



THE LONDON SCHOOL  
OF ECONOMICS AND  
POLITICAL SCIENCE ■

# Causal Inference in Spatial Environmental Economics

**Lorenzo Sileci**

LSE, Department of Geography and the Environment  
Grantham Research Institute for Climate Change and the Environment

A thesis submitted to the Department of Geography and the Environment of the London  
School of Economics and Political Science for the degree of Doctor of Philosophy

London, December 2022

*In memory of Lorenzo Lunghi*

*And if the world has ceased to hear you,  
Say to the silent earth: I flow.  
To the rushing water, speak: I am.*

*(R.M. Rilke)*

## **Declaration**

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 50,000 words, excluding the bibliography.

## **Statement of Inclusion of Co-authored Work**

I hereby confirm that the entirety of the Introduction, Chapter 1 and Chapter 2 are the result of my own work only. Chapter 3, “The Unintended Consequences of Colombia’s Covid-19 Lockdown on Forest Fires”, is co-authored with my supervisor, Dr. Charles Palmer (LSE), Monica Amador-Jiménez and Naomi Millner (University of Bristol), and Toby Pennington (University of Exeter), and was published in 2020 on *Environmental and Resource Economics*. In particular, I have devised the empirical strategy, performed the data collection, cleaning, and analysis, produced all tables and figures, contributed to the manuscript by writing the entirety of the Data and Methods and Results sections, and contributed to the paper by collaborating on the Discussion and Conclusions sections. Chapter 4, “Carbon emissions reductions from Indonesia’s moratorium on forest concessions are cost-effective yet contribute little to Paris pledges”, is the result of equal contribution between myself and my two supervisors, Dr. Charles Palmer (LSE) and Prof. Ben Groom (University of Exeter), and was published in 2022 on *Proceedings of the National Academy of the Sciences*.

## Acknowledgements

No doctoral journey is possible without a collective effort, and mine naturally makes no exception. I therefore wish to use this space to thank everyone who stood in my corner over these years, conscious that the boundaries of a written page are much too narrow when it comes to expressing gratitude.

Thanks first, foremost, and in rigorous alphabetical order, to my supervisors Ben Groom and Charlie Palmer, for always treating me as their intellectual peer, for the controlled chaos and the creative destruction, and for their assiduous supervision. Thank you for being the best team I could wish for. You made my journey truly enjoyable, which is quite something when the journey is a PhD.

Thanks to Simon Dietz, Eugenie Dugoua, and Sefi Roth, for setting stellar standards within our Department. I am grateful to Marco Sanfilippo, Matteo Fiorini and Leonardo Baccini for giving me the opportunity to join their project at the European University Institute, in my hometown of Florence. A wholehearted thank you to Elisabetta Pietrostefani, for taking me onboard her spectacular course as a GTA. Thanks to Paul Ferraro and Hendrik Wolff for agreeing to review my thesis.

I am deeply grateful to have been able to pursue this adventure in a joyous, friendly, and collective environment that smoothed the lows and enhanced the highs of the ever tortuous doctoral path. Special thanks to new and old friends Chiara, Glen, Yang, Julien, Antonio, Bea, Caps, Jeff, Till, Romano, the rest of the Geography PhD cluster, and the Italians: Alberto, Alessandro, Armando, and Giovanni. An equally emphatic thank you goes to Achilleas, Arianna, and Giacomo, for sharing this adventure from afar, but ever so close. You all deserve the success you experience.

Thank you to Dan, my flatmate and very best friend in London, for the late-night chats, the music-infused cooking sessions, the citizenship tests, the Stroud Green market strolls, the pints at the Fullback. Thank you for choosing to live with me for six years, even when it meant sleeping on a boat in the English winter.

Thanks to my extended family: Albi, Guga, Leo, Neri, Seve, and Zano, for always being there for each other, spontaneously, since day one. To the rest of the Florence squad – Fede, Franci, Vitto, Marta, Nene, Caro, Ire, Lazze, Cesa, Giuso, and Matte K: you have all kept me going more than you can imagine. Thanks for all the great times we've had and for the ones we haven't had yet.

I am forever grateful to my parents, for giving me and my brother Giovanni their love, but not their thoughts; for we had our own thoughts. Thank you for being the bows from which you sent us forth as living arrows.

Finally, thank you, Elena.

You are the thread that keeps all from unravelling.

You are my rock, you are my island, you are my home.

## Abstract

This thesis consists in four papers in spatial and environmental economics, in which causal inference methods are employed to analyse two topics: carbon pricing and deforestation. The dissertation is comprised of two parts. The first evaluates the impacts of the 2008 carbon tax implemented in British Columbia, Canada; the second analyses illegal deforestation in Colombia and REDD+ policies in Indonesia. Chapter one studies the effect of British Columbia's carbon tax on road transportation CO<sub>2</sub> emissions and on carbon leakage due to cross-border fuel shopping in the USA. Using the synthetic control method and its extensions, we find that the tax is associated with a decrease in transportation CO<sub>2</sub> emissions. However, this effect is not statistically significant, with no role detected for cross-border fuel shopping. In chapter two, we analyse the impact of the tax on PM<sub>2.5</sub> concentrations arising from transportation. We detect a statistically significant effect of the carbon tax on air pollution co-benefits, which is heterogeneously distributed across metropolitan areas. Less polluted, less dense and richer areas see greater reductions in air pollution, identifying a post-tax increase in inequality with respect to pollution exposure. Reductions are driven by a switch in commute mode towards low emissions means of transport, principally public transit. Health gains from the tax are large, and regressively co-vary with income. Chapter three focuses on the effects of Colombia's 2020 Covid-19 lockdown on forest fires. We find that the lockdown is associated with an upsurge in cumulative fires, which is correlated with the presence of armed groups. Chapter four evaluates the effect of the 2011 Indonesian Moratorium on oil palm, timber, and logging concessions. We find that dryland forest inside the Moratorium experienced, at most, a 0.65% rate of forest cover retention compared to non-Moratorium areas, while no effect is detected for carbon-rich peatland. The implied effective carbon price is below US\$ 5/tCO<sub>2</sub>-eq. Moreover, the Moratorium only contributes 3-4% towards Indonesia's 2015 Paris commitment of a 29% reduction in deforestation by 2030.

# Contents

<b>Introduction</b>	<b>17</b>
<b>I The Impact of Carbon Pricing on Emissions</b>	<b>30</b>
<b>1 Carbon Pricing with Permeable Borders: The Impact of British Columbia’s Carbon Tax on Road Transportation CO<sub>2</sub> Emissions</b>	<b>31</b>
1.1 Introduction . . . . .	32
1.2 The 2008 British Columbian carbon tax . . . . .	35
1.2.1 Background . . . . .	35
1.2.2 Review of the literature . . . . .	35
1.3 Data . . . . .	37
1.3.1 Effect of the 2008 carbon tax on CO <sub>2</sub> emissions . . . . .	37
1.3.2 Carbon leakage to Washington State, USA . . . . .	39
1.4 Empirical Strategy . . . . .	40
1.4.1 Synthetic Control Method . . . . .	40
1.4.2 Extensions of the Synthetic Control Method . . . . .	42
1.5 Results . . . . .	44
1.5.1 Synthetic control method: naive implementation . . . . .	44
1.5.2 Implementation fulfilling contextual requirements . . . . .	50
1.5.3 Extensions of the synthetic control method . . . . .	53
1.6 Potential confounder: carbon leakage from cross-border travel . . . . .	55
1.6.1 Empirics . . . . .	60
1.7 Discussion . . . . .	65
1.8 Conclusion . . . . .	67
1.A Full Sample Results . . . . .	71
1.B Alternative specification of the demand model . . . . .	72
<b>2 Carbon Pricing with Regressive Co-benefits: Air Quality and Health Effects from British Columbia’s Carbon Tax are Positive but Unequally Distributed</b>	<b>73</b>
2.1 Introduction . . . . .	74
2.2 Data and Descriptive Statistics . . . . .	77
2.3 Empirical Strategy . . . . .	82
2.3.1 Two-way fixed effects difference-in-differences (TWFE-DID) . . . . .	82
2.3.2 Synthetic control method and synthetic difference-in-differences . . . . .	84
2.3.3 Heterogeneous Treatment Effects . . . . .	86

