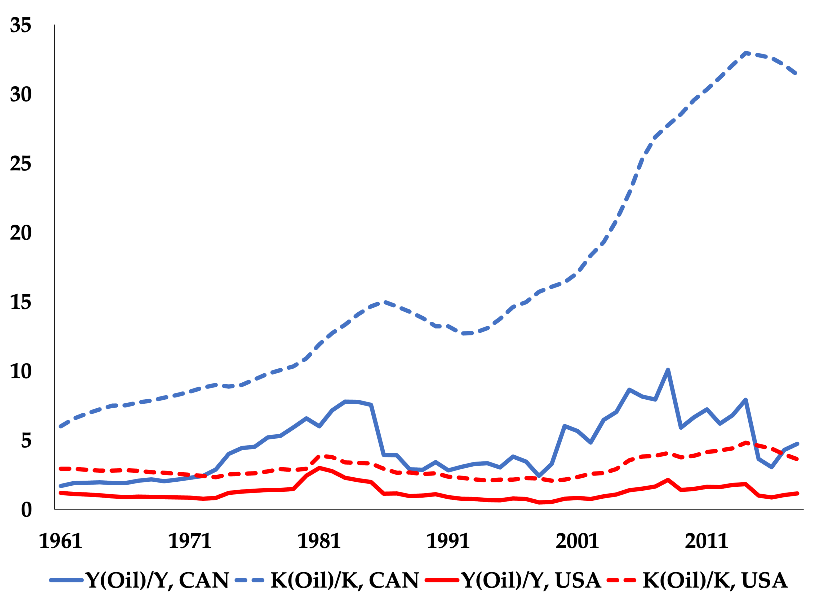

For Canada, during the period of analysis, the Oil sector's share of value added hovers around a trend of about 5 per cent, while the proportion of capital allocated to the Oil sector has increased from 6 per cent to 31 per cent. This growth is mostly observed after 2001, suggesting that the Oil sector is using a relatively greater amount of input to produce relatively the same amount of output. Since the Oil sector has low labour requirements, the divergence between the growth of capital and value added mechanically accounts for the lower TFP.

For the United States, on the other hand, the Oil sector's share of value added remains around 1 per cent and the proportion of capital allocated to the Oil sector remains at around 3 per cent. The relative constant percentages for the importance of the Oil sector in the US explains why TFP behaves similarly to Net-of-Oil TFP.

2.4.2 Canada and Norway

Norway's oil production accounts for 20 per cent of its GDP.[14] At the same time, Norway is a small, open, very rich economy. The importance of the Oil sector, together with its developed economy makes it an excellent candidate to compare with Canada. We find that, even though the Norwegian economy has fared better than the Canadian one in the last half a century, the evolution of TFP and the Net-of-Oil TFP of both countries exhibit a remarkably similar story during the last two decades. To conduct our analysis, we again use the same analysis of Canada, utilizing the relevant data for Norway from the Statistisk sentralbyrå.[15]

Figure 2.6a displays the TFP and Net-of-Oil TFP evolution for Canada and Norway from 1970 to 2018, normalized to 100 in 1970.[16] As expected, all series increase with some fluctuations, and the Canadian curve is smoother. Until the 2000s, Norwegian TFP grows faster than Canadian TFP. The pattern that emerges from 1970 to 2000 is that of two similar lines for Norway (in orange, substantially higher) and two similar lines for Canada (in blue, substantially lower). Since 2000, however, we note that TFP growth in both countries is stagnant.

Figure 2.6: TFP and Net-of-Oil TFP, Canada and Norway

(A) **Since 1970** (B) **Since 2001**

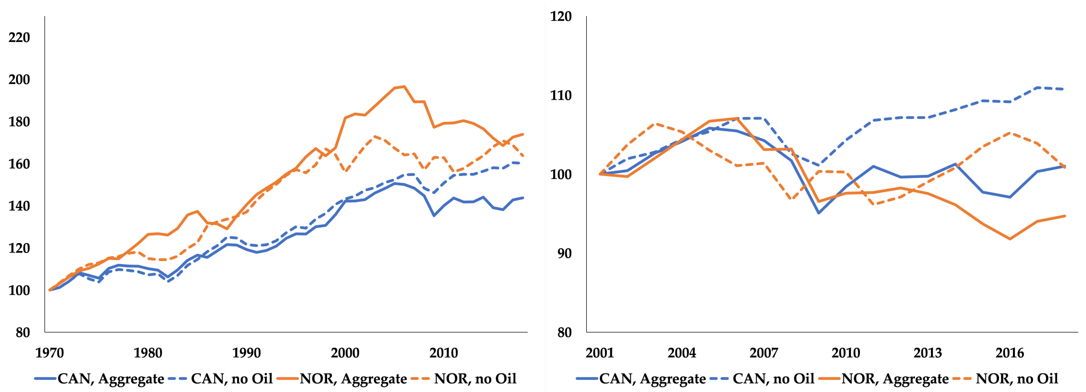

Figure 2.6b displays the TFP and Net-of-Oil TFP evolution for Canada and Norway from 2001 to 2018. Both countries have a sluggish evolution in their overall TFP (solid lines) with Norwegian TFP decreasing by 0.32 per cent (orange line) and Canadian TFP showing a slight increase of 0.06 per cent (blue line). Notably, Net-of-Oil TFP exhibits more favourable outcomes (dotted lines). Norwegian Net-of-Oil TFP remains relatively stable, with a modest increase of 0.05 per cent (orange line) wile Canadian Net-of-Oil TFP experiences increases of 0.60 per cent (blue line). In summary, excluding the Oil

[14] To keep the comparison figures the same with the previous section, this average is also from 1997 to 2018.

[15] See Appendix A for details.

[16] Canadian TFP has grown slower than Norway's during this period, and Canadian output per person is lower than its Norwegian counterpart.

sector results in an annual growth rate that is 0.37 per cent higher in Norway and 0.54 per cent higher in Canada. The primary factor behind this outcome is the allocation of capital.

Figure 2.7 reproduces Figure 2.5 but changing the United States for Norway, depicting the evolution of the ratio of capital invested in the Oil sector compared to the total capital (dashed orange line) and the ratio of value added contributed by the Oil sector as a fraction of the total value added in Norway (solid orange line).

FIGURE 2.7: Importance of oil, Canada and Norway

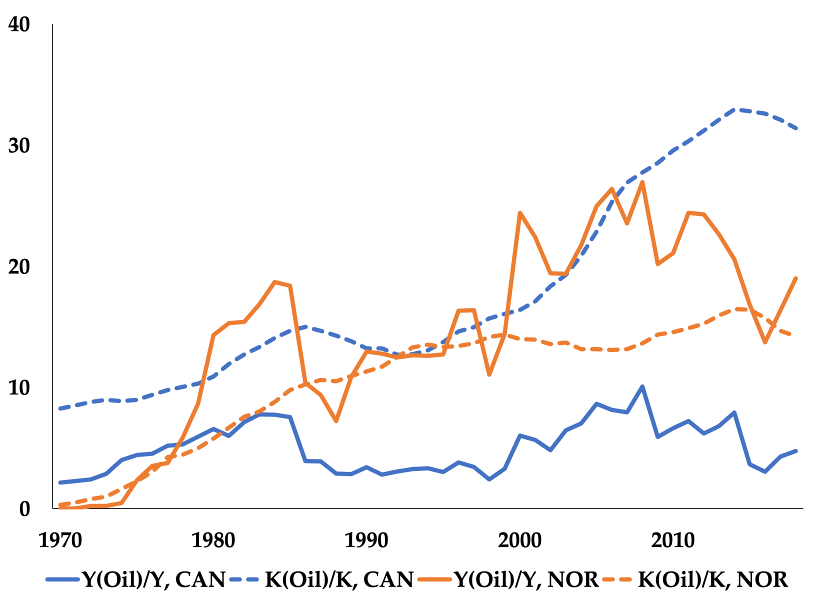

In the context of Norway, the Oil sector's contribution to value added experiences a notable growth until the mid-2000s, followed by a consistent decline leading up to 2018. During this period, there is an increase in the allocation of capital to the Oil sector, albeit with a slight decrease towards the end of the period. From 2001 to 2018, the importance of the Oil sector in terms of value added diminishes, while its importance in capital allocation sees a modest rise. In the case of Canada, the Oil sector's share of value added remains relatively stable, with occasional fluctuations. However, the proportion of capital allocated to the Oil sector sees a significant increase, particularly during the late 1990s and early 2000s. Over the period from 2001 to 2018, the importance of the Oil sector in value added remains fairly constant, while its importance in capital allocation increases substantially.

In both instances, the observed trend is characterized by a decline or stagnation in value added (the numerator in TFP) while capital (which appears in the denominator of TFP) increases. Consequently, the Oil sector plays a crucial role in explaining the lower growth rate of TFP compared to Net-of-Oil TFP in both Norway and Canada.

2.5 More details about the Canadian Oil sector

In this section we show, first, that TFP in the Canadian Oil sector is a rarity compared to the other sectors of the economy. Then, we propose two potential explanations for the observed fall in TFP: a composition effect, and a misalloaction of factors effect. While the evidence we provide aligns well with the former, our evidence does not allow us to confirm or deny the latter. Last, we explain why it is unlikely that Canadian data alone will suffice to disentangle the question of misalloaction.

2.5.1 Canadian Oil sector TFP: a rarity

In Figure 2.8a we document a significant difference in the evolution of TFP between the Oil sector and the rest of the economy. While TFP for the overall economy grows, TFP for the Oil sector experiences a secular decline. Thus, the Oil sector in Canada stands out as a distinctive sector with unique characteristics.

FIGURE 2.8: TFP: Oil vs rest of the economy

(A) **Oil vs aggregate** (B) **Oil vs other sectors**

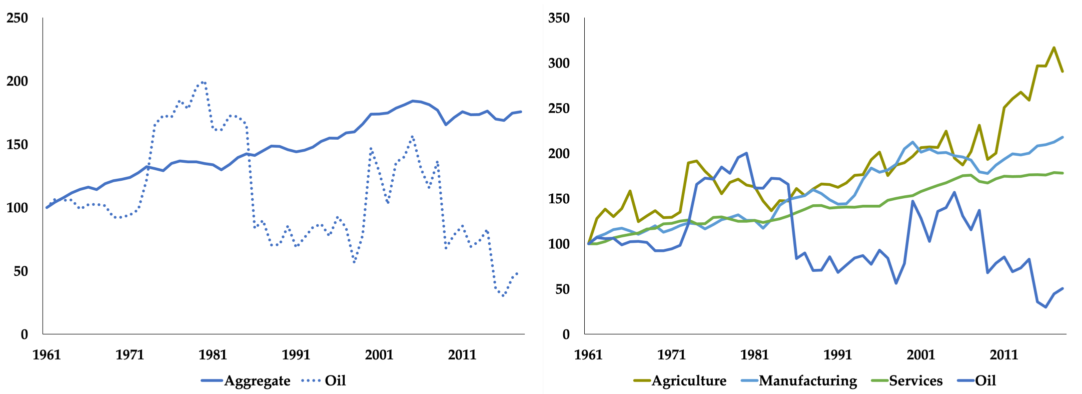

To further support this perspective, Figure 2.8b illustrates the trends of TFP in the Oil sector, as well as in Agriculture, Manufacturing, and Services sectors. The volatile, declining pattern of TFP in the Oil sector is not observed in any of the other sectors. On the contrary, other sectors' TFP grows in tandem with the overall economy's TFP over the entire period.

2.5.2 A composition effect

Oil TFP may be falling due to a composition effect. For the sake of argument assume that the oil sector has two sub-sectors. The traditional oil sub-sector is highly productive and operates at capacity; the oil sands sub-sector has low productivity and decreasing returns to scale. If the world oil price increases, the traditional oil sub-sector cannot expand its production; however, the oil sands sub-sector will expand its output using more factors of production, even if the marginal unit produced will be increasingly unproductive. As a result, the effect of a price increase will decrease sectoral TFP.

The assumptions about productivity in the two sub-sectors is consistent with the heavy discount on a barrel of crude bitumen as well as the massive costs of starting an oil sands project as noted in Heyes, Leach, and Mason (2018) and Leach (2022). When the oil price boom occurred, it became profitable to invest heavily in the sector. As capital flowed in, the sectoral productivity mechanically fell as resources were allocated to the less productive type of oil.

In Figure 2.9a we plot the percentage of capital and investment allocated to the Oil sands as a fraction of the total capital and investment in the Oil sector. The proportion of capital in the Oil sands ranged from 5 to 8 per cent between 1961 and 2001, but it surged to about 30 per cent afterwards. Similarly, the proportion of investment in the Oil sands, which ranged at similar values between until the mid-1990's albeit more erratically, reaching a peak of 44 per cent post-2001.