2.3.4	Mechanisms . . . . .	87
2.4	Results . . . . .	89
2.4.1	Main Specification . . . . .	89
2.4.2	Robustness Checks . . . . .	91
2.4.3	Heterogeneous Treatment Effects . . . . .	93
2.5	Mechanisms . . . . .	95
2.6	Health Gains . . . . .	98
2.7	Conclusions . . . . .	101
2.A	Additional Descriptive Statistics . . . . .	106
2.B	Additional Regression Results . . . . .	110
2.B.1	Composition of Main SDID Control . . . . .	110
2.B.2	DAs in the Vancouver CMA . . . . .	111
2.B.3	DAs matching NAPS Monitoring Stations . . . . .	112
2.B.4	Post-treatment period limited to 2013 . . . . .	113
2.C	Additional Mechanisms Tables . . . . .	114
2.D	Additional Health Results . . . . .	115
2.D.1	Estimates using RR from <a href="#">Krewski et al. (2009)</a> . . . . .	115
2.D.2	Health-income relationships . . . . .	116

## II Causal Analyses of Tropical Deforestation 117

<b>3</b>	<b>The Unintended Impact of Colombia’s Covid-19 Lockdown on Forest Fires</b>	<b>118</b>
3.1	Introduction . . . . .	119
3.2	Background to Deforestation, Conflict and Covid-19 in Colombia . . . . .	120
3.3	Data and Methods . . . . .	123
3.3.1	Data Description . . . . .	123
3.3.2	Methodology . . . . .	126
3.4	Results . . . . .	129
3.5	Discussion . . . . .	135
3.6	Conclusion . . . . .	138
3.A	Additional Figures . . . . .	142
3.B	Additional Results . . . . .	144
<b>4</b>	<b>Carbon emissions reductions from Indonesia’s Moratorium on forest concessions are cost-effective yet contribute little to Paris pledges</b>	<b>148</b>
4.1	Introduction . . . . .	149
4.2	Results . . . . .	152
4.2.1	Forest cover trends in Indonesia . . . . .	152
4.2.2	National-level effects of the Moratorium . . . . .	154
4.2.3	Testing for leakage . . . . .	155
4.2.4	Meeting Indonesia’s NDC commitments and the effective carbon price . . . . .	156
4.3	Discussion . . . . .	159
4.3.1	Impacts of Indonesia’s Moratorium on forest loss and emissions	159
4.3.2	Comparing estimates of impact . . . . .	159

4.3.3	Beyond the Moratorium . . . . .	160
4.3.4	Delivering carbon value for money and meeting the NDC commitments . . . . .	161
4.4	Materials and Methods . . . . .	162
4.4.1	Data . . . . .	162
4.4.2	Empirical approach . . . . .	164
4.4.3	Estimating carbon emissions . . . . .	170
4.A	Background to the Moratorium . . . . .	178
4.A.1	The Indonesia-Norway REDD+ partnership . . . . .	178
4.A.2	The 2011 Moratorium on palm oil, timber and logging concessions	179
4.A.3	How the Indonesia-Norway partnership estimated forest-based emissions reductions in 2017 . . . . .	184
4.B	Main and Additional Results . . . . .	187
4.B.1	ATT results: year-on-year . . . . .	187
4.B.2	Estimated avoided forest cover loss and carbon emissions reductions . . . . .	189
4.B.3	Robustness checks (2017 endpoint) . . . . .	190
4.B.4	“Intact primary” forest results . . . . .	196
4.B.5	Spatially heterogeneous effects of the Moratorium, 2011-17 . . . . .	200
4.C	Summary Statistics . . . . .	201
4.D	Methods . . . . .	205
4.D.1	Summary of empirical approach . . . . .	205
4.D.2	Model selection and balance tests . . . . .	213
4.E	Robustness . . . . .	234
4.E.1	Sensitivity to matching procedures . . . . .	234



# List of Figures

1.1	(A) Trends in road transportation and (B) total CO <sub>2</sub> emissions for British Columbia, an average of all other Canadian provinces and an average of all US States between 1990 and 2017. . . . .	38
1.2	Path plot of road transportation CO <sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia . . . . .	45
1.3	Gap plot of road transportation CO <sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia . . . . .	46
1.4	Path plot of road transportation CO <sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia, in-time placebo with treatment in 2002 . . . . .	48
1.5	In-space placebo . . . . .	49
1.6	MSPE ratio test . . . . .	49
1.7	Leave-one-out test . . . . .	50
1.8	Path plot of road transportation CO <sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia, fulfilling contextual requirements . . . . .	52
1.9	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on road transportation CO <sub>2</sub> emissions in British Columbia. . . . .	54
1.10	DID, SCM and SDID units weights. . . . .	55
1.11	Model parameterisation with low individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price. . . . .	59
1.12	Model parameterisation with high individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price. . . . .	59
1.13	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on gasoline stations (NAICS 447) taxable retail sales in Whatcom County, WA. . . . .	63
1.14	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on auto dealers and gasoline stations (NAICS 441-447) taxable retail sales in Whatcom County, WA. . . . .	63
1.15	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on the food sector (NAICS 445) taxable retail sales in Whatcom County, WA. . . . .	64

1.16	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on all NAICS retail trade codes (44-45) taxable retail sales in Whatcom County, WA. . . . .	64
1.A.1	DID, SCM and SDID estimates of the effect of the 2008 carbon tax on road transportation CO <sub>2</sub> emissions in British Columbia, full control pool. . . . .	71
1.B.1	Model parameterisation with low individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price. . . . .	72
1.B.2	Model parameterisation with high individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price. . . . .	72
2.2.1	Satellite PM <sub>2.5</sub> (Meng <i>et al.</i> , 2019) vs PM <sub>2.5</sub> from NAPS monitoring stations. Both measures are in $\mu\text{g}/\text{m}^3$ . The correlation coefficient is 0.597. . . . .	79
2.2.2	Satellite PM <sub>2.5</sub> (Meng <i>et al.</i> , 2019) vs Satellite PM <sub>2.5</sub> (van Donkelaar <i>et al.</i> , 2019). Both measures are in $\mu\text{g}/\text{m}^3$ . The correlation coefficient is 0.729. . . . .	79
2.2.3	Spatial distribution of PM <sub>2.5</sub> and relevant covariates within the Vancouver CMA. Top row: PM <sub>2.5</sub> , population density and road density; Middle row: NTLs, dwelling value and median income; Bottom row: high emission commute mode %, public transport % and cycling/walking %. . . . .	81
2.3.1	Trends in satellite PM <sub>2.5</sub> (Meng <i>et al.</i> , 2019), British Columbia vs control provinces, between 2000 and 2016. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. . . . .	83
2.3.2	Trends in satellite PM <sub>2.5</sub> (van Donkelaar <i>et al.</i> , 2019), British Columbia vs control provinces, between 2000 and 2018. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. . . . .	83
2.4.1	Graphical results from DID, SCM and SDID for PM <sub>2.5</sub> concentrations, with Meng <i>et al.</i> (2019) data. . . . .	89
2.4.2	Graphical results from DID, SCM and SDID for PM <sub>2.5</sub> concentrations, with van Donkelaar <i>et al.</i> (2019) data. Time weights $\lambda_t$ are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period. . . . .	91
2.4.3	Graphical SDID results: heterogeneity by quintiles of baseline characteristics for the treated sample. Panel A) Quintiles of baseline PM <sub>2.5</sub> ; B) Quintiles of baseline population density; C) Quintiles of road density; D) Quintiles of baseline nightlights values. 95% confidence intervals are shaded in grey. . . . .	94

2.6.1	Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from <a href="#">Lepeule <i>et al.</i> (2012)</a> , for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs. . . . .	100
2.A.1	Availability of PM <sub>2.5</sub> readings in the National Atmospheric Surveillance Program database between 2000 and 2018. Lighter colours indicate higher availability of readings, and hence, monitoring stations which were added earlier. . . . .	106
2.A.2	Gasoline vs Other fuel car sales in BC 2011-2021. . . . .	107
2.A.3	Minor fuel groups car sales in BC 2011-2021. . . . .	107
2.A.4	Passenger cars vs Truck and SUV sales, large Canadian Provinces, 1990-2021. . . . .	108
B.1	Composition of the synthetic unit of <a href="#">Figure 2.4.1</a> . . . . .	110
B.2	Composition of the synthetic unit of <a href="#">Figure 2.4.2</a> . . . . .	110
B.3	Graphical results from DID, SCM and SDID for PM <sub>2.5</sub> concentrations, with <a href="#">Meng <i>et al.</i> (2019)</a> data, dataset restricted to DAs in the Vancouver CMA. . . . .	111
B.4	Graphical results from DID, SCM and SDID for PM <sub>2.5</sub> concentrations, with <a href="#">Meng <i>et al.</i> (2019)</a> data, dataset restricted to DAs matching NAPS monitoring stations' locations. . . . .	112
B.5	Graphical results from DID, SCM and SDID for PM <sub>2.5</sub> concentrations, with <a href="#">Meng <i>et al.</i> (2019)</a> data, dataset restricted to 2013. . . . .	113
D.1	Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from <a href="#">Krewski <i>et al.</i> (2009)</a> , for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs. . . . .	115
D.2	Bivariate distribution of health gains using the RR from <a href="#">Lepeule <i>et al.</i> (2012)</a> and median income for the Vancouver (top panel), Victoria (middle panel) and Abbotsford (bottom panel) CMAs. . . . .	116
3.3.1	Location of fire hotspots in Colombia, 14 March–28 May 2020. <i>Sources:</i> <a href="#">Hansen <i>et al.</i> (2013)</a> , NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/Radiometer Suite (VIIRS) . . . . .	124
3.3.2	Cumulative number of forest fires by year in Colombia, 2012–2020. <i>Note:</i> 95% Confidence interval around the historical mean shaded in grey. <i>Source:</i> NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/Radiometer Suite (VIIRS) . . . . .	125
3.3.3	Daily number of forest fires by year in Colombia, 2012–2020. <i>Source:</i> NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/Radiometer Suite (VIIRS) . . . . .	125
3.4.1	Actual series, historical mean, SCM and ASCM, for all of Colombia. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	130

3.4.2	In-time Placebo Test for SCM and ASCM, with lockdown beginning on 19 February 2020. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	130
3.4.3	Results for the Amazon region, conditional on armed group presence. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	132
(a)	Areas with FARC-EP dissidents. . . . .	132
(b)	Areas with no known presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries. . . . .	132
3.4.4	Results for the Andes region, conditional on armed group presence. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	133
3.A.1	Colombia’s biotic regions. Areas within regions denote the boundaries of Colombian municipalities. . . . .	142
3.A.2	Mobility by destination in Colombia, 2020. . . . .	142
3.A.3	Mobility, walking or driving, Colombia, 2020. . . . .	143
3.A.4	Known locations of FARC-EP dissidents and the Gulf Clan in Colombia by municipality, 2019. . . . .	143
3.B.1	Alternative dependent variable, Fire Radiative Power (FRP). <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	144
3.B.2	Results for the Caribbean region, conditional on armed group presence. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	145
(a)	Areas with Gulf Clan neo-paramilitaries. . . . .	145
(b)	Areas with FARC-EP dissidents and Gulf Clan neo-paramilitaries. . . . .	145
(c)	Areas with no presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries. . . . .	145
3.B.3	Results for the Orinoco region, conditional on armed group presence. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	146
3.B.4	Results for the Pacific region, conditional on armed group presence. Only areas with Gulf Clan neo-paramilitaries. <i>Note:</i> 95% Confidence Interval around the ASCM series shaded in grey. . . . .	146
3.B.5	Location of fire hotspots in Puerto Boyacá, 14 March–28 May 2020. . . . .	147
4.2.1	Forest cover trends inside and outside the Moratorium, 2000-2018: non-concession dryland grid cells (A); non-concession peatland grid cells (B); concession dryland grid cells (C); concession peatland grid cells (D). Shaded areas denote treatment period. Grid cells in (A) and (B) also exclude forest in protected areas. . . . .	153

4.2.2	Cumulative avoided forest loss ('000 ha) and avoided carbon dioxide emissions (MtCO <sub>2</sub> -eq): dryland forest DD, 2012-2018 (A); dryland forest DDD, 2012-2018 (B); peatland forest DD, 2012-2017 (C); peatland DDD, 2012-2017 (D). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI. . . . .	154
4.2.3	Regression discontinuity LATE (with <a href="#">Calonico <i>et al.</i> (2014)</a> bandwidth), 2005-2018. Scatterplots of all the observations within and outside the Moratorium's boundaries are shown in <a href="#">section 4.E</a> . Error bars denote the 95% CI. . . . .	156
4.A.1	Share of pre-2011 palm oil, timber and logging concessions (combined) by district. . . . .	180
4.A.2	Extent of Moratorium in Indonesia, 2011. The 2011 Moratorium boundaries denote the treatment areas for our empirical analysis, and are assumed constant throughout our treatment periods. Sources: Authors' elaboration on forest cover data ( <a href="#">Hansen <i>et al.</i>, 2013</a> ). Moratorium shapefile from <a href="#">Ministry of Forestry of the Republic of Indonesia (2011)</a> . . . . .	182
4.B.1	Cumulative avoided forest loss ('000 ha) and avoided carbon dioxide emissions (MtCO <sub>2</sub> -eq): Dryland forest > 1,000m DDD, 2012-2018 (A); Peatland > 1,000m DDD, 2012-2017 (B). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI. . . . .	195
4.B.2	Cumulative avoided "intact primary" forest loss ('000 ha) and avoided carbon dioxide emissions (MtCO <sub>2</sub> -eq): dryland forest DD, 2012-2018 (A); dryland forest DDD, 2012-2018 (B); peatland forest DD, 2012-2017 (C); peatland DDD, 2012-2017 (D). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI. . . . .	199
4.B.3	Spatially heterogeneous effects of the Moratorium on dryland and peatland forests, 2011-17 . . . . .	200
4.D.1	Forest cover trends inside and outside the Moratorium, 2005-2018: non-concession dryland grid cells. . . . .	216
4.E.1	Regression discontinuity LATE with <a href="#">Calonico <i>et al.</i> (2014)</a> bandwidth (A); regression discontinuity LATE with <a href="#">Imbens and Kalyanaraman (2012)</a> bandwidth (B). . . . .	237
4.E.2	Forest cover in (A) 2004 and (B) 2010, with respect to distance from Moratorium boundaries . . . . .	238