FIGURE 2.9: Oil sands capital and WTI price

(A) **Oil sands K and I** (B) **WTI price**

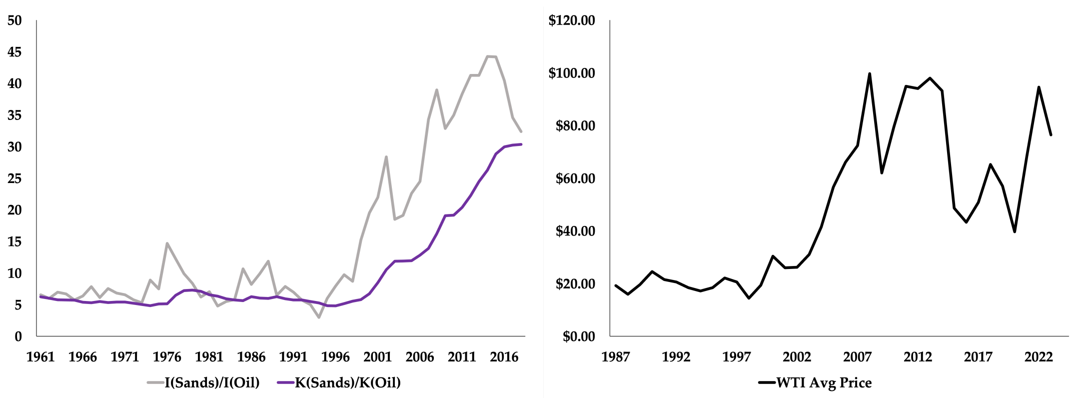

The increase of capital allocated to the Oil sands is consistent with two occurrences: the technological advancement allowing for the opening of the first commercial Steam-assisted gravity drainage (SAGD) project at Foster Creek in 1996 and the potential profitability of exploitation.[17] Figure 2.9b plots the evolution of the Oil price, measured as the West Texas Intermediate. It was roughly $20 per barrel in the 1990s but surged to approximately $100 per barrel by 2007 and has since fluctuated but at significantly higher prices than before 2001.

[17] More details on the history of the Oil sands and technology used can also be found in the 2008 report prepared for the US Congress (Humphries, 2008).

2.5.3 Capital reallocation

Another possibility regarding the fall in Oil sector TFP is that there has been a surge in resource misallocation.[18] Given the rapid growth in the percentage of the aggregate capital stock installed in the Oil sector, it is worth asking whether this came at the expense of more productive sectors.

The Oil sector is a particularly capital intensive sector, with a labour share of 0.11. To be able to show that there is an increase in misallocation, we would have to demonstrate that in the absence of the oil price boom, capital would have been installed in more productive sectors. However, no other sector is as capital intensive as the Oil sector, making it difficult to establish an appropriate counterfactual. It is also not obvious that aggregate investment in the capital stock would have been comparable had the oil price boom never occurred.

The closest comparable sector to look at is Mining excluding oil and gas (NAICS code 212). In Figure 2.10a we show that in the rest of the mining sector (RoM), TFP is relatively flat. It spikes during the 2000s commodity boom (as does Oil). However, Figure 2.10b shows that the RoM share of the aggregate capital stock has historically been below oil, and remains fairly constant during the period when the Oil sector's share explodes. Hence, it does not seem that capital flowed in that direction.

FIGURE 2.10: TFP and Capital, Oil and Other Mining

(A) **TFP** (B) **Percentage of Aggregate Capital Stock**

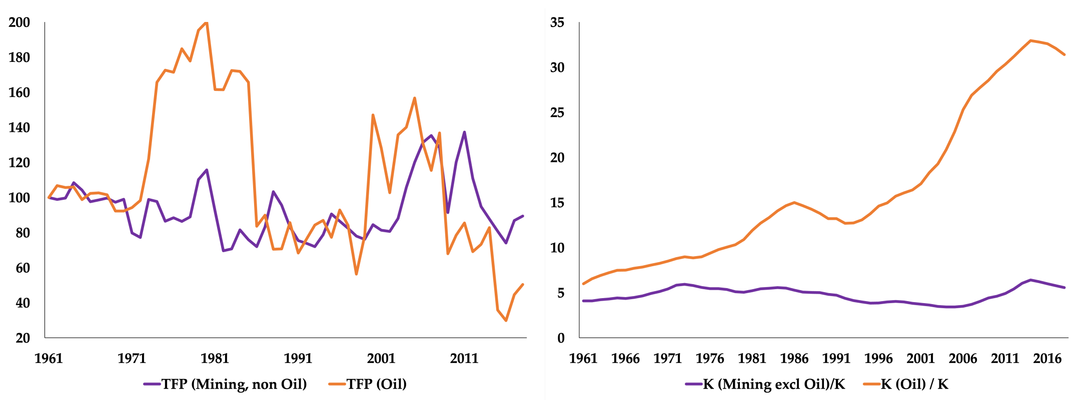

Alternatively, it is worth investigating whether misallocation happens as capital is removed from the Manufacturing sector into the Oil sector — Figure 2.11 shows that the decrease in the former coincides with an increase in the latter. Before we dig into the details, it is worth noting that the decline in capital allocated to the Manufacturing sector is consistent with the canonical pattern of structural transformation during this

[18]The literature on misallocation, started by Restuccia and Rogerson (2008), Guner, Ventura, and Xu (2008 and covered in Restuccia and Rogerson (2017), shows how capital misallocation can have large impacts on aggregate productivity.

time period (see, for instance, Herrendorf, Rogerson, and Valentinyi, 2013). The overall fall in capital installed in Manufacturing can be seen in Canada (it fell from 16.97% in 2001 to 8.06% in 2018) but also, even if less dramatically, in the United States (from 10.37% in 2001 to 8.54% in 2018).

FIGURE 2.11: Percentage of Aggregate Capital Stock by sector

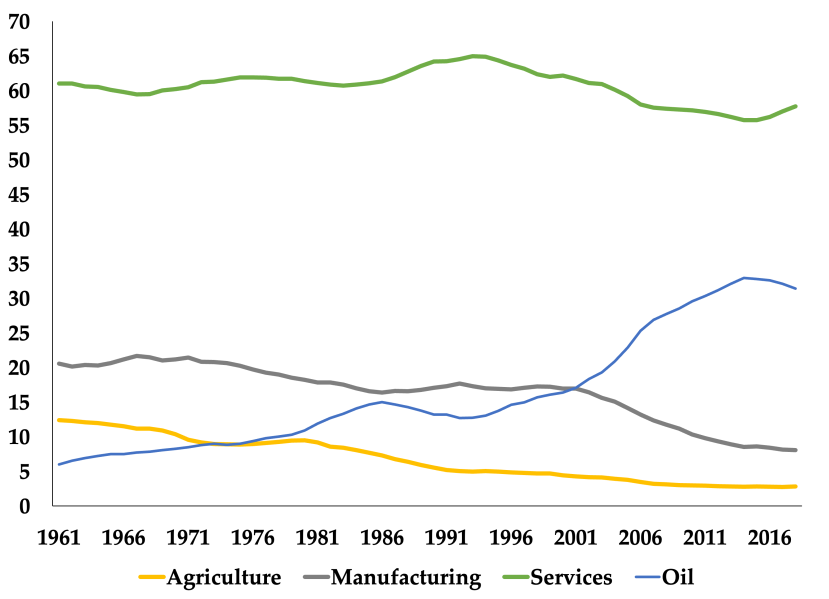

The data from Statistics Canada splits capital into Engineering Construction (EC), Intellectual Property Products (IPP), and Machinery & Equipment (ME). Figures 2.12a and 2.12b plot the evolution of each type of capital as a percentage of the Oil sector's and the Manufacturing sector's total capital stock. The first thing to note is that the Oil sector's capital stock is dominated by EC, and this importance grows during the oil price boom. By comparison, EC in the Manufacturing sector is small and stays relatively constant during that time period. On the other hand, the capital stock in the Manufacturing sector largely consists of ME, which falls during the capital boom in the Oil sector. However, the ME in the Oil sector stays relatively flat.

These findings indicate that each sector is reliant on a different type of capital, and that the type of capital that grows in importance in the Oil sector during the boom is different than the type of capital that falls in the Manufacturing sector during that same period.

Still, the comparison we make may be masking more substantive patterns because the denominator in the the two figures are different — the aggregate capital in each sector. In the three panels of Figure 2.13 we plot the evolution of the different types of capital split between Manufacturing and Oil sector, and the latter further divided into Oil sands and traditional oil.

FIGURE 2.12: Capital Composition, Oil and Manufacturing

(A) **Percentage of Oil Capital**

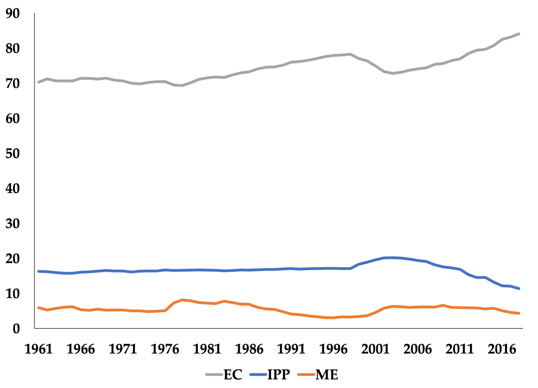

(B) **Percentage of Manufacturing Capital**

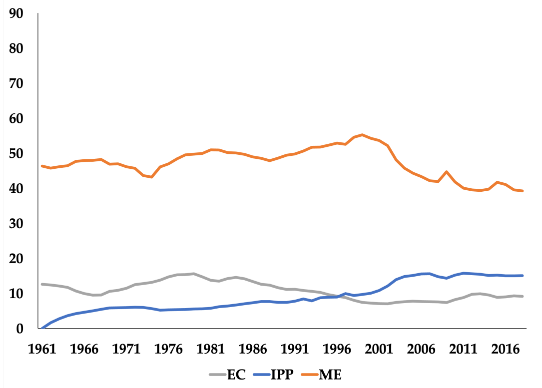

FIGURE 2.13: Capital types in Oil and Manufacturing

(A) **EC capital**

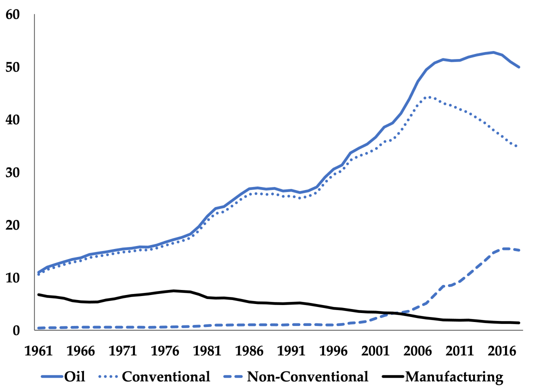

(B) **IPP capital**

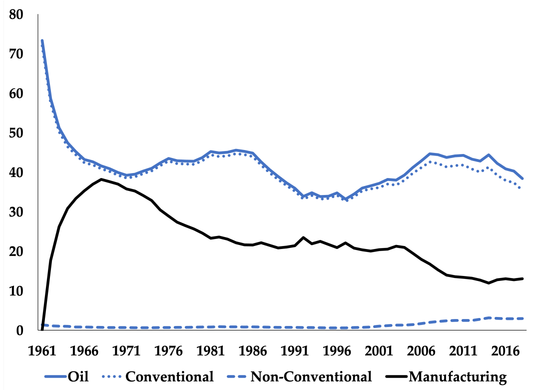

(C) **ME capital**

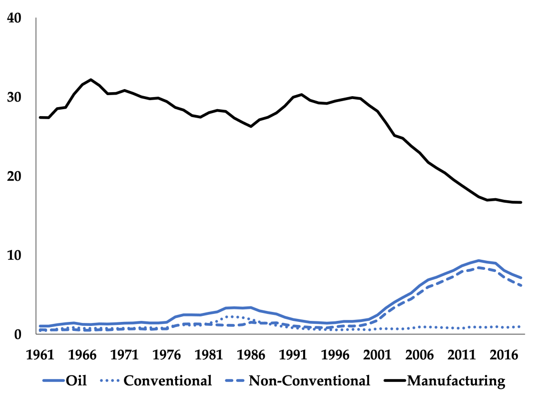

Figure 2.13a shows that EC is largely used in the Oil sector, initially only in the traditional sector, and it spikes in the Oil sands after the 2000s. During this time period, the percentage of EC in the Manufacturing sector is very small throughout. Hence, it is

unlikely that there is a movement of capital between the two sectors.

Figure 2.13b shows that IPP was mostly installed in the Oil sector in the 1960s. In the early 1970s it then stabilized at around 41 per cent in Oil and 31 per cent in Manufacturing. Since then, we can observe a secular increase in Oil to about 44 per cent and a secular decrease in Manufacturing to about 12 per cent. The decrease in Manufacturing is more pronounced after 2001.

Last, Figure 2.13c shows that ME increases dramatically in the Oil sector after 2001 (from 2 per cent to 7 per cent), and it falls substantially in the Manufacturing sector during the same time period (from 28 per cent to 17 per cent). The increase in the Oil sector is largely caused by the increase in the Oil sands.

The picture that emerges from comparing the types of capital installed in the Manufacturing sector to those installed in the Oil sector is that capital (especially IPP and, more prominently, ME) could have been reallocated from the Manufacturing sector to the Oil sector — and notably, to the Oil sands.

That being said, the potential evidence of factor reallocation does not necessarily mean that there is misallocation. It could well be that the units of capital in the Oil sands became less productive than when they were installed in manufacturing — or the other way around. To be able to disentangle this question, we need more disaggregated data, at the firm, and probably also household (owner) level. In the absence of more disaggregated data, we cannot conclusively state that the growth in the capital stock allocated to the Oil sector comes at the expense of more productive activities in the Manufacturing (or any other non-Oil) sector. It could well be that the capital was installed to take advantage of high oil prices and would not have otherwise been used. Our aggregate analysis cannot provide conclusive evidence either way.

2.6 Concluding remarks

Our research shows that the Oil sector is the primary reason for the lack of TFP growth in Canada (and Norway) and that it does not generate a similar lack of TFP growth in the United States. Hence, our result concludes that the difference in productivity growth between Canada and the United States can be entirely attributed to the Oil sector.