4.E.3	Deforestation in the proximity of Moratorium boundaries, <a href="#">Calonico et al. (2014)</a> bandwidth, pre-treatment years. . . . .	239
4.E.4	Deforestation in the proximity of Moratorium boundaries, <a href="#">Calonico et al. (2014)</a> bandwidth, post-treatment years. . . . .	240
4.E.5	Deforestation in the proximity of Moratorium boundaries, <a href="#">Imbens and Kalyanaraman (2012)</a> bandwidth, pre-treatment years. . . . .	241
4.E.6	Deforestation in the proximity of Moratorium boundaries, <a href="#">Imbens and Kalyanaraman (2012)</a> bandwidth, post-treatment years. . . . .	242
4.E.7	Regression discontinuity robustness checks . . . . .	243

# List of Tables

1.1	Covariate balance between British Columbia and synthetic British Columbia . . . . .	45
1.2	Composition of synthetic British Columbia . . . . .	46
1.3	Summary of $\hat{\tau}$ point estimates and relative standard errors from seven different estimation methods, control pool restricted to Canadian provinces. . . . .	54
1.4	Determinants of border crossings: log-log . . . . .	61
1.A.1	Summary of $\hat{\tau}$ point estimates and relative standard errors from seven different estimation methods, full control pool. . . . .	71
2.4.1	Summary of $\hat{\tau}$ point estimates and standard errors from all estimation methods, dependent variable from <a href="#">Meng <i>et al.</i> (2019)</a> . . . . .	90
2.4.2	Summary of $\hat{\tau}$ point estimates and standard errors from all estimation methods, dependent variable from <a href="#">van Donkelaar <i>et al.</i> (2019)</a> . . . . .	92
2.4.3	SDID results for 10th and 90th percentiles, with 90-10 ratios. . . . .	95
2.5.1	TWFE-DID results for high emissions commute mode . . . . .	97
2.5.2	TWFE-DID results for low emissions commute mode . . . . .	97
2.A.1	Summary Statistics, 2000-2007 . . . . .	109
2.A.2	Summary Statistics, 2008-2018 . . . . .	109
B.1	Summary of $\hat{\tau}$ point estimates and standard errors from all estimation methods, dependent variable from <a href="#">Meng <i>et al.</i> (2019)</a> , dataset restricted to DAs in the Vancouver CMA. . . . .	111
B.2	Summary of $\hat{\tau}$ point estimates and standard errors from all estimation methods, dependent variable from <a href="#">Meng <i>et al.</i> (2019)</a> , dataset restricted to DAs matching NAPS monitoring stations' locations. . . . .	112
B.3	Summary of $\hat{\tau}$ point estimates and standard errors from all estimation methods, dependent variable from <a href="#">Meng <i>et al.</i> (2019)</a> , dataset restricted to 2013. . . . .	113
C.1	TWFE-DID results for public transport . . . . .	114
C.2	TWFE-DID results for zero emissions commute mode . . . . .	114
4.2.1	The Moratorium's contribution to Indonesia's NDC commitments and the effective carbon price, 2011-2017 . . . . .	158
4.B.1	DD dryland forest ATT . . . . .	187
4.B.2	DDD dryland forest ATT . . . . .	187
4.B.3	DD peatland forest ATT . . . . .	188
4.B.4	DDD peatland forest ATT . . . . .	188

4.B.5	DD dryland forest avoided forest cover loss and carbon emissions reductions . . . . .	189
4.B.6	DDD dryland forest avoided forest cover loss and carbon emissions reductions . . . . .	189
4.B.7	DD peatland forest avoided forest cover loss and carbon emissions reductions . . . . .	189
4.B.8	DDD peatland forest avoided forest cover loss and carbon emissions reductions . . . . .	190
4.B.9	DD dryland forest trimmed sample results (2017 endpoint) . . . . .	190
4.B.10	DDD dryland forest trimmed sample results (2017 endpoint) . . . . .	190
4.B.11	DD peatland forest trimmed sample results (2017 endpoint) . . . . .	191
4.B.12	DDD peatland forest trimmed sample results (2017 endpoint) . . . . .	191
4.B.13	Coarsened Exact Matching dryland forest results (2017 endpoint) . . . . .	192
4.B.14	Coarsened Exact Matching peatland forest results (2017 endpoint) . . . . .	192
4.B.15	DDD dryland forest ATT: caliper = 0.001 . . . . .	193
4.B.16	DDD peatland forest ATT: caliper = 0.001 . . . . .	193
4.B.17	DDD dryland: robustness (2017 endpoint) . . . . .	194
4.B.18	DDD peatland: robustness (2017 endpoint) . . . . .	194
4.B.19	DD dryland “intact primary” forest ATT . . . . .	196
4.B.20	DDD dryland “intact primary” forest ATT . . . . .	196
4.B.21	DD peatland “intact primary” forest ATT . . . . .	197
4.B.22	DDD peatland “intact primary” forest ATT . . . . .	197
4.B.23	DD dryland “intact primary” forest, avoided forest cover loss and carbon emissions reductions . . . . .	197
4.B.24	DDD dryland “intact primary” forest, avoided forest cover loss and carbon emissions reductions . . . . .	198
4.B.25	DD peatland “intact primary” forest, avoided forest cover loss and carbon emissions reductions . . . . .	198
4.B.26	DDD peatland “intact primary” forest, avoided forest cover loss and carbon emissions reductions . . . . .	198
4.C.1	Dryland forest non-concession grid cells, Moratorium . . . . .	201
4.C.2	Dryland forest non-concession grid cells, non-Moratorium . . . . .	201
4.C.3	Dryland forest concession grid cells, Moratorium . . . . .	202
4.C.4	Dryland forest concession grid cells, non-Moratorium . . . . .	202
4.C.5	Peatland forest non-concession grid cells, Moratorium . . . . .	203
4.C.6	Peatland forest non-concession grid cells, non-Moratorium . . . . .	204
4.C.7	Peatland forest concession grid cells, Moratorium . . . . .	204
4.C.8	Peatland forest concession grid cells, non-Moratorium . . . . .	205
4.D.1	Results of model selection (endpoint 2017) . . . . .	220
4.D.2	Placebo test in time 1: DD dryland, 2005-2010 outcome, same matching . . . . .	221
4.D.3	Placebo test in time 2: DD dryland, 2005-2010 outcome, matching up to 2004 . . . . .	221
4.D.4	Placebo test in time 1: DD peatland, 2005-2010 outcome, same matching . . . . .	221



4.D.5	Placebo test in time 2: DD peatland, 2005-2010 outcome, matching up to 2004 . . . . .	222
4.D.6	Placebo test in time 3: DDD dryland, placebo treatment in 2005, same matching . . . . .	222
4.D.7	Spatial placebo: Dryland, 2011-2017 outcome . . . . .	223
4.D.8	Spatial placebo: Peatland, 2011-2017 outcome . . . . .	223
4.D.9	Dryland forest, untrimmed, matching up to 2010 . . . . .	224
4.D.10	Dryland forest, untrimmed, matching up to 2004 . . . . .	225
4.D.11	Dryland forest, > 30% forest cover in 2005, matching up to 2010 . . . . .	225
4.D.12	Dryland forest, > 30% forest cover in 2005, matching up to 2004 . . . . .	226
4.D.13	Dryland forest, > 60% forest cover in 2005, matching up to 2010 . . . . .	226
4.D.14	Dryland forest, > 60% forest cover in 2005, matching up to 2004 . . . . .	227
4.D.15	Peatland forest, untrimmed, matching up to 2010 . . . . .	227
4.D.16	Peatland forest, untrimmed, matching up to 2004 . . . . .	228
4.D.17	Peatland forest, > 30% forest cover in 2005, matching up to 2010 . . . . .	228
4.D.18	Peatland forest, > 30% forest cover in 2005, matching up to 2004 . . . . .	229
4.D.19	Peatland forest, > 60% forest cover in 2005, matching up to 2010 . . . . .	229
4.D.20	Peatland forest, > 60% forest cover in 2005, matching up to 2004 . . . . .	230
4.D.21	Dryland forest, cells above 1000m, matching up to 2010 . . . . .	230
4.D.22	Peatland forest, cells above 1000m, matching up to 2010 . . . . .	231
4.D.23	Dryland forest, all-elevation cells, matching up to 2010 . . . . .	231
4.D.24	Peatland forest, all-elevation cells, matching up to 2010 . . . . .	232
4.D.25	Dryland “intact primary” forest, matching up to 2010 . . . . .	232
4.D.26	Peatland “intact primary” forest, matching up to 2010 . . . . .	233

# Introduction

This dissertation contains four independent empirical studies, which examine two main topics: the impact of carbon pricing on emissions from the transportation sector in a developed economy, and drivers and determinants of deforestation and forest conservation in two tropical contexts.

All papers are characterised by the study of place-based policies and shocks, i.e. discontinuities in institutional and regulatory frameworks with a precise geographical scope. Furthermore, all chapters analyse the effects of these shocks with an explicit spatial dimension: we leverage large, highly disaggregated data to enquire about general impacts, but also carefully examine spatial spillovers, heterogeneous effect across geographies and correlations with other geolocated features.

Across all chapters, there is certainly a methodological common denominator: the thesis indeed draws from, and attempts to contribute to, the literature in applied spatial and environmental economics, leveraging state of the art data and econometric techniques in order to provide causal estimates of the phenomena under investigation. The work contained in this dissertation is inspired by the “credibility revolution in empirical economics” (Angrist and Pischke, 2010), and ultimately rooted in the *potential outcomes framework* (Rubin, 2005; Athey and Imbens, 2017). All studies consist indeed in natural experiments, which enquire about the effect of a *treatment*, be it a policy or an exogenous shock, on the outcome of interest. Since the outcome path of the observational unit(s) under investigation in the absence of a policy, or its *counterfactual* outcome, is by definition unobservable<sup>1</sup>, all the analyses attempt to reconstruct it using econometric techniques, in order to estimate the “true” causal effect of the shock on the treated entities.

The shocks under investigation in this thesis are: (i) the introduction, in 2008, of a Province-wide carbon tax in British Columbia, Canada (in the first two chapters); (ii) the 2020 Covid-19 lockdown and consequent mobility restrictions in Colombia (in the third chapter); (iii) the 2011 REDD+ agreement between the Government of Norway and the Government of Indonesia in order to reduce emissions from deforestation and forest degradation on the Indonesian territory (in the fourth chapter). We leverage this array of sharp policy-induced and contextual changes in order to examine a set of fundamental problems in empirical environmental economics, which is described in detail in the following sections.

---

<sup>1</sup>A fact known as the “fundamental problem of causal inference” (Holland, 1986).

## Part I: The Impact of Carbon Pricing on Emissions

In the first chapter of the thesis, we primarily attempt to answer the question: is a market-based instrument such as a carbon tax a sufficient condition to drive down CO<sub>2</sub> emissions from the transportation sector?

The past century has withstood the development of an enormous body of theoretical work on the issue of carbon pricing, which has permeated the field since the 1920s (Pigou, 1920) and gained traction in the 1970s with the work of Baumol (1972), Weitzman (1974), and Nordhaus (1977). The two principal market-based instruments devised by economists in order to introduce a price on carbon are carbon taxes, or levies proportional to the amount of carbon contained in a marketed good, and emission trading schemes (ETSs), i.e. systems where a target fixed quota of periodical emissions is allocated or auctioned among emitters, who are then allowed to trade permits between themselves.

Notwithstanding the economic consensus towards a preference for carbon taxes (Weitzman, 2015; Nordhaus, 2008), ETSs have been more widely implemented and have usually covered larger shares of the world's total emission flows (World Bank, 2022). This is possibly due to the relative political infeasibility of imposing additional taxation due to citizens' beliefs and tax aversion (Douenne and Fabre, 2022). One of the main strategies to gain public support for carbon taxes is domestic revenue recycling (Carattini *et al.*, 2019), which is the defining feature of the British Columbian carbon tax (Murray and Rivers, 2015); however, recent research has shown that the impact of tax and rebate programmes on carbon tax approval rates is low, and that partisan politics plays a much more defining role in shaping citizens' beliefs (Mildenberger *et al.*, 2022). To give an idea of the relative paucity of long-standing carbon tax schemes, only 15 pricing mechanisms were already in place in 2008 (covering approximately 5% of the global GHG emission flows, and mostly through the EU ETS) (World Bank, 2022).

It is thus unsurprising that, with a few stark exceptions (e.g. Andersson, 2019), the *ex post* empirical evidence on the effectiveness of carbon taxes on CO<sub>2</sub> emissions flows is scattered. Furthermore, the case of Sweden examined by Andersson (2019) is a peculiarity rather than the norm: the Swedish carbon tax, currently at \$130/tCO<sub>2</sub>, has indeed been the highest in the record from its implementation until the recent introduction of Uruguay's carbon tax in January 2022. While it is certainly of utmost importance to assess the effectiveness of carbon taxation in general, the first part of this paper is instead concerned with estimating the eventual environmental benefits of a carbon tax levied at a much lower rate, which was lauded as the "grand experiment in climate policy" (Murray and Rivers, 2015), and was the first carbon tax to be introduced in North America. With the British Columbian carbon tax having been rolled out on a Canadian federal basis in 2019, but still sitting well below the recommended 2030 carbon price corridor 14 years after its implementation (World

Bank, 2022; High-Level Commission on Carbon Prices, 2017), it is fundamental to understand its short and medium-term impact on emissions, given that similar schemes will need to proliferate across developed economies in the next decades in order to attain decarbonisation. In the first part of the first chapter we thus attempt to identify the causal impact of the 2008 British Columbian carbon tax on road transportation CO<sub>2</sub> emissions. We conduct the main analysis at the Provincial aggregate level, using the synthetic control method for comparative case studies (Abadie and Gardeazabal, 2003; Abadie *et al.*, 2015), for which we carefully assess robustness to contextual requirements (Abadie, 2021). Furthermore, we introduce recent extensions to synthetic controls, such as the synthetic difference-in-differences from Arkhangelsky *et al.* (2021), de-meaned synthetic controls (Doudchenko and Imbens, 2016; Ferman and Pinto, 2021) and matrix completion methods (Athey *et al.*, 2021) for the first time in the literature on the empirical evaluation of carbon taxes.