While we believe that our result sheds light on the underlying cause of differential evolution of productivity between the two countries, it should be used with caution: we find that the Oil sector can explain, in an accounting sense, the lack of Canadian TFP growth. Our findings do not, however, make a judgement on the desirability of this result, nor get into the debate surrounding an industry that confronts a host of challenges, such as carbon emissions.

Namely, it is perfectly plausible that the Canadian economy is responding optimally by exploiting a resource when its value is very high. Whether this represents the best course of action is a question that requires further exploration. Consequently, we defer the answer to this crucial issue to future research.

We view the question of whether the explosion in the Oil sector capital stock constitutes a case of resource misallocation as an important one and worthy of future research. It would be valuable to know whether capital moved from more productive sectors like Manufacturing in response to changes in oil prices, or if the observed trend is a result of an investment boom that was directed to more profitable uses of capital. To properly answer this question, one will likely require firm level data.

Similarly, since our data spans only until 2018, our work is silent about the well-documented post-Covid productivity decline.[19] Still, there has been a surge in oil prices that may be behind part of the sluggish behaviour of Canadian productivity. Understanding how oil prices have interacted with the post-Covid productivity decline is another promising area of future research.

References

Alexopoulos, M. and Cohen, J. (2018). 'Canadian Productivity Growth, Secular Stagnation, and Technological Change'. International Productivity Monitor , volume 35, 113–137.

Baldwin, J. R. and Willox, M. (2016). 'The Industry Origins of Canada's Weaker Labour Productivity Performance and the Role of Structural Adjustment in the Post-2000 Period'. International Productivity Monitor , volume 31, 19–36.

Beine, M., Bos, C., and Coulombe, S. (2012). 'Does the Canadian Economy Suffer from Dutch Disease?' Environmental and Resource Economics, volume 34, 468–492.

Boadway, R., Coulombe, S., and Tremblay, J. (2012). 'The Dutch Disease and the Canadian Economy: Challenges for Policy-Makers'. Thinking Outside the Box, a Conference in Celebration of Thomas J. Courchene.

Boothe, P. and Roy, R. (2008). 'Business Sector Productivity in Canada: What Do We Know?' International Productivity Monitor , volume 16, 3–13.

Brandt, A., Englander, J., and Bharadwaj, S. (2013). 'The Energy Efficiency of Oil Sands Extraction: Energy Return Ratios from 1970 to 2010'. Energy, volume 55, 693–702.

Carney, M. (2012). 'Dutch Disease'. Bank of Canada speeches and appearances.

Cette, G., Fernald, J., and Mojon, B. (2016). 'The Pre-Great Recession Slowdown in Productivity'. European Economic Review, volume 88.

Conesa, J. C., Kehoe, T. J., and Ruhl, K. J. (2007). 'Modelling Great Depressions: The Depression in Finland in the 1990's'. In Kehoe, T. J. and Prescott, E. C., editors,

[19] See Tombe (2023) for a recent description of this decline.

'Great Depressions of the Twentieth Century', pages 427–475. Federal Reserve Bank of Minneapolis.

Conesa, J. C. and Pujolas, P. (2019). 'The Canadian Productivity Stagnation, 2002-2014'. Canadian Journal of Economics, volume 52, no. 2, 561–583.

Currie, B., Scott, L. W., and Dunn, A. W. (2021). 'The Future of Productivity: Clear Choices for a Competitive Canada'. Deloitte's Future of Canada series.

Douglas, C., Di Matteo, L., Fuss, J., Geloso, V., Globerman, S., Johan, S., Jones, L., Mintz, J., Murphy P., R., Sobel, R. S., Tombe, T., and Whalen, A. (2021). Achieving the FourDay Work Week: Essays on Improving Productivity Growth in Canada. Fraser Institute

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). 'The Next Generation of the Penn World Table'. American Economic Review, volume 105, no. 10, 3150–3182.

Fernald, J. (2015). 'Productivity and Output Before, During and After the Great Recession'. In Parker, J. and Woodford, M., editors, 'NBER Macroeconomics Annual 2014', pages 1–51. University of Chicago Press, Chicago.

Fong, F. (2019). 'The Solution to Canada's Plummeting Productivity'. CPA Pivot Magazine.

Gagnon, N., Hall, C., and Brinker, L. (2009). 'A Preliminary Investigation of Energy Return on Energy Investment for Global Oil and Gas Production'. Energies, volume 2, no. 3, 490–503.

Gu, W. (2018). 'Accounting for Slower Productivity Growth in the Canadian Business Sector after 2000: the Role of Capital Measurement Issues'. International Productivity Monitor, volume 34, 21–39.

Guner, N., Ventura, G., and Xu, Y. (2008). 'Macroeconomic Implications of Size-dependent Policies'. Review of Economic Dynamics, , no. 11, 721–744.

Hall, C., Lambert, J., and Balgoh, S. (2014). 'EROI of Different Fuels and the Implications for Society'. Energy Policy, volume 64, 141–152.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2013). 'Two Perspectives on Preferences and Structural Transformation'. American Economic Review, volume 103, no. 7, 2752–2589.

Heyes, A., Leach, A., and Mason, C. (2018). 'The Economics of Canadian Oil Sands'. Review of Environmental Economics and Policy, volume 12, no. 2, 242–263.

Humphries, M. (2008). 'North American Oil Sands: History of Development, Prospects for the Future'.

Karabarbounis, L. and Neiman, B. (2014). 'The Global Decline of the Labor Share'. The Quarterly Journal of Economics, volume 129, no. 1, 61–103.

Keay, I. (2009). ‘Resource Specialization and Economic Performance: A Canadian Case Study, 1970-2005’. Canadian Public Policy, volume 35, no. 3, 291–313.

Kehoe, T. J. and Prescott, E. C. (2007). ‘Great Depressions of the Twentieth Century’. In Kehoe, T. J. and Prescott, E. C., editors, ‘Great Depressions of the Twentieth Century’, pages 1–20. Federal Reserve Bank of Minneapolis.

Leach, A. (2022). ‘Canada’s Oil Sands in a Carbon-constrained World’. Canadian Foreign Policy Journal, volume 28, no. 3, 285–304.

OECD (2023). ‘Economic Survey of Canada - 6 March 2023’. OECD Economic Snapshots.

Olewiler, N. (2017). ‘Canada’s Dependence on Natural Capital Wealth: Was Innis Wrong?’. Canadian Journal of Economics, volume 50, no. 4, 927–964.

Papyrakis, E. and Raveh, O. (2014). ‘An Empirical Analysis of a Regional Dutch Disease: The Case of Canada’. Environmental and Resource Economics, volume 58, 179–198.

Poisson, A. and Hall, C. (2013). ‘Time Series EROI for Canadian Oil and Gas’. Energies, volume 6, 5940–5959.

Restuccia, D. and Rogerson, R. (2008). ‘Policy Distortions and Aggregate Productivity with Heterogeneous Establishments’. Review of Economic Dynamics, volume 11, no. 4, 707–720.

Restuccia, D. and Rogerson, R. (2017). ‘The Causes and Costs of Misallocation’. Journal of Economic Perspectives, volume 31, no. 3, 151–174.

Robson, W. (2022). ‘Faster Productivity Growth Would Solve Many Problems’. Financial Post.

Shao, L. and Tang, R. (2021). ‘Allocative Efficiency and Aggregate Productivity Growth in Canada and the United States’. Bank of Canada Staff Working Papers, volume 21-1.

Sharpe, A. (2010). ‘Can Sectoral Reallocations of Labour Explain Canada’s Abysmal Productivity Performance?’ International Productivity Monitor , volume 19, 40–49.

Solow, R. M. (1957). ‘Technical Change and the Aggregate Production Function’. The Review of Economic Studies, volume 39, no. 3, 312–320.

Tombe, T. (2023). ‘As productivity plunges, Ontario and Alabama now have the same per capita GDP’. The Hub.

2.7 Appendix

A Data details

In this Appendix we describe the data sources we use to compute all the figures in our paper (Sections 2.3, 2.4, and 2.5). Note that all the references to “Table” in this section allude to on-line tables obtained from Statistics Canada, the Bureau of Economic Analysis, and Statistisk sentralbyrå (Statistics Norway). Specifically, we examine the sectoral information for all business sectors of the economy that report capital stock, hours worked, and value added.

We use on-line Tables 36-10-0217-01, 36-10-0208-01 and 36-10-0096-01 from Statistics Canada for aggregate GDP, hours worked and compensation. The latter refers to row "Business sector" in on-line Table 36-10-0208-01, which consists of the entire economy net of public administration, non-profit institutions and the rental value of owner-occupied dwellings. Capital ($K_{i,t}$) in each sector i for each year t is constructed using data (in current prices) on investment ($I_{i,t}$), geometric depreciation ($\delta_{i,t}K_{i,t}$) and geometric end-year net stock ($K_{i,t+1}$) so that in year t and sector i

$$K_{i,t} = K_{i,t+1} + \delta_{i,t}K_{i,t} - I_{i,t} \tag{A1}$$

The capital series is then deflated using the deflator for each year implied by the aggregate value added in current and 2012 prices. The aggregate capital series is constructed by subtracting the government sector and non-profits from the the investment, depreciation and end-year net stock of Total Industries.

The Canadian, provincial-level data is also taken from Statistics Canada, on-line Tables 36-10-0211-01 (aggregate GDP, hours worked and compensation), 36-10-0402-01 (Oil sector GDP), 36-10-0489-01 (hours worked and compensation in the Oil Sector), 36-10-0096-01 (capital).

The data for the United States is produced analogously using on-line data from the Bureau of Economic Analysis (BEA). Using on-line Tables on Value Added by Industry, we obtain value added for each industry in current prices and compute value added in 2012 prices using the tables for Chain-Type Price Indexes for Value Added by Industry. We combine the values from the current tables for the years 1997-2018 with the historic tables that cover 1961-1997. Where discrepancies for the year 1997 exist, we use the values from the current tables. Using the current price and 2012 price values for value added, we compute the aggregate deflator, which is then used to produce real valued estimates for value added in each industry.

For Capital, we combine on-line Tables 3.1ESI, 3.4ESI and 3.7ESI on net-stock, depreciation and investment of Private Fixed Assets by Industry respectively in the same way described above. We then deflate the value of the capital stock using the aggregate deflator for value added.

For hours worked we use on-line Tables 6.9B, 6.9C and 6.9D. Due to discrepancies between on-line Tables 6.9C and 6.9D, we use the data for 1998-2018 from on-line Table 6.9D, and 1987-1997 from on-line Table 6.9C. The tables lack data for hours worked in oil and gas extraction. To get around this, we compute average hours worked in mining by dividing total hours worked in mining by the number of Full-Time and Part-Time employees in mining (from on-line Tables 6.4B, 6.4C and 6.4D), and then multiply this value by the number of Full-Time and Part-Time employees in oil and gas extraction:

$$Hours_{oil} = Hours_{mining} \times \frac{Employees_{oil}}{Employees_{mining}} \tag{A2}$$

Similar to hours worked, the values for 1998-2018 come from on-line Table 6.4D. Finally, for compensation we use on-line Tables 6.2B, 6.2C and 6.2D in the same way.

The data used for Norway comes from the Annual National Accounts. Data on value added is taken from on-line Table 9 (Value added by kind of main activity at basic values). Data on employee compensation comes from on-line Table 13 (Compensation of employees by kind of main activity) and hours worked are taken from on-line Table 15 (Total hours worked by kind of main activity, Employees and self-employed). The capital stock is taken from on-line Table 37 (Fixed assets by kind of main activity). Data on value added, the capital stock and employee compensation are reported in current prices. They are then converted into real terms using the GDP deflator taken from the OECD. From each table, the aggregate series refers to the "Total Industry" row and the Oil sector series refers to the "Oil and gas extraction" row.

The parameters of the production function and TFP for all three countries are computed as described in Section (2.3). Data on working age population in each country is taken from the OECD.

For each country, the Net-of-Oil aggregates are constructed by taking the aggregate values for value added, the capital stock, hours worked and employee compensation and subtracting the corresponding values from the Oil sector. For monetary values (GDP and capital), we take aggregate variable in current prices and subtract the corresponding variable in oil extraction in current prices. This value is then deflated using the aggregate GDP deflator to obtain the Net-of-Oil variables in real terms. Net-of-Oil hours worked is obtained by subtracting hours worked in the Oil sector from aggregate hours worked. We then recompute the labour share and TFP using the Net-of-Oil values instead of the aggregate values.

B Alternative exclusions in GAD exercise

This appendix section presents alternative growth accounting decompositions that differ from those in the main text. Specifically, we demonstrate that the correction of the TFP series that arises when we exclude the Oil sector does not occur when we exclude other sectors such as Agriculture, Manufacturing, Services, or Mining-other-than-Oil. To illustrate this point, in Figure 2.14, we remove the Agriculture sector and observe that TFP remains stagnant during the 2000s. Thus, we can conclude that the Agriculture sector alone cannot account for the lack of TFP growth during this period.

FIGURE 2.14: Robustness: Net-of-Agriculture

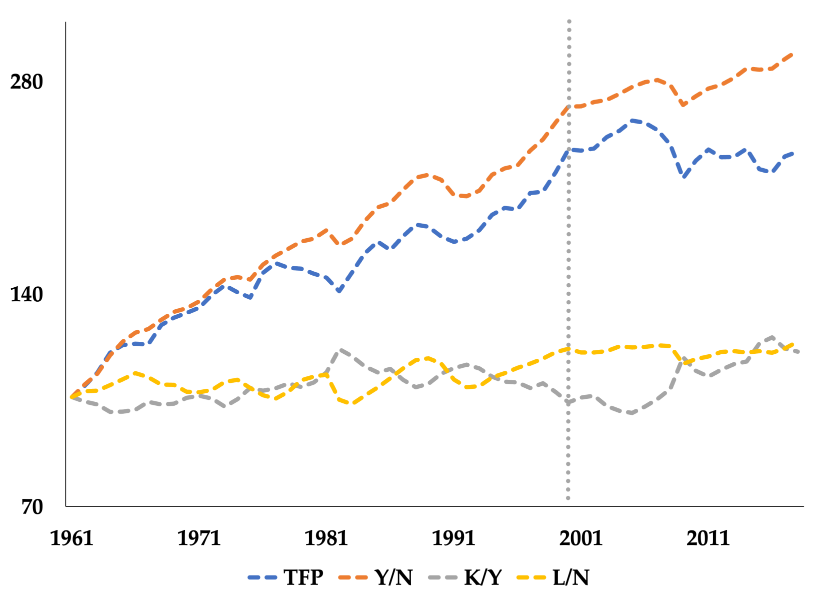

In Figure 2.15, we conduct a similar analysis, but this time we exclude the Manufacturing sector. As before, we arrive at the same outcome, namely, that TFP remains stagnant during the 2000s even after removing the Manufacturing sector from consideration.

FIGURE 2.15: Robustness: Net-of-Manufacturing

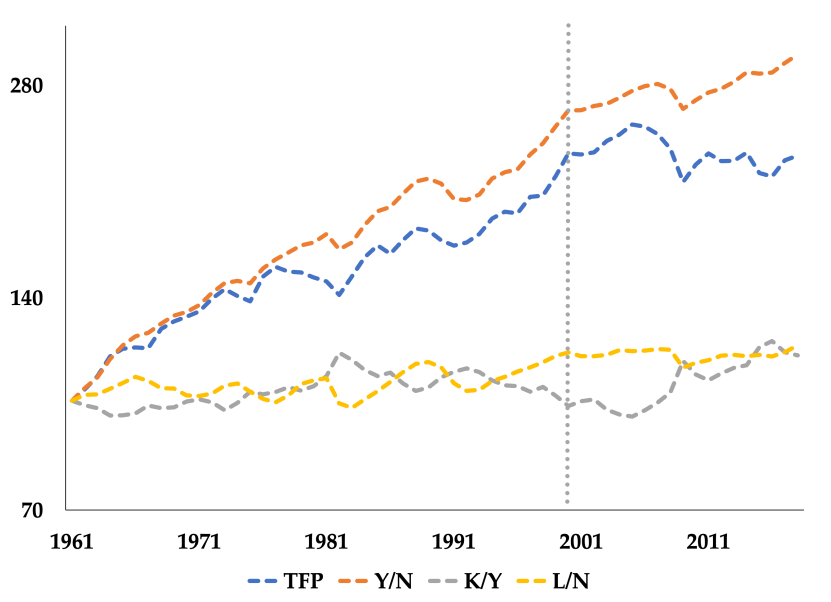

In Figure 2.16, we repeat the same analysis by excluding the Services sector. Although we reach the same conclusion that removing this sector alone cannot explain the stagnant

TFP, we obtain a much more volatile depiction. This is due to the fact that the Services sector constitutes a significant proportion of GDP, resulting in greater measurement error when it is excluded from the analysis. Additionally, the declining Hours per WAP line can be attributed to structural transformation, where the Services sector has grown significantly during this period. Removing it from consideration causes the total number of hours worked in the economy to fall mechanically.

FIGURE 2.16: Robustness: Net-of-Services

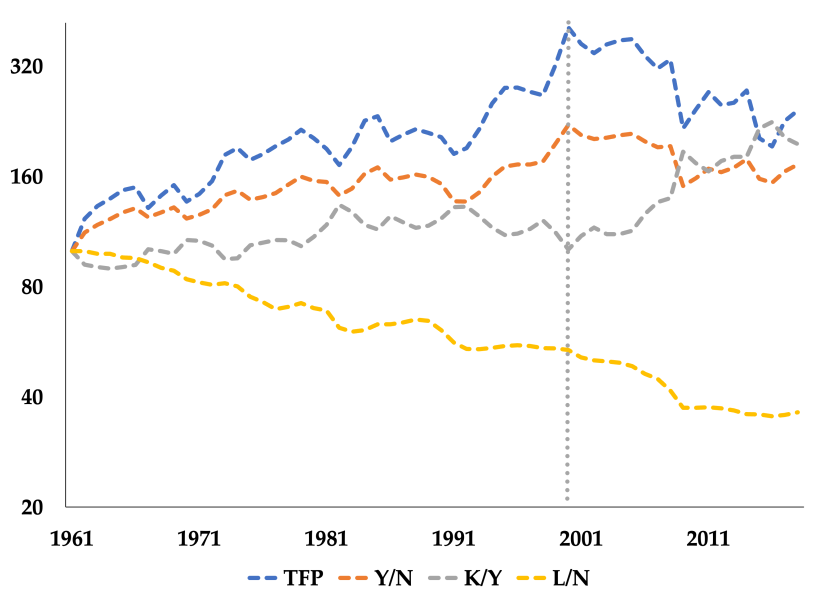

Finally, Figure 2.17 demonstrates that the remaining mining sector, apart from Oil, is also not accountable for the stagnant TFP.

C Adjusting for PPP in GAD exercise

In our analysis comparing Canada to the United States, we do not adjust for exchange rate fluctuations. It is well established that oil prices greatly affected Canadian exchange rates, particularly during the commodity boom in the 2000's. To verify the robustness of our results, we repeat the decomposition using the Penn World Table (Feenstra, Inklaar, and Timmer, 2015) values for Canadian and American Real GDP respectively. In particular, we use "Output-side real GDP at chained PPPs" which allows us to "compare relative productive capacity across countries and over time." We then recompute the aggregate GDP deflator by taking the ratio between aggregate GDP in current prices and the PWT values, and deflate all the monetary variables accordingly.

As Figures 2.18a and 2.18b show, adjusting for PPP does not meaningfully impact the results of the aggregate growth decomposition for Canada. When oil is included, the lack of TFP growth in the 2000's persists.

FIGURE 2.17: Robustness: Net-of-other-Mining

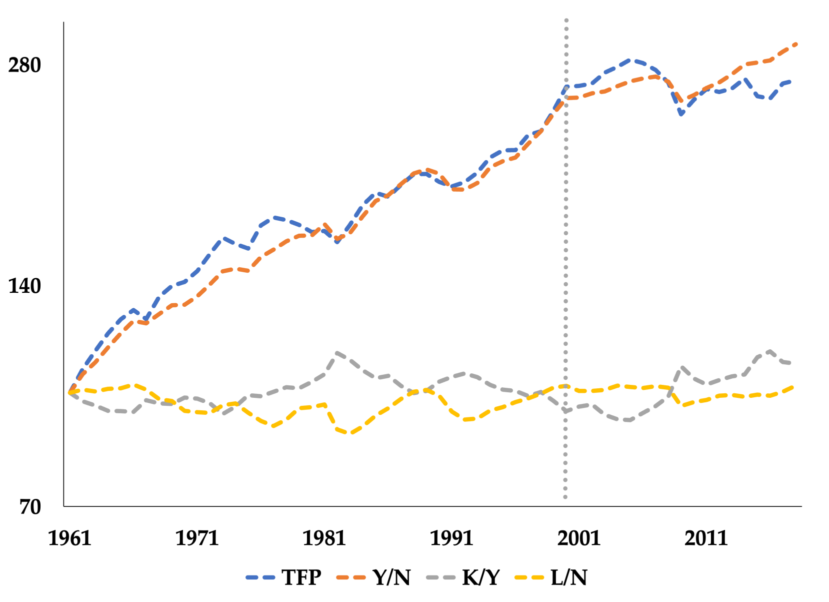

FIGURE 2.18: PPP adjusted Growth Accounting Decomposition

(A) **With Oil**

(B) **Without Oil**

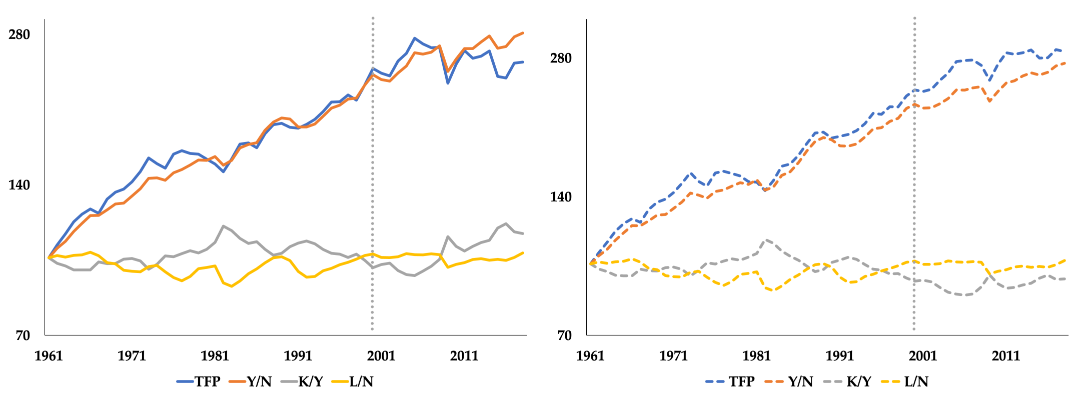

In Figures 2.19a and 2.19b, we repeat the comparison of the evolution of TFP growth between Canada and the United States using these PPP adjusted measures.

While the overall trend post-2010 of Net-of-Oil TFP growth is more subdued when accounting for PPP, the main result is strengthened. When we remove oil from the growth accounting Canadian TFP evolved comparably to the United States post-2001. In fact, coming out of the Great Recession, Canada's Net-of-Oil TFP grew significantly more than that of the United States before slowing down towards the end of our sample.

FIGURE 2.19: PPP adjusted TFP and Net-of-Oil TFP, Canada and United States

(A) **Since 1961** (B) **Since 2001**

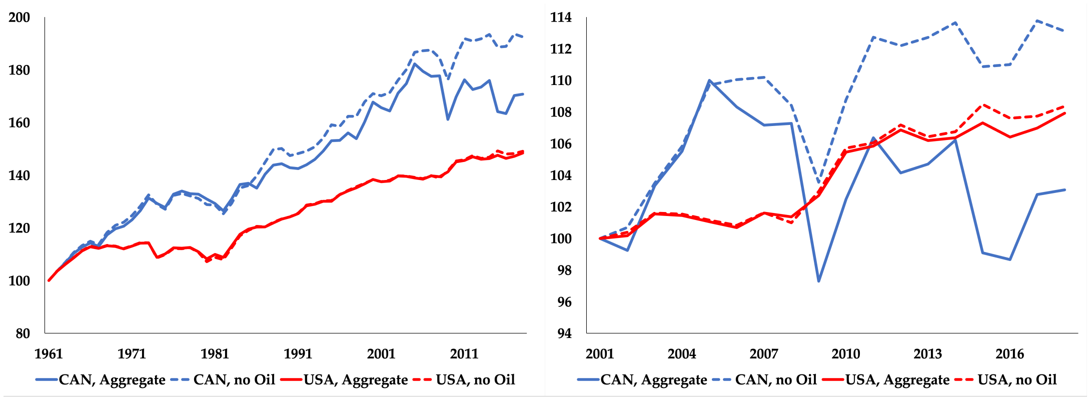

D Alternative α in GAD exercise

In our analysis of TFP performance between the whole economy and the Net-of-Oil economy, we use different values of α, which is the parameter that governs the capital share of the economy. It is natural to wonder if the lack of TFP growth being accounted for by netting out the Oil sector is due to the use of different parameter values. To address this question, we perform an alternative decomposition using the benchmark α value on the Net-of-Oil economy, and the resulting picture is very similar to the one in the main text. Figure 2.20 shows the decomposition.

FIGURE 2.20: Robustness: Alternative α

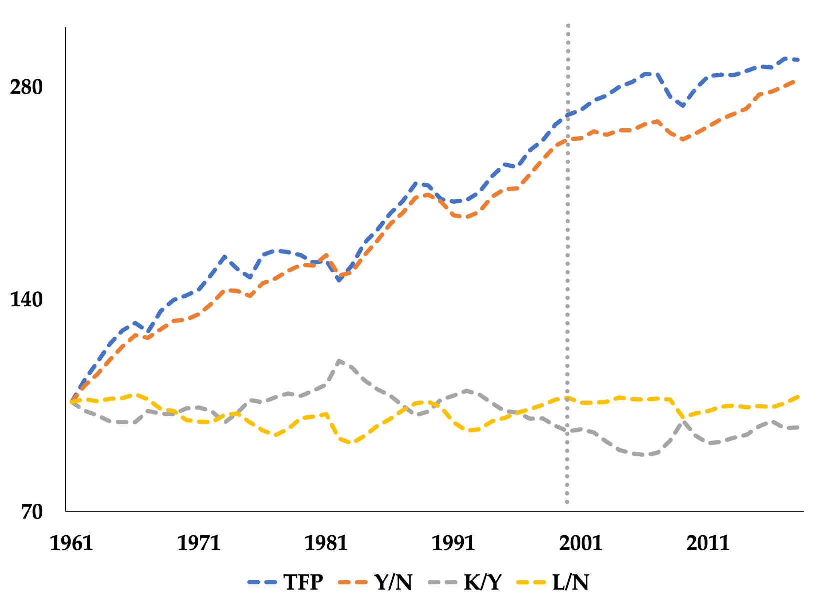

Another natural concern with our exercise has to do with the evolution of the labour share. As documented by Karabarbounis and Neiman (2014) and many others, the labour share of income may be declining in the data. In our exercise, the labour share is used to determine parameter α in the production function, which plays a crucial role in the measurement of TFP. In the benchmark exercise, this parameter is calculated using (one minus) the average of the series for the labour share. As a robustness check, we compute TFP using a different value of α in every year, ensuring that it is consistent with that year's labour share. We plot the results in Figures 2.21a and 2.21b.