A corollary to the main question is the investigation of eventual *carbon leakage* arising from the tax: the possibility that, given the peculiar geographic location of British Columbia<sup>2</sup> and the attractive fuel price gradient at the US border, the carbon tax wedge pushed BC residents to cross the US customs to shop for fuel. This issue is only anecdotally mentioned in press articles, and in passing in peer reviewed publications (Rivers and Schaufele, 2015; Lawley and Thivierge, 2018; Andersson, 2019), but has never been examined empirically in the context of carbon taxation. We draw from the classic literature on product differentiation (Hotelling, 1929) and from a recent literature in industrial organisation (e.g. Chandra *et al.*, 2014; Friberg *et al.*, 2022) in order to first establish theoretical predictions of cross-border fuel shopping, and then empirically calibrate them with a novel US county-level dataset on retail trade revenues in the State of Washington, relying again on the synthetic difference-in-differences estimator. This line of empirical analysis constitutes the second part of Chapter 1.

Another feature of the introduction of carbon pricing is its ability to generate *environmental co-benefits*, such as the reduction of local pollutants (particulate matter, sulfur and nitrogen oxides, and black carbon) due to their complementarity with GHGs (Timilsina, 2022). Indeed, this class of pollutants is often emitted via the same channel as CO<sub>2</sub>, that is, through the combustion of fossil fuels (Li *et al.*, 2018). Therefore, the second chapter of Part I asks the following question: does a carbon tax generate environmental co-benefits through the reduction of small-scale particulate matter PM<sub>2.5</sub>?

While there is ample numerical evidence on the effect of carbon taxation on air pollution co-benefits (e.g. Li *et al.*, 2018; Woollacott, 2018; Zhang *et al.*, 2021; Vandyck *et al.*, 2018), this channel has never been examined in an observational setting using *ex post* data. However, in order to correctly calibrate cost-benefit analyses of carbon taxation, it is necessary to include both global benefits arising from reductions in CO<sub>2</sub> flows, and local benefits arising from improvements in air quality. In order to

---

<sup>2</sup>Whose population is mostly concentrated in a region with rapid access to the main metropolitan areas of the US State of Washington (Chandra *et al.*, 2014).

do so, we exploit highly disaggregated satellite-derived data on  $\text{PM}_{2.5}$ , in conjunction with Canadian census geometries at a fine geographical scale<sup>3</sup>, and examine the impact of the 2008 carbon tax on air pollution in the main metropolitan areas of British Columbia, employing the synthetic difference-in-differences methodology. In addition, we inspect the heterogeneity of this impact across space: it is indeed plausible that areas with lower access to alternative means of transport bear the burden of the carbon tax vis-à-vis areas with greater substitution options. Coupled with the extensively documented pollution disparities within North American metropolitan areas (e.g. [Jbaily et al., 2022](#); [Currie et al., 2020](#); [Sager and Singer, 2022](#)), this heterogeneity could give rise to a spatial dimension of *carbon tax regressiveness*, in addition to the well-known vertical income regressiveness shown e.g. by [Douenne \(2020\)](#). This enquiry constitutes the first part of Chapter 2.

An important exercise in order to understand the magnitude of the eventual gains or losses from improving or deteriorating air quality is to convert the estimates of pollution change into estimates of morbidity and mortality change. Here, we rely on the insights of the environmental health and epidemiology literatures (e.g. [Krewski et al., 2009](#); [Lepeule et al., 2012](#)), summarised by [Fowlie et al. \(2019\)](#), and obtain highly disaggregated maps of health benefits/costs by inter-relating mortality change with the Value of a Statistical Life. Nevertheless, environmental co-benefits have been shown to be multifaceted and extending far beyond the immediate effects on mortality (see [Aguilar-Gomez et al., 2022, for a review](#)), suggesting that our estimates will at most constitute a lower bound on the non-market impacts of a carbon tax. Furthermore, the interaction of these estimates with the distribution of the population of metropolitan areas can provide a picture of which social stratus reaps the eventual health benefits due to carbon taxation. The quantification of health effects constitutes the second part of Chapter 2 and closes Part I of this dissertation.

## Part II: Causal Analyses of Tropical Deforestation

The second part of this thesis consists in two empirical evaluations concerning deforestation and forest conservation in the tropical countries of Colombia and Indonesia. The main motivation behind the search for causality in biodiversity economics is that the discipline has lagged behind other fields of environmental economics in terms of the evaluation of policy impacts and of shocks to forest ecosystems ([Ferraro and Pattanayak, 2006](#)). However, especially but not only in light of the 2021-2030 UN Decade on Ecosystem Restoration, it is fundamental to understand the drivers and determinants of forest loss and to assess the impact of policies aimed at conserving Earth's largest reserves of terrestrial biodiversity and carbon sinks.

The third chapter of this thesis exploits one of the most salient worldwide exogenous shocks of the past decades, the Covid-19 pandemic, in order to evaluate whether it has

---

<sup>3</sup>The Dissemination Area, corresponding to US census tracts.

resulted in unintended consequences on the natural forests of Colombia’s major biotic regions. Motivated by the presence of collaborating researchers in the field, who informed us that the emergence of Covid-19 was accompanied by an unexpected increase in forest fires, we asked the following question: are the mobility restrictions connected to Colombia’s Covid-19 lockdown connected with the observed upsurge in forest fires?

In order to do so, we obtain near-real time, highly disaggregated data on the location and intensity of fire hotspots observed by satellites in the Colombian territory for 2012-2020, and exploit techniques which, in 2020, were being used by epidemiologists to track “excess mortality” due to the Covid-19 pandemic using historical death averages. We integrate these methods with cutting edge tools for policy evaluation used in classical econometrics (Abadie and Gardeazabal, 2003; Ben-Michael *et al.*, 2021), so as to improve upon the historical counterfactuals and obtain closer measurements of the *potential outcome* path of forest fires in the absence of the Covid-19 lockdown.

Moreover, aware of the strong interrelation between forest governance and militarisation of conflict areas in the Colombian context, we test whether the increase in forest fires is correlated with the presence of paramilitary organisations in Colombian municipalities. Indeed, the mobility restrictions imposed upon the civilian population might have resulted in a negative shock on forest monitoring by government entities and local communities which could have been exploited by organisations engaging in illegal forest clearing.

In the conclusive chapter of the dissertation, we instead enquire about the cost-effectiveness and environmental contribution of one of the largest REDD+<sup>4</sup> schemes worldwide, the 2011 Indonesian moratorium on new palm oil, timber and logging concessions in primary and peatland forests.