FIGURE 2.21: Changing labour share GAD

(A) **With Oil** (B) **No Oil**

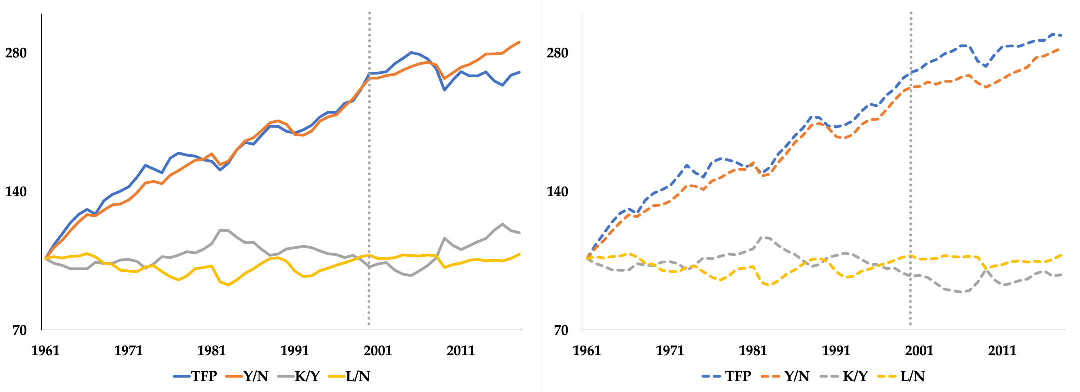

Our results show that the growth accounting exercises — with a flat TFP when oil is included and a growing TFP when oil is excluded — are very similar to those of a constant parameter α.

Therefore, given the two exercises in this Appendix, we conclude that the difference in TFP growth between the whole economy and the Net-of-Oil economy is not an artefact of the parameter value (or values) α used in our analysis.

Chapter 3

What are we deflating for?

What are we deflating for?

Abstract

Conclusions drawn from comparing the evolution of GDP per capita among developed economies relative to the United States can differ markedly depending on whether current-PPP or constant-PPP metrics are used. Constant prices are useful to compare a country's output over time; PPP measures are useful to compare two countries' output in a given period. However, constant-PPP metrics introduce a "triangular bias" that appears to favour the United States. Our numerical experiments demonstrate that different base currency-years can alter the qualitative interpretation of the relative performance of two countries. The existence of the triangular bias suggests using current price-PPP comparisons to compare the evolution of GDP per capita among developed countries.

3.1 Introduction

This paper first shows that, since the beginning of the 21st century, the economic performance of rich countries relative to the United States is grim when computing the ratio of real gross domestic product in purchasing power parity (GDP in PPP). The picture is noticeably less grim when we compute the ratio of nominal GDP in PPP instead. Note that the inflationary problem that arises with nominal values disappears when we calculate the ratio of nominal GDP in PPP, as the secular increase in price is cancelled out in both the numerator and the denominator. The noticeable discrepancy between the two ratios calls into question which metric best describes the relative evolution of two industrialized economies.

In this paper we argue that the relative evolution of two developed economies is better captured when using the ratio of nominal GDPs in PPP. Hence, the grim evolution of developed economies compared to the United States is more nuanced than usually described.

The comparison of economic activity for any two countries and any two time periods requires using a basket of prices and different quantities from each country and time period. When comparing the relative performance of a country over two different periods, we need to calculate the economic activity using the same basket of prices, which is referred to as "real GDP." When comparing the relative performance of two countries in a given period, we need to calculate economic activity using the same basket of prices, which is referred to as "purchasing power parity" (PPP). In both cases we use prices not observed in the period or place that the economic decisions took place, but they can be viewed as first approximations.

The problem arises when measuring the value of economic activity in a country in a given time period using the basket of prices of a different country and a different time period. This is the "Two Approximations Problem" (henceforth, TAP) which is at the heart for our findings.

Comparing the evolution of two countries' GDP using real GDP in PPP necessarily falls into the TAP. If, instead, we compare the evolution of two countries' GDP using nominal GDP in PPP we will be only doing one approximation. In a numerical illustration, we show how large the TAP can be. Moreover, we show that using the ratio of real GDP in PPP can paint a very different picture than using the ratio of nominal GDP in PPP.

In the data, the TAP appears to systematically overstate the relative performance of the US economy. We show that the relative decline in GDP per capita is consistently larger when real PPP measures are used as opposed to nominal PPPs. Given that the use of PPPs as a common base unit is, by construction, meant to eliminate differences in price and exchange rate fluctuations across countries, we argue in favour of using nominal measures when performing intertemporal cross-country comparisons.

To highlight the ways in which changes in the price base across time and countries can result in a misreading of data, we develop a simple numerical two-period model to compare two economies. We show that, for a given parametrization of the model, the choice in how we measure real GDP can produce drastically different interpretations. In the extreme case, we show that an economy measured in real (non-PPP) terms appears to grow, but contracts when measured in real PPPs. In subsequent experiments, we show that the choice of base currency, choice of base year, and choice of deflator all meaningfully influence the interpretation of the evolution of the two economies relative to one another, to the point that seemingly small changes in measurement choices completely alter the conclusions we would draw from observing the data.

The problems with using PPPs as a unit of measurement for conducting cross-country analysis is well documented, though most of the focus has been on issues in comparing poor countries. For the purpose of cross-country comparisons, the price index of a common basket of goods and services is constructed that adjusts for price disparities across countries. This PPP exchange rate is then used to arrive at a common unit or currency for the purpose of comparison (Lafrance and Schembri,2002). The challenges in the calculation of PPPs, particularly for poor countries, are well documented (for an overview, see Ravallion, 2020). For one, the construction of a representative basket is not always feasible across countries, particularly when the economic structure of the countries is not comparable (Deaton and Heston, 2010). Furthermore, revisions to the measurement of PPPs can highlight the extent to which past estimates of economic activity may be understated (Deaton and Aten, 2017). For example, methodological changes that place a higher weight on traded goods affects the comparability of PPP measures with domestic deflators (Ravallion, 2018). The challenges in obtaining measures of real economic activity using PPPs is not a new observation. Dowrick and Quiggin (1997) point out that making comparisons of GDP per capita using standard constant price measures results in significant bias. As a result, observed convergence of real GDP may actually be an observed convergence of prices, with the direction of the bias depending on if the price base is derived from a rich or a poor country. We expand on this by highlighting the issues that arise when using PPPs for comparing the performance of comparably developed countries where the price levels and economic structures are comparable. Using our numerical model, we show that different measures of real GDP can deliver opposite observations of the relative performance of two economies.

The flaws of using GDP as a metric of well-being and for cross-country analysis is also well documented. Early attempts at compiling estimates of real GDP across countries (such as Kravis et al., 1978) acknowledge significant inconsistencies in measurement of economic accounts across countries. For instance, self-reported measures of GDP from autocratic regimes may exaggerate economic performance (Martínez, 2022). Measurement in a single, developed economy such as the United States presents challenges in producing consistent measures of GDP over time and across methods of accounting, due to the need to combine data from various sources which in some cases may have been collected for entirely different purposes (for a thorough summary, see Landefeld et al., 2008). Researchers have also considered many approaches to addressing the fact that accounting approaches that in theory yield identical measures of GDP, such as expenditure and income, tend to produce statistical discrepancies (for examples, see Stone et al., 1942; Weale, 1992; and Aruoba et al., 2016). Chang and Li (2018) finds that re-estimating models using income qualitatively affects the results in a significant number of cases. Kohli (2004) argues that real GDP underestimates the growth in real income across several countries. However, real measures of income are not immune to the measurement problems when producing multiple approximations to perform cross-country comparisons. We highlight this by showing that the interpretation of the evolution of Net National Income (NNI) per capita is prone to the same or stronger distortions from the triangular bias. For example, the conclusions drawn from observing NNI per capita of Germany relative to the United States is entirely dependent on whether we focus on the real PPP or nominal PPP measure. In real terms, it would appear that the average German's income has fallen by 20 percentage points relative to the US, while in nominal terms one would conclude that the relative incomes are unchanged.

The rest of the paper proceeds as follows: Section 3.2 presents the evolution of GDP per capita across several industrialized nations to highlight how using real vs nominal measures in PPPs affects the cross-country comparisons. Section 3.3 expands on the intuition behind the triangular bias and demonstrates via a simple numerical example how the triangular bias distorts comparisons of GDP. Finally, Section 3.4 summarizes the paper and highlights areas for future research

3.2 Cross-country comparisons using PPPs

In this section, we highlight the differences in economic performance relative to the United States among the other G7 countries. We begin by highlighting how the evolution in GDP per capita when measured in real PPPs seems to paint an alarming picture before highlighting how this discrepancy is a result of our proposed "Two Approximations Problem." All data is taken from the Organization for Economic Co-operation and Development

(OECD) tables. In the OECD tables, the value of GDP with the domestic currency and in PPP is equivalent (only) for the United States. Because of this, we assume PPP values are using prices in the United States.

3.2.1 GDP per capita

In recent years there has been considerable concern over the performance of rich economies relative to the United States. Up to the end of the 20th century, the gap in real GDP per capita was constant. Starting around the year 2000, it seems that the United States started to pull away from their G7 counterparts. Figure 3.1 shows the evolution of GDP per capita in real PPPs relative to the United States. For each country, period between 2000 and 2010 marks the beginning of a divergence in the evolution of GDP per capita.

The gap varies across countries but is on the order of $10,000 to $20,000 a year (in PPP$2020). However, when comparing the evolution of GDP per capita relative to the United States in both nominal and real terms, a slightly different story emerges. Figure 3.2 plots the ratio of each countries GDP per capita relative to the United States in both real and nominal PPPs since 2000. On close inspection it appears that between 2000 and 2010, the trends of both real and nominal GDP were nearly parallel. In Canada's case, the distinction between real and nominal GDP is non-existent. Post 2010, with the exception of Japan, the decline in relative GDP per capita is starker in real than in nominal terms for each country. The differences from 2010 to 2023 are summarized in Table 3.1.

Across all countries, the difference in the observed decline in relative GDP per capita when using nominal vs real PPP measures is between 4 and 10 percentage points. In the case of Japan, real GDP understates the decline, while for all other countries it overstates the decline.

Table 3.1: Ratio of GDP per capita

	Nominal			Real			
	2010	2023	Difference	2010	2023	Difference	Difference of Differences
CAN	0.8256	0.7792	0.0463	0.8269	0.7248	0.1021	0.0558
DEU	0.8332	0.8246	0.0086	0.9520	0.8494	0.1026	0.0940
FRA	0.7425	0.7052	0.0373	0.8672	0.7607	0.1065	0.0692
GBR	0.7513	0.7043	0.0471	0.8425	0.7562	0.0863	0.0393
ITA	0.7202	0.6990	0.0212	0.8432	0.7213	0.1219	0.1007
JPN	0.7277	0.6042	0.1235	0.7068	0.6359	0.0708	-0.0526

3.2.2 Alternative Comparisons

The observed divergence between real and nominal economic activity is not strictly a result of using GDP to compare standards of living. While it remains the most common measure of economic activity used in economics, the literature has also argued in favour

FIGURE 3.1: GDP per capita in 2020 PPPs

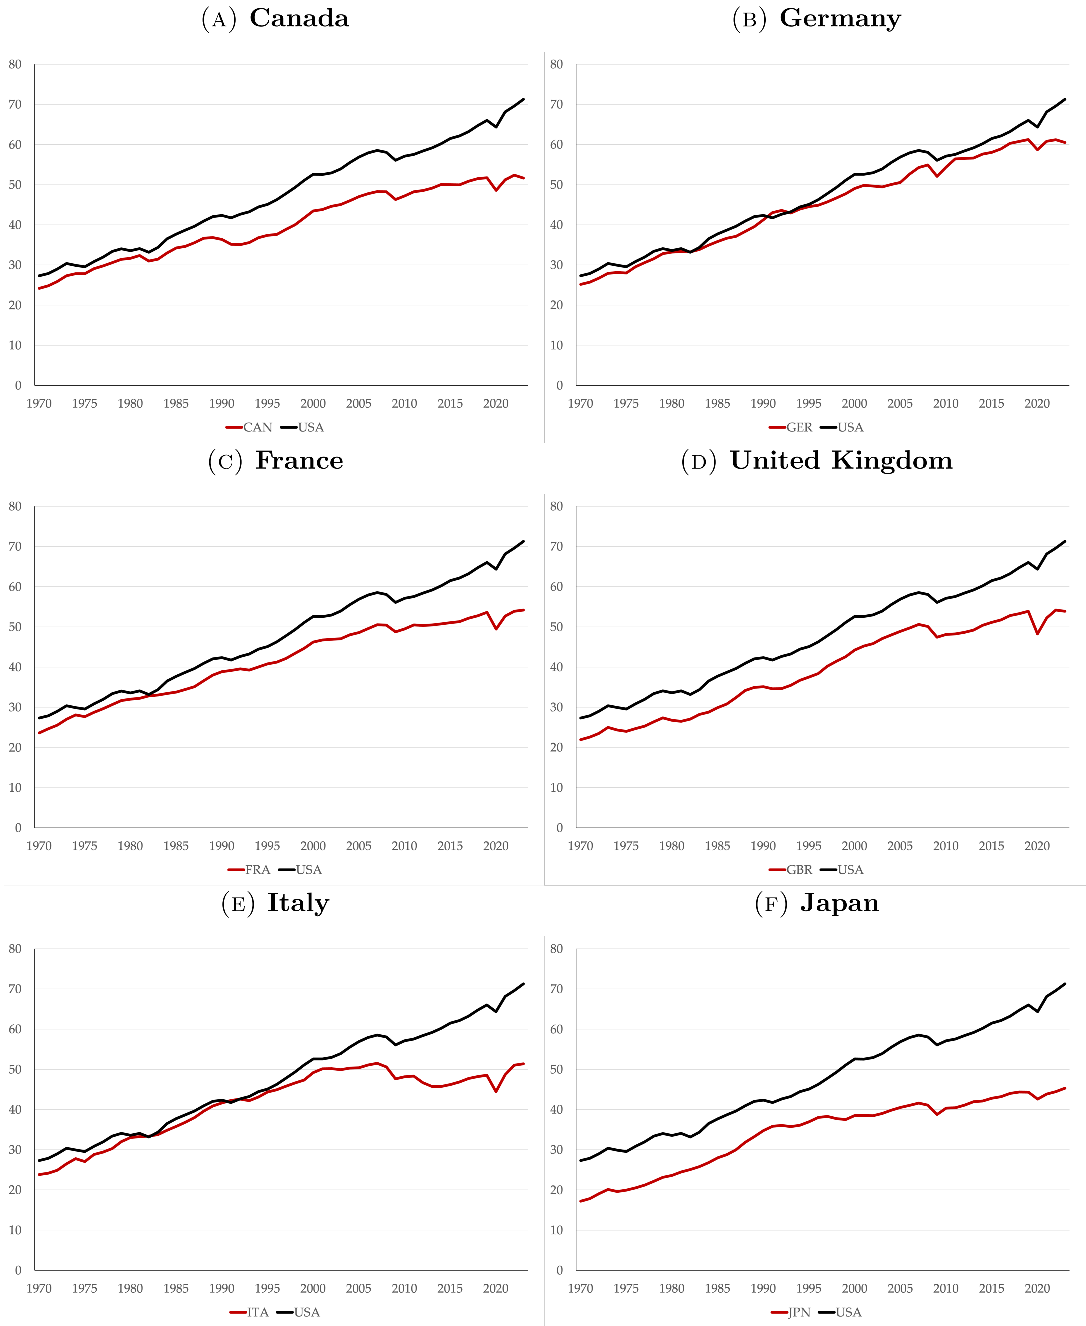

Source: OECD and authors calculations

of alternative metrics to compare living standards (Jones and Klenow, 2016; Dědeček and Dudzich, 2022). A number of researchers have argued in favour of using various measures of income to assess national well-being rather than (or in tandem with) GDP (see Macdonald, 2010; Chamberlin, 2011; Nolan et al., 2019), though not without controversy (see

FIGURE 3.2: Ratio of GDP per capita relative to the US

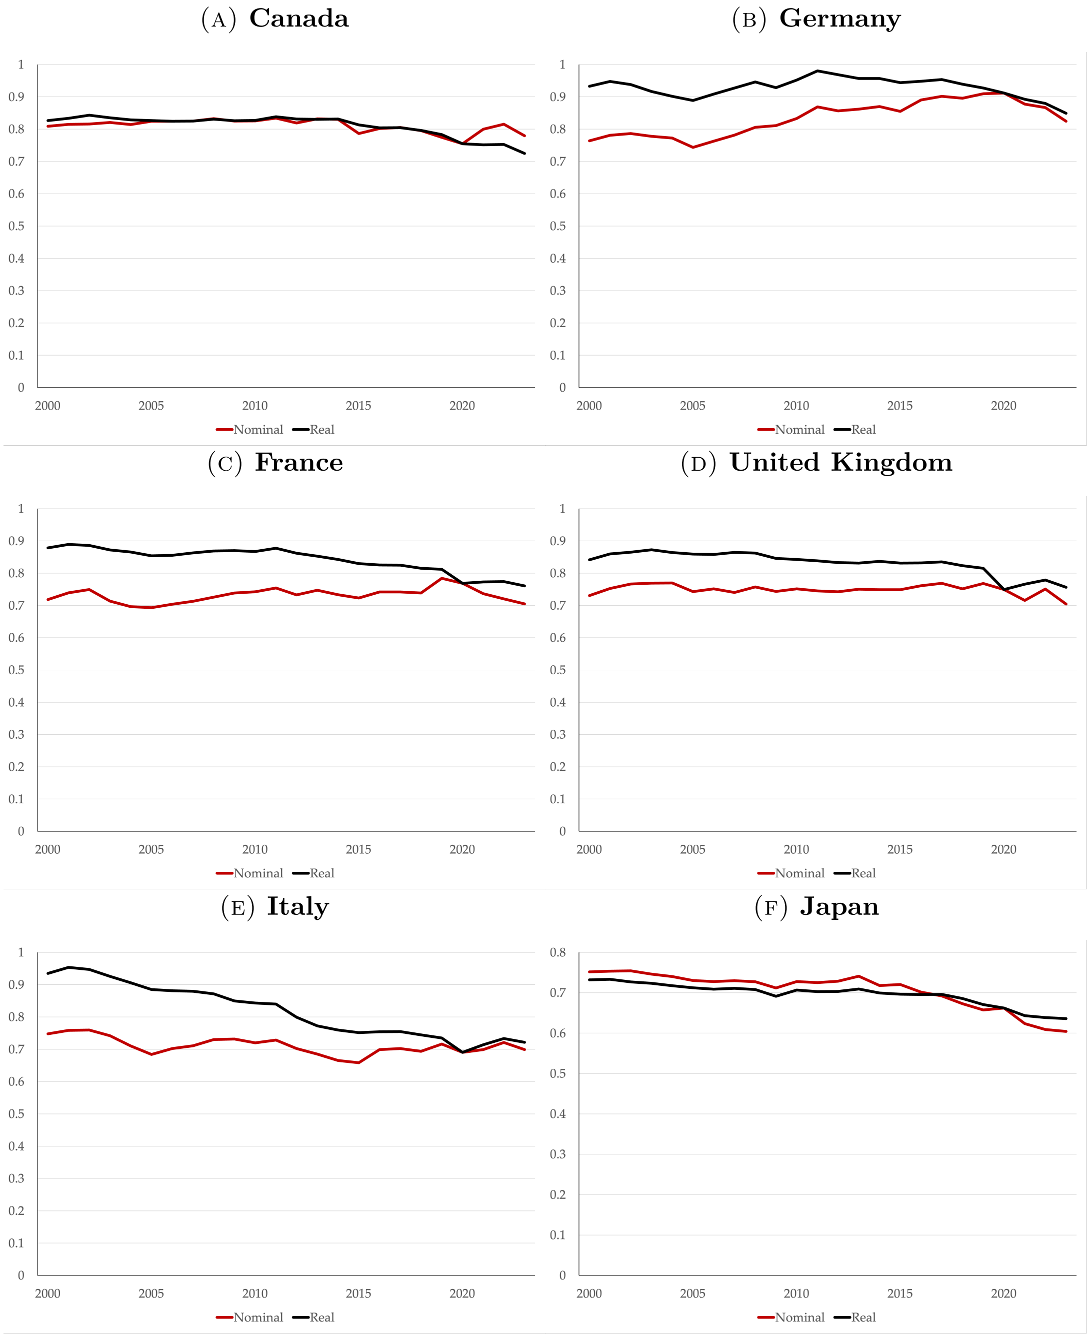

Source: OECD and authors calculations

Feldstein, 2017 for the argument against income as a measure of standards of living).

In our analysis, the same problem that distorts the measurement of real GDP across countries affects the measurement of national income as well. To further emphasize our observation we turn our attention to an alternative metric of economic well-being that is often measured in PPPs for cross-country comparisons: Net National Income (NNI). NNI is computed as gross national income net of depreciation of fixed capital assets. The OECD produces estimates of NNI per capita in both current (nominal) and constant (real) (2020) PPPs. Figure 3.3 shows the evolution of NNI per capita of the G7 countries relative to the United States in nominal and real terms since 2000.

FIGURE 3.3: Ratio of NNI per capita relative to the US

(A) **Canada**

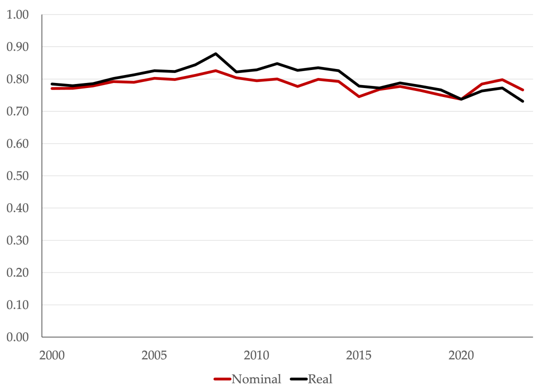

(B) **Germany**

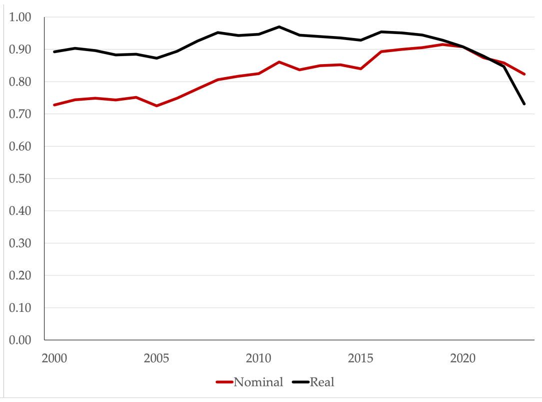

(C) **France**

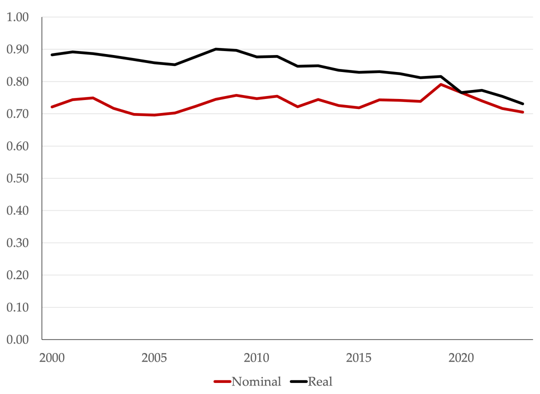

(D) **United Kingdom**

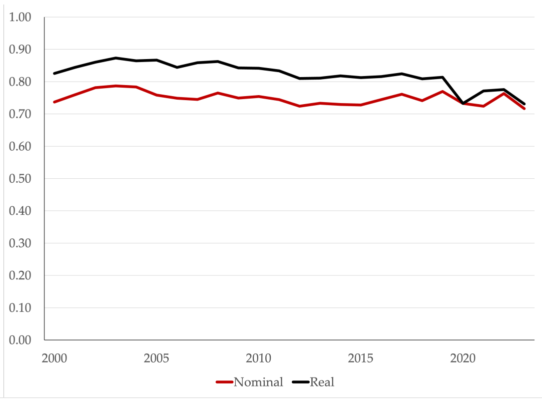

(E) **Italy**

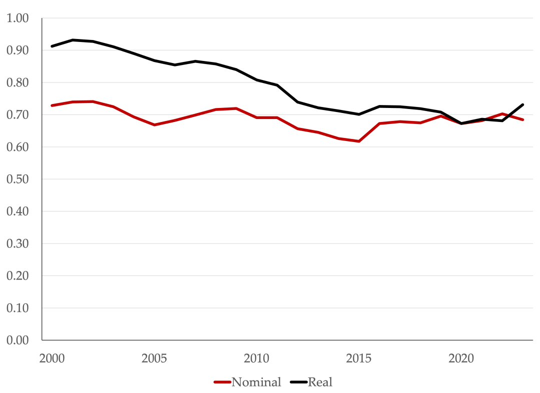

(F) **Japan**

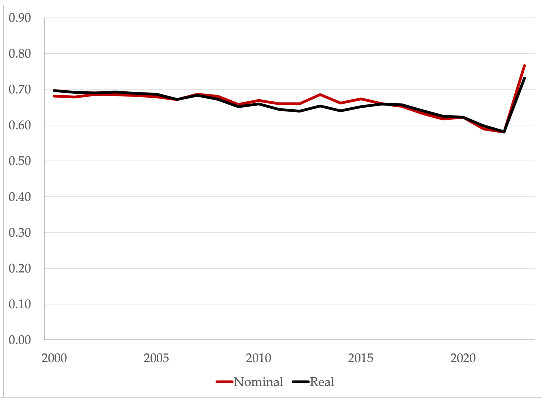

Source: OECD and authors calculations

The evolution of the ratios of national income paint a similar story as GDP for most of the G7 countries. With the exception of Canada and Japan, we can observe similar parallel trends between real and nominal measures between 2000 and 2010. After 2010 we observe a narrowing gap between the two measures. In the cases of Canada and Japan, the difference between real and nominal measures of income are less pronounced. However, the same story of differences in the observed decline in relative incomes between 2010 and 2023 emerges in Table 3.2.