Indonesia’s GHG emissions from deforestation, forest degradation, peatland decomposition, and peat fires accounted for around a quarter of global emissions from the forestry sector during 2000-2016 (Government of Indonesia, 2018; FAO, 2021). After the establishment of the REDD+ framework at COP13 in 2007 (UNFCCC, 2008), the country struck an agreement with the Government of Norway, one of the prominent backers of REDD+ globally, in order to fund disbursements based on the achievement of emissions reductions from deforestation, up to \$1 billion (Government of Norway, 2010). One of the main policies resulting from the agreement is the aforementioned moratorium, which intends to prevent Indonesian district governments from granting new licenses for the conversion of primary forest and peatland into areas designated for palm oil and timber production, and selective logging.

In 2017, Indonesia announced that it had reduced emissions from deforestation by 11.2 MtCO<sub>2</sub>-eq, compared to a 10-year historical baseline calculated by averaging deforestation emissions between 2006 and 2016. As a consequence, Norway announced that a \$56.2 million payment, based on a carbon price of \$5/tCO<sub>2</sub>-eq, would be issued to Indonesia within the two countries’ REDD+ partnership. However, the historical

---

<sup>4</sup>Reducing emissions from deforestation and forest degradation.

counterfactual upon which Norway’s payment was based does not appropriately reflect the causal impact of the Moratorium in reducing emissions, since it conflates treatment and control *areas* (the baseline is calculated on *all* primary forest and peatland, not only those within the moratorium’s remit) and treatment and control *periods* (2011-2016 are included in the baseline but the moratorium was already in place).

In order to address these concerns, we observe that the appropriate quantification of a policy effect from the moratorium rests on the ability to identify a suitable counterfactual, i.e. a *control group* which reproduces potential forest cover trends in the absence of the moratorium. Since the moratorium area was not randomly assigned, it is highly likely that its characteristics differ substantially from those of non-moratorium forested areas. A simple difference in means between the rates of deforestation observed within its boundaries and outside of them would thus produce biased estimates of the policy’s impact. Therefore, we leverage a large literature in causal inference in the social sciences (e.g. Heckman *et al.*, 1997; Heckman *et al.*, 1998), which is steadily being introduced in biodiversity economics and conservation biology (e.g. Andam *et al.*, 2008; Joppa and Pfaff, 2010; Schleicher *et al.*, 2020) to perform statistical matching between moratorium and non-moratorium areas, thereby obtaining comparable treatment and control cohorts, which can be analysed in a difference-in-differences framework.

We augment the traditional difference-in-differences strategy with the inclusion of a third difference (Gruber, 1994; Chabé-Ferret and Subervie, 2013), which removes remaining divergences in pre-treatment deforestation trends after matching, in order to obtain the most precise estimate of the moratorium’s impact at our avail<sup>5</sup>. The *spatial dimension* of the policy impact is instead analysed by looking at whether deforestation was displaced outside the fixed boundaries of the moratorium (a similar type of *carbon leakage* to that examined in Chapter 1). Here, we exploit the fact that the boundary between the moratorium and non-moratorium areas is unlikely to trace a sharp discontinuity in forest and morphological characteristics: within a reasonable distance from the border, moratorium and non-moratorium areas are expected to be closely matched in terms of forest cover, deforestation trends, and terrain features relevant to deforestation. We thus exploit the regression discontinuity estimator (Lee and Lemieux, 2010; Calonico *et al.*, 2014), in order to identify the effect of this area-based programme on deforestation leakage.

Given that the performance of REDD+ initiatives is evaluated on the basis of their ability to reduce GHG emissions, and not deforestation, we then convert our estimates of policy impact (in terms of forest cover change) into CO<sub>2</sub>-equivalent changes, and into monetary terms using the stated \$5/tCO<sub>2</sub>-eq carbon price proposed by Norway to Indonesia in their agreement. We are thus able to determine whether the Norwegian government is getting carbon value for its payments, and whether

---

<sup>5</sup>At the time when the analysis was performed, the extensions of the synthetic control method to multiple treated units and the synthetic difference-in-differences method were not yet peer-reviewed nor were the statistical routines yet implemented.

Indonesia is on track to meet its 2015 Paris Agreement commitments through the 2011 moratorium.

## Methodological Considerations

During the writing of this thesis, there have been extraordinary developments in the field of applied econometrics, which need to be taken into account when tackling observational studies. In the first three chapters, we principally exploit the recent extensions to the category of estimators related to the synthetic control method (SCM) of [Abadie and Gardeazabal \(2003\)](#): mainly, the synthetic difference-in-differences estimator of [Arkhangelsky \*et al.\* \(2021\)](#), but also the augmented SCM by [Ben-Michael \*et al.\* \(2021\)](#), de-measured synthetic controls ([Doudchenko and Imbens, 2016](#); [Ferman and Pinto, 2021](#)), and matrix completion methods for programme evaluation introduced by [Athey \*et al.\* \(2021\)](#). While the SCM was originally developed in order to analyse treatments or shocks at an *aggregate* level, or imposed upon a single treated unit, at a single point in time, its extension are more flexible and can accommodate multiple treated units and staggered interventions.

An attractive feature of this class of estimators is certainly their ability to match treatment and control cohorts on the basis of their pre-intervention outcome paths, which is especially useful in the context of applied climate econometrics. In fact, especially when dealing with outcome variables derived from satellite observations (as in Chapters 2 and 3), but also with emissions estimates from governmental bodies (Chapter 1), it is fundamental that the control units' outcome paths closely resemble those of treated units. It is indeed difficult to predict these trends with accuracy using only socio-economic and morphological characteristics, as is the case in standard observational studies in economics. By matching on lags of the outcome variable in the estimation, synthetic methods can thus augment the confidence in pre-treatment matching quality, thereby ensuring that the policy impacts are estimated with precision. Nevertheless, given the relative ease with which these methods can be rolled out, a high degree of rigour is required on the researcher in order to avoid incurring in “perfunctory applications” of synthetic methods ([Abadie, 2021](#)). Appropriate falsification tests and pre-implementation checks need to be conducted in order to ensure the feasibility and robustness of the estimators, and avoid publishing misleading results which could wrongly inform the policymaker.

Chapter 4 takes a slightly different approach, motivated by the unavailability, at the time<sup>6</sup>, of the extensions of synthetic controls introduced by [Arkhangelsky \*et al.\* \(2021\)](#). Nonetheless, we rely on the latest advances in the field of biodiversity econometrics (see [Schleicher \*et al.\*, 2020](#), for a review), and employ statistical matching in order to obtain the best possible balance between treatment and control observations on the basis of observable characteristics. We introduce matched triple differences ([Chabé-Ferret and Subervie, 2013](#)) for the first time in an applied paper on forest

<sup>6</sup>Chapter 4 is chronologically the first to have been developed, in 2018. However, due to its large supplementary information, we have decided to close the thesis with it to ensure more fluid reading.



dynamics, in order to take into account the remaining unobservable factors which may have influenced deforestation trends differently across the treated and control groups. A potential re-evaluation of the moratorium could however exploit the synthetic difference-in-differences methodology to assess its performance vis-à-vis the matched triple difference estimator, and contribute to the burgeoning literature on programme evaluation in the field of biodiversity economics and conservation biology.

Causal inference in the field of biodiversity economics is possibly an even steeper challenge than in standard socio-economic studies, due to the heavily interlinked and complex nature of coupled human and natural systems (Ferraro *et al.*, 2019). Abiding to this stream of the literature, in Chapter 4 and in this thesis as a whole we test hypotheses through multiple channels, confirm the robustness of findings by adopting alternative methodologies, and try to maintain a generally conservative attitude, highlighting the shortcomings and limitations of each estimation strategy. We report several null results, which are informative about the uncertainty around our causal estimates and about the specific subsets of our studies in which significant findings arise. Only through careful practice in causal interpretation, indeed, can the credibility of policy analysis be ameliorated (Athey and Imbens, 2017) and ultimately, policymaking in human-environment systems improved.