TABLE 3.2: Ratio of Net National Income

	Nominal			Real			
	2010	2023	Difference	2010	2023	Difference	Difference of Differences
CAN	0.7950	0.7663	0.0287	0.8286	0.7311	0.0975	0.0688
DEU	0.8248	0.8233	0.0016	0.9464	0.7311	0.2153	0.2138
FRA	0.7471	0.7051	0.0420	0.8761	0.7311	0.1450	0.1030
GBR	0.7541	0.7165	0.0376	0.8413	0.7311	0.1102	0.0726
ITA	0.6907	0.6844	0.0063	0.8075	0.7311	0.0764	0.0701
JPN	0.6688	0.7663	0.0975	0.6594	0.7311	0.0711	-0.0258

The differences in the observed evolution of relative NNI between real and nominal measures in PPP are quantitatively substantial. In the case of Germany the observed difference is more than 21 percentage points: in real terms it would appear Germans have become significantly poorer relative to an American in the last 10 years, while in nominal terms it would appear German income relative to the US has been constant. From these results we can see that the choice of whether or not to further deflate two series that have already undergone an initial transformation to approximate a common price base has the effect of making an accurate comparison impossible. At best, the conclusions drawn from comparing any country to the US are exaggerated (as would be the case for Canada or France, for example), and at worst the conclusions may not reflect reality at all (as would be the case for Germany).

3.3 The Two Approximations Problem

In this section we explain our notion of the "Two Approximations Problem," the result of using two approximations to make an intertemporal comparison of two economies. We explain the intuition behind why this bias occurs before moving on to an illustrative, numerical example that highlights how the bias occurs.

3.3.1 Intuition

The cause of this discrepancy can be understood through a form of triangular bias. Consider a world consisting of two countries and two periods. Suppose country 1 is our representative country (the United States) and country 2 is our comparison case (Canada).

Each economy produces a basket of three goods: agriculture (A), manufacturing (M), and services (S), and for ease of illustration assume there is no trade. Let $p_j^i(t)$ denote the price in country i of good j in period t, and let $q_j^i(t)$ denote country i's production of good j in period t. Then nominal GDP in each country in period 1 would be given by

$$GDP^1(1) = p_A^1(1)q_A^1(1) + p_M^1(1)q_M^1(1) + p_S^1(1)q_S^1(1)$$
$$GDP^2(1) = p_A^2(1)q_A^2(1) + p_M^2(1)q_M^2(1) + p_S^2(1)q_S^2(1)$$

Suppose now we wanted to make a period 1 comparison of GDP in each country. In this first approximation, we would use our representative country's prices to measure country 2's GDP to arrive at

$$GDP_{PPP}^2(1) = p_A^1(1)q_A^2(1) + p_M^1(1)q_M^2(1) + p_S^1(1)q_S^2(1)$$

In this step we are already introducing some measurement issues by imposing non-equilibrium prices on country 2 and holding quantities fixed.

Next, suppose we wanted to compare the GDP of country 2 in period 2 to that of period 1. There are a number of ways to perform this comparison by either holding as fixed period 1 prices, period 1 quantities, or an average of the two. Suppose we measured real GDP in country 2 in period 2 by holding period 1 prices as fixed. Then we would arrive at

$$GDP_{real}^2(2) = p_A^2(1)q_A^2(2) + p_M^2(1)q_M^2(2) + p_S^2(1)q_S^2(2)$$

Similar to the cross-country comparison, we are imposing period 1 equilibrium prices on period 2 equilibrium quantities.

The problem then compounds when we try to do both simultaneously. Suppose we wanted to make a cross-country, intertemporal comparison of GDP in both countries, using country 1's prices from period 1. Then country 2's GDP in period 2 would be measured as

$$GDP_{real\ PPP}^2(2) = p_A^1(1)q_A^2(2) + p_M^1(1)q_M^2(2) + p_S^1(1)q_S^2(2)$$

The first approximation (converting to PPPs) is meant to eliminate differences in price fluctuations and distortions from nominal exchange rates in each period. To some extent, this is already a "real" measure of comparison for the two economies in question. The second approximation (using constant prices) is meant to eliminate the effect of price fluctuations over time *within* an economy and as a result is the appropriate metric to use when studying the evolution of a single economy intertemporally. However, the effect of using two approximations to compare two economies over time distorts the picture and leads us from making something approximating an apples-to-apples comparison (using PPPs) to an apples-to-bananas comparison (using real PPPs).

3.3.2 Numerical example

To illustrate how making cross-country comparisons using real PPPs can lead to drawing incorrect conclusions, we perform a numerical exercise. Using a simple model, we demonstrate that making intertemporal cross-country comparisons can yield dramatically different results depending on the choices we make in measuring the path of the economy. For simplicity, we develop a two period model of an economy consisting of a representative household and three production sectors. The household is endowed with one unit of labour that it supplies inelastically across sectors, so that in equilibrium the wage across sectors is equalized. For ease of interpretation of the evolution of prices in this economy, we normalize the wage to 1.

The household chooses consumption of three goods, $\{c_A, c_M, c_S\}$ by solving

$$\begin{aligned} &\max \left(\mu_A \left(c_A - \bar{c_A}\right)^{\rho} + \mu_M \left(c_M\right)^{\rho} + \mu_S \left(c_S + \bar{c_S}\right)^{\rho}\right)^{\frac{1}{\rho}} \\ \text{subject to: } & p_A c_A + p_M c_M + p_S c_S = w \equiv 1 \end{aligned}$$

with $\mu_S = 1 - \mu_A - \mu_M$.

This problem yields the following two conditions from the first order conditions:

$$\begin{aligned} \frac{p_A}{p_M} &= \frac{\mu_A}{\mu_M} \left(\frac{(c_A - \bar{c_A})}{c_M} \right)^{\rho-1} \\ \frac{p_A}{p_S} &= \frac{\mu_A}{\mu_S} \left(\frac{(c_A - \bar{c_A})}{(c_S + \bar{c_S})} \right)^{\rho-1} \end{aligned}$$

The only factor of production in this economy is labour, and the only distinction across sectors is the level of productivity Z_i, so that the production technology is described by

$$c_i = Z_i l_i$$

for $i = \{a, m, s\}$.

Perfect competition in labour markets implies that in equilibrium, the price for the output of each of the three sectors is

$$p_i = 1/Z_i$$

for $i = \{a, m, s\}$.

Given a set of parameters for ρ, c_A, c_S, μ_A and μ_M, different productivities Z_i will yield both different allocations and different prices across sectors within an economy. To highlight this, we select the following parameters for two countries, called US and CAN. We set $\rho = 0.5$, $c_A = 0.01$, $c_S = 0.01$, $\mu_A = 0.3$ and $\mu_M = 0.2$, and solve for the consumption allocations using the following set of productivities:

Table 3.3: Productivities

	Period 1			Period 2		
	Z_A	Z_M	Z_C	Z_A	Z_M	Z_C
US	1	1	1	1.5	0.8	1
CA	2	0.5	1	0.5	2	1

The solution to the model yields the following consumption allocations:

Table 3.4: Consumption allocations

	Period 1			Period 2		
	c_A	c_M	c_C	c_A	c_M	c_C
US	0.2268	0.1053	0.6679	0.4740	0.0612	0.6075
CAN	0.7860	0.0021	0.5628	0.0506	0.4309	0.6833

Since the wage w is normalized, it is always the case that

$$GDP = p_A c_A + p_M c_M + p_S c_S = 1$$

so

$$GDP_t^X(p_t^X) = 1$$

for the equilibrium vector of prices $p_t^X = (p_A^X(t), p_M^X(t), p_S^X(t))$, in each country $X = \{US, CAN\}$ and each period $t = \{1, 2\}$.

We interpret values in PPP to mean that the price levels are from the US, and we interpret "real" as keeping period $t = 1$ prices as fixed. Table 3.5 summarizes the evolution of GDP in each country across three measures: real GDP using domestic prices, nominal PPP (using US prices), real PPP (using US prices for period 1)

Table 3.5: GDPs, United States is PPP, period 1 prices

	Real GDP		Nominal PPP		Real PPP	
	US	CAN	US	CAN	US	CAN
Period 1	1	1	1	1.3709	1	1.3709
Period 2	1.1427	1.5705	1	1.2557	1.1427	1.1649

We begin by considering the evolution of the two economies independently. If we measure real GDP using period 1 prices in each country, we find that US GDP grows by

14.27 percentage points (pp) while CAN GDP grows by 57.05 pp. However, if we compare the two economies using the real PPP measure, CAN GDP falls by 20.60 pp. This encapsulates the immense challenge in making cross-country comparisons. It is not simply that this method of approximating the "real economy" introduces some measurement error but still tells a qualitatively consistent story, but rather that the bias introduced is so strong it completely reverses the story. Similarly, if we consider the evolution of the CAN/US ratio, we get two slightly different pictures. The comparison in nominal PPP terms suggests that CAN GDP relative to US GDP fell from 1.3709 to 1.0194, a decrease of 35.15pp. The comparison in nominal PPP terms suggests that CAN GDP relative to US GDP fell from 1.3709 to 1.2557, a decrease of 11.52pp. The implied difference, a spread of 23.63pp, illustrates the chasm between the conclusions drawn: an economist looking at the nominal PPP data sees bad news, an economist focusing on the real PPP data sees a catastrophe, and an economist looking at the two economies separately sees encouraging signs.

Choice of base prices

Suppose now that instead of using US prices as our measure of PPP we choose to use CAN prices. After re-computing the values, the evolution of GDP across measures between the two countries is summarized in Table 3.6. If we were to compute the CAN/US ratios

TABLE 3.6: GDPs, Canada is PPP, period 1 prices

	Real GDP		Nominal PPP		Real PPP	
	US	CAN	US	CAN	US	CAN
Period 1	1	1	0.9918	1	0.9918	1
Period 2	1.1427	1.5705	1.5861	1	0.9669	1.5705

using this set of data, we would find that relative real GDP in PPPs grew from 1.0083 to 1.6243, an increase of 61.60pp, while in nominal terms it went from 1.0083 to 0.6305, a decline of 37.78pp. Note that the only change made is which economy's prices we fixed as our standard, and yet we observe a completely different story in terms of the relative performance of each economy, particularly in real terms. In the previous example, we found that in real PPPs the CAN economy contracted relative to the US, while now we find that it actually increased. This suggests that the measurement errors are not solely the result of which metric of real or nominal we choose, but that the choice of our reference country will also qualitatively change the results.

Choice of base year

Suppose we repeated the previous example but instead used period 2 as our reference year. The resulting measures of GDP are summarized in Table 3.7. Just by switching the reference year we get an entirely new picture of the evolution of the two economies. When considered separately, real GDP in US grew by 5.19pp and fell in CAN by 54.40pp. If we look at real GDP in PPPs instead, we would conclude that GDP in CAN grew by

Table 3.7: GDPs, using period 2 prices

	Real GDP		Nominal PPP		Real PPP	
	US	CAN	US	CAN	US	CAN
Period 1	0.9507	2.1458	1	1.3709	0.9507	1.1144
Period 2	1	1	1	1.2557	1	1.2557

12.68pp. Given this measurement approach, we also get that the CAN/US ratio in real PPPs goes from 1.1722 to 1.2557, an increase of 8.35pp. In this scenario, the only change we made was to fix our reference prices to period 2 instead of period 1. This alone is enough to arrive at a dramatically different view of the evolution of the two economies over time.

Chain Volume Indices

In practice, the transformation from nominal to real GDP is often performed using a chain volume index. The OECD makes their transformation using the Chain Fisher volume index. We can compute this same index in the model using the following steps. First, we compute the Laspeyres quantity (LQ) index by imposing prices of a fixed year on quantities for period t and taking the ratio to our base year. If we fix period $t = 1$ prices, then for any period t, the LQ index is

$$LQ(t) = \frac{p_A(1)q_A(t) + p_M(1)q_M(t) + p_S(1)q_M(s)}{p_A(1)q_A(1) + p_M(1)q_M(1) + p_S(1)q_S(1)}$$

Next, we compute the Paasche quantity (PQ) index by taking the ratio of GDP in period t to a measure of GDP using period t prices and base period quantities. Again, using $t = 1$ as our base period, this means the PQ index in any period t is

$$PQ(t) = \frac{p_A(t)q_A(t) + p_M(t)q_M(t) + p_S(t)q_S(t)}{p_A(t)q_A(1) + p_M(t)q_M(1) + p_S(t)q_S(t)}$$

Finally, we compute the Fisher index FQ by taking the geometric mean of LQ and PQ so that

$$FQ(t) = \sqrt{LQ(t) \times PQ(t)}$$

Real GDP in period t ($GDP_{real}(t)$) is then derived by multiplying GDP in the base year ($t = 1$) by $FQ(t)$. So in this method

$$GDP_{real}(1) = GDP(1) \times FQ(1) = GDP(1)$$
$$GDP_{real}(2) = GDP(1) \times FQ(2)$$

There are a number of ways we can proceed to compute real GDP in this case. Suppose we started by looking at each economy independently and computing the Fisher Index.

Given our set of parameters, prices, and quantities, we would obtain the following values

Table 3.8: Fisher Indices, using USA

	US	CAN
$FQ(1)$	1	1
$FQ(2)$	1.0963	0.8555

Using these, we can compute real GDP (in home prices) and real GDP in PPPs. To compute real GDP in PPPs for CAN, we use GDP computed in nominal PPPs (by using US prices with CAN quantities) and then multiply period 1 GDP by the Fisher Index for CAN. The results are summarized in Table 3.9 If we focus our analysis on CAN: real GDP

Table 3.9: Real GDP in US-Fisher Indices

	Real GDP		Real PPP	
	US	CAN	US	CAN
Period 1	1	1	1	1.3709
Period 2	1.0963	0.8555	1.0963	1.1728

in home prices fell by 14.45pp, while real GDP in PPPs fell by 19.81 pp. The difference in the two measures is 5.36pp. Recall that the CAN/US ratio in nominal PPPs declined by 11.52pp. In real terms (using the Fisher Quantity Index), real GDP appears to have fallen by 30.11pp (a difference of 18.59pp).

Suppose now instead of using the Fisher index derived from home prices for CAN we use the Fisher Index derived from using PPPs (US prices). FQ for US is unchanged, while for CAN the new values are summarized in Table 3.10 The new measure of real

Table 3.10: Fisher Indices, Canada prices

	CAN
$FQ(1)$	1
$FQ(2)$	0.9785

GDP in PPPs is summarized in Table 3.11 In this method, real GDP in PPPs falls by

Table 3.11: Real GDP in Canada-Fisher Indices

	Real PPP
	CAN
Period 1	1.3709
Period 2	1.3414

2.95pp. The CAN/US ratio in real PPPs goes from 1.3709 to 1.2235, a fall of 14.73pp.

Across these exercise we have shown that the interpretation of the evolution of two economies is entirely dependent on the choices in measurement. We have demonstrated that across different measures of real economic activity we can get diametrically opposing results, that the choice of base year will also completely change our measurement both quantitatively and qualitatively, and that the approximations used by practitioners to convert nominal to real will also significantly alter the quantitative results.

What this highlights is that the choice of base year and base price base dramatically affect the interpretation of how the economy behaves: from one lens the “Canadian” economy grows in real terms while from another lens it actually contracts. This is dependent both on the choice of base year and which measure of real economic activity we choose to focus on. These choices also translate to the comparison between real and nominal values for making a cross-country comparisons. Our interpretation of this simple numerical example is to say one can derive whatever result they like by choosing the appropriate year and set of base prices to make their assessment. Ultimately this renders any cross-country intertemporal comparison meaningless as any interpretation is a function of the choice of deflator used.

3.4 Concluding Remarks

In this paper we present evidence that the observation that GDP per capita among G7 countries has been lagging behind the United States at an alarming rate in the last 10 years is possibly a mirage created by the choice of how we measure real economic activity. We show that for each G7 country, the conclusions one derives from observing relative GDP per capita depends on whether we focus on the ratio of nominal or real measurements. We then propose that any attempt to make a cross-country intertemporal comparison is flawed due to what we term the “Two Approximations Problem:” the idea that in order to arrive at a metric that eliminates price fluctuations across time and both across **and** within countries necessarily biases one country in favour of another. We provide further evidence by showing that the same distortions occur when we compare alternative measures of economic well-being, such as Net National Income.

Using a simple numerical model we show how the TAP can distort the conclusions we derive from observing two economies in two periods. We show that two measures of real economic activity using different price bases deliver opposite results for both within country analyses and deliver quantitatively different results for cross-country analyses. This highlights the continued flaws in performing cross-country comparisons. We reinforce this idea by examining how the measurement of the two economies changes when we alter our base year, our country of reference, and how we approximate a measure of real economic activity. We repeat our initial exercise first by switching which country’s prices we were adopting as our measure of PPP, and second by using a different base year to measure real economic activity. In both scenarios we find that the comparison of real GDP per PPP reversed from the initial exercise. Next, we change the base year and once again find that this choice dramatically altered the conclusions a researcher would derive from observing the data. Finally, we compute the Fisher Quantity indices for each economy under two different approaches and find that while this approach produces qualitatively similar results to our initial exercise, it significantly affects the magnitudes.

While the literature on the use of PPPs as a base unit has already highlighted the problems with comparing poor countries, there is a deeper problem that is only obvious when making the comparison among seemingly similar, developed countries. Fixing a base year prices within a country to evaluate intertemporal economic activity is straightforward, and using a common base unit to compare countries in a given year makes intuitive sense. However, the path towards arriving at a metric and unit of real economic activity that delivers a truly “apples-to-apples” comparison continues to elude economists and requires further research.

In the absence of such a metric, it is our view that cross-country comparisons using PPPs are better done on a nominal basis. Given that the purpose of PPPs is to arrive at an approximation of a “real exchange rate” that normalizes the price of a comparable basket, the additional step of further deflating to eliminate price changes within each economy can only distort the picture. As we point out in Section 3.2, the performance of G7 economies relative to the United States is truly alarming only when the comparison is made using “real” PPPs. The nominal PPP comparison still delivers the result that each economy is falling behind, but quantitatively the difference is cut in half.

References

Aruoba, S. B., Diebold, F. X., Nalewaik, J., Schorfheide, F., and Song, D. (2016). ‘Improving GDP measurement: A measurement-error perspective’. Journal of Econometrics, volume 191, no. 2, 384–397.

Chamberlin, G. (2011). ‘Gross domestic product, real income and economic welfare’. Economic & Labour Market Review, volume 5, no. 5, 5–25.

Chang, A. C. and Li, P. (2018). ‘Measurement Error in Macroeconomic Data and Economics Research: Data Revisions, Gross Domestic Product, and Gross Domestic Income’. Economic Inquiry, volume 56, no. 3, 1846–1869.

Deaton, A. and Aten, B. (2017). ‘Trying to Understand the PPPs in ICP 2011: Why Are the Results So Different?’ American Economic Journal: Macroeconomics, volume 9, no. 1, 243–264.

Deaton, A. and Heston, A. (2010). ‘Understanding PPPs and PPP-Based National Accounts’. American Economic Journal: Macroeconomics, volume 2, no. 4, 1–35.

Dowrick, S. and Quiggin, J. (1997). ‘True Measures of GDP and Convergence’. The American Economic Review, volume 87, no. 1, 41–64.

Dědeček, R. and Dudzich, V. (2022). 'Exploring the limitations of GDP per capita as an indicator of economic development: a cross-country perspective'. Review of Economic Perspectives, volume 22, no. 3, 193–217.

Feldstein, M. (2017). 'Underestimating the Real Growth of GDP, Personal Income, and Productivity'. Journal of Economic Perspectives, volume 31, no. 2, 145–164.

Jones, C. I. and Klenow, P. J. (2016). 'Beyond GDP? Welfare across Countries and Time'. American Economic Review, volume 106, no. 9, 2426–2457.

Kohli, U. (2004). 'Real GDP, real domestic income, and terms-of-trade changes'. Journal of International Economics, volume 62, no. 1, 83–106.

Kravis, I. B., Heston, A. W., and Summers, R. (1978). 'Real GDP Per Capita for More Than One Hundred Countries'. The Economic Journal, volume 88, no. 350, 215–242.

Lafrance, R. and Schembri, L. (2002). 'Purchasing-Power Parity: Definition, Measurement, and Interpretation'.

Landefeld, J. S., Seskin, E. P., and Fraumeni, B. M. (2008). 'Taking the Pulse of the Economy: Measuring GDP'. Journal of Economic Perspectives, volume 22, no. 2, 193–216.

Macdonald, R. (2010). 'Real Gross Domestic Income, Relative Prices, and Economic Performance Across the Oecd'. Review of Income and Wealth, volume 56, no. 3, 498–518.

Martínez, L. R. (2022). 'How Much Should We Trust the Dictator's GDP Growth Estimates?' Journal of Political Economy, volume 130, no. 10, 2731–2769.

Nolan, B., Roser, M., and Thewissen, S. (2019). 'GDP Per Capita Versus Median Household Income: What Gives Rise to the Divergence Over Time and how does this Vary Across OECD Countries?' Review of Income and Wealth, volume 65, no. 3, 465–494.

Ravallion, M. (2018). 'An exploration of the changes in the international comparison program's global economic landscape'. World Development, volume 105, 201–216.

Ravallion, M. (2020). 'On Measuring Global Poverty'. Annual Review of Economics, volume 12, no. Volume 12, 2020, 167–188.

Stone, R., Champernowne, D. G., and Meade, J. E. (1942). 'The Precision of National Income Estimates'. The Review of Economic Studies, volume 9, no. 2, 111–125.

Weale, M. (1992). 'Estimation of data measured with error and subject to linear restrictions'. Journal of Applied Econometrics, volume 7, no. 2, 167–174.

Conclusion

This thesis pushes the frontier on three important areas of the economics literature: how the gains and losses of a global transition to a Net Zero world are distributed across households in resource-rich economies; how the expansion of individual industries can distort measurement of aggregate productivity; and how comparing real economic activity across countries can result in contradictory conclusions being drawn.

In Chapter 1, I develop a quantitative model calibrated to Canadian data to first assess how the gains in lifetime consumption due to the boom in oil prices between 1997 and 2020 were distributed across the age, income, and regional distributions. I find that the highest gains were among young, low-income households living in the fossil fuel region, whose lifetime welfare increased by 15%. The increase in oil prices produced both stronger wage growth in the region, and higher profits distributed back to households, which was particularly beneficial for the young who benefit from higher lifetime labour earnings, and the low-income households who particularly benefit from these transfers.

In a forward looking exercise, I forecast the impact of the transition to a Net Zero world where the global demand for fossil fuels falls by half, and domestic clean energy production expands to account for 90% of total energy inputs. I find that while the fall in demand for fossil fuels reduces aggregate welfare, the expansion of the clean energy sector can reduce these losses across the distribution of households by between 15 to 54%. The primary reason for this is as the relative price for clean energy falls, the aggregate energy input bundle becomes less expensive. This induces an increased demand for non-energy inputs, creating a channel for the expansion to spillover across regions. However, the strongest welfare gains due to the growth in clean energy production are at the top of the income distribution.

These results imply two important policy channels that merit further study. First, these results are suggestive of a role for redistributive policies due to the concentration of welfare gains at the upper end of the income distribution. Research exploring the optimal design of a tax and transfer policy that compensates the households left worse off from the transition would produce valuable insights for policy makers. Further research of labour market responses and the frictions for the workers displaced by the transition may also carry implications for the optimal policy response. Second, studying how policies aimed at enhancing productivity across the economy can generate the necessary growth in the rest of the economy to effectively negate the contraction of the fossil fuel sector. My paper shows how productivity growth in one sector can spill over across sectors and regions to dampen the losses experienced in the contraction of another sector. Studying these channels in more detail, and the existing barriers impeding productivity growth would produce richer insights with significant policy implications.

In Chapter 2, Pau Pujolas and I deliver new insights into the stagnation in Canadian TFP growth. We find that the entirety of the slowdown in productivity growth can be accounted for by the growth of the oil sector. This is a result of the massive capital investment that occurred during the early 2000s, which effectively lowers measured productivity. To demonstrate this, we show that when we remove the oil sector from the measurement of the aggregate economy, the productivity stagnation disappears.

We reinforce these results by studying the evolution of TFP in Alberta and Ontario. We show how the Net-of-Oil measures of TFP completely eliminate an observed decline in Alberta's aggregate productivity, while having no impact on Ontario. In a series of cross-country comparisons, we show that Net-of-Oil TFP in Canada grew at comparable rates to the United States since 2001, and that we observe similar reversals in measured TFP in Norway. The key feature of our findings lies in the comparison with the United States. In both Canada and the U.S., the oil sector accounts for a constant fraction of GDP over time. However, in Canada the percentage of the aggregate capital stock in the oil sector exploded from approximately 5% to more than 30% This disproportionate growth in the capital allocated to the sector without an observed growth in the share of aggregate GDP accounted for by the sector is what generates the drag on aggregate productivity.

Future research considering whether this period is a case of capital misallocation is an important avenue to better understand the evolution of the Canadian economy since 2001. Similarly, further study of the observed post-Covid productivity decline and oil sector dynamics is warranted. Finally, a more comprehensive theory of how the stock of natural resources ought to be considered in the national accounts would produce massive new insights into the performance of the Canadian economy in the last two decades.

In Chapter 3, Pau Pujolas and I document the discrepancy in conclusions drawn from studying the evolution of GDP per capita in G7 countries relative to the United States in current- and constant-PPPs. We highlight that, apart from Japan, the conversion from nominal to real valued PPPs tends to result in significantly larger declines in the ratio of GDP per capita across each country. This result similarly holds for alternative measures of economic well-being, such as Net National Income. Using a simple numerical model, we show that making cross-country, intertemporal comparisons using real metrics is fraught with measurement issues. The results are strongly influenced by the base set of prices and base year chosen and can lead to contradictory conclusions being drawn. We attribute this to being the result of performing two simultaneous approximations to arrive at constant price measures of economic activity in a common currency. We argue in favour of using nominal PPP measures when performing cross-country comparisons.

The three papers in this thesis, while touching on three separate questions, form a cohesive body of work pushing forward the frontier of the macroeconomics literature. The thesis as a whole addresses questions of economic welfare embodied in consumption, productivity, and income, using cutting edge quantitative models and deeper insights of the national accounts and how the measurement of the evolution of macroeconomic aggregates ultimately influences the conclusions we arrive at as researchers. The first paper studies how changes in global demand for fossil fuels and productivity growth in clean energy production influence lifetime consumption across heterogeneous households. The second paper further studies how the economic incentives driven by rising oil prices distorted our understanding of Canadian productivity growth. The third paper further enhances our understanding of the measurement of macroeconomic data to highlight the challenges in making comparisons of economic well-being across economies, even when they are comparable in their level of economic development.