## Summary of Findings

In this section, we provide a brief summary of the main findings for each chapter in the dissertation.

Chapter 1 essentially reports null results: while the synthetic control and its extensions identify a negative effect of the BC carbon tax on road transportation emissions, (i) the synthetic control result is not robust to traditional falsification procedures, and (ii) its extensions fail to identify a statistically significant effect. It is highly likely that the main analysis is substantially underpowered, and more investigation at a finer geographical scale will be needed in order to assess its impact on CO<sub>2</sub> emissions. Moreover, the tax has not led to statistically significant levels of carbon leakage, possibly due to a low signal-to-noise ratio which did not alter consumers' perceptions of pump price changes.

In Chapter 2, we exploit a disaggregated dataset, which is able to pick up a small, but statistically significant effect of the carbon tax on PM<sub>2.5</sub> emissions, providing empirical evidence about carbon tax co-benefits for the first time. However, the spatial distribution of co-benefits is heterogeneous, with less dense, richer areas benefitting disproportionately from the policy vis-à-vis denser and relatively poorer locations. Coupled with the mechanism underlying the result, i.e. commute mode switching towards public and active transport, the result highlights a secondary, spatial dimension of carbon tax regressiveness which needs to be taken into account when designing environmental policy. Moreover, monetary health gains due to mortality reductions

associated with the improve in air quality are also unevenly distributed across space, and co-vary with income, reinforcing the spatial regressiveness hypothesis.

Chapter 3 identifies an unintended consequence of the Colombian Covid-19 lockdown on the natural environment, manifesting through a statistically significant increase in cumulative forest fires for the whole country. This increase is highly correlated with the presence of armed groups on the Colombian territory, suggesting that the reductions in mobility may have created a perverse incentive to clear forest due to reduced governmental and community-based monitoring and enforcement.

Finally, Chapter 4 determines that the moratorium impact is variegated but small: on dryland forest, the policy resulted in a retention of at most 0.65% higher forest cover compared to comparable untreated dryland forest, while the effect is not statistically significant on carbon-rich peatland. Carbon emissions reductions range from 67.8 to 86.9 MtCO<sub>2</sub>-eq, implying a lower effective carbon price than what Norway paid Indonesia: at \$5/tCO<sub>2</sub>-eq, the payment should have been in the \$339-434.5 million range, compared to the \$56.2 million effectively disbursed. Moreover, the moratorium is projected to only contribute 3-4% towards Indonesia's Paris commitment of a 29% emissions reduction before 2030.

Read in conjunction, these results suggest that the policies we analyse in this thesis have been only mildly effective in obtaining their stated effect of mitigating global pollutants. Moreover, unintended effects from climate policy or from regulatory shocks are pervasive, and can manifest through perverse incentives: on the one hand, a carbon tax can have doubly regressive effects, over income and pollution distributions; on the other hand, a national security emergency can give rise to illegal exploitation of natural resources due to a reduction in monitoring and enforcement. In light of the necessity of preserving biodiversity during the 2021-2030 UN Decade for Ecosystem Restoration and of getting global emissions on a feasible trajectory to reach Net Zero in 2050, a greater level of ambition in climate policy and forest conservation will be required. Only through rigorous impact evaluations can the causal effect of regulation be identified. Policymakers will thus need to rely on empirically sound research in order to be able to rank interventions in cost-benefit analyses.

# References

- Abadie, Alberto (2021). “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects”. *Journal of Economic Literature* 59.2, pp. 391–425.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2015). “Comparative Politics and the Synthetic Control Method”. *American Journal of Political Science* 59.2, pp. 495–510.
- Abadie, Alberto and Gardeazabal, Javier (2003). “The Economic Costs of Conflict : A Case Study of the Basque Country”. *American Economic Review* 93.1, pp. 113–132.
- Aguilar-Gomez, Sandra, Dwyer, Holt, Graff Zivin, Joshua S, and Neidell, Matthew J. (2022). *This is Air: The "Non-Health" Effects of Air Pollution*. Working Paper 29848. National Bureau of Economic Research.
- Andam, Kwaw S., Ferraro, Paul J., Pfaff, Alexander, Sanchez-Azofeifa, G. Arturo, and Robalino, Juan A. (2008). “Measuring the effectiveness of protected area networks in reducing deforestation”. *Proceedings of the National Academy of Sciences* 105, pp. 16089–16094.
- Andersson, Julius J. (2019). “Carbon Taxes and CO2 Emissions: Sweden as a Case Study”. *American Economic Journal: Economic Policy* 11.4, pp. 1–30.
- Angrist, Joshua D. and Pischke, Jörn-Steffen (2010). “The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics”. *Journal of Economic Perspectives* 24.2, pp. 3–30.
- Arkhangelsky, Dmitry, Athey, Susan, Hirshberg, David A., Imbens, Guido W., and Wager, Stefan (2021). “Synthetic Difference-in-Differences”. *American Economic Review* 111.12, pp. 4088–4118.
- Athey, Susan, Bayati, Mohsen, Doudchenko, Nikolay, Imbens, Guido, and Khosravi, Khashayar (2021). “Matrix Completion Methods for Causal Panel Data Models”. *Journal of the American Statistical Association* 116.536, pp. 1716–1730.
- Athey, Susan and Imbens, Guido W. (2017). “The State of Applied Econometrics: Causality and Policy Evaluation”. *Journal of Economic Perspectives* 31.2, pp. 3–32.
- Baumol, William (1972). “On Taxation and the Control of Externalities”. *American Economic Review* 62.3, pp. 307–22.
- Ben-Michael, Eli, Feller, Avi, and Rothstein, Jesse (2021). “The Augmented Synthetic Control Method”. *Journal of the American Statistical Association* 116.536, pp. 1789–1803.
- Calonico, Sebastian, Cattaneo, Matias D., and Titiunik, Rocio (2014). “Robust non-parametric confidence intervals for regression-discontinuity designs”. *Econometrica* 82.6, pp. 2295–2326.

- Carattini, Stefano, Kallbekken, Steffen, and Orlov, Anton (2019). “How to win public support for a global carbon tax”. *Nature* 565.
- Chabé-Ferret, Sylvain and Subervie, Julie (2013). “How much green for the buck? Estimating additional and windfall effects of French agro-environmental schemes by DID-matching”. *Journal of Environmental Economics and Management* 65.1, pp. 12–27.
- Chandra, Ambarish, Head, Keith, and Tappata, Mariano (2014). “The Economics of Cross-border Travel”. *The Review of Economics and Statistics* 96.4, pp. 648–661.
- Currie, Janet, Voorheis, John, and Walker, Reed (2020). *What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality*. en. Tech. rep. w26659. Cambridge, MA: National Bureau of Economic Research, w26659. (Visited on 08/24/2022).
- Doudchenko, Nikolay and Imbens, Guido W. (2016). *Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis*. Working Paper 22791. National Bureau of Economic Research.
- Douenne, Thomas (2020). “The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy”. *The Energy Journal* 41.3, pp. 231–254.
- Douenne, Thomas and Fabre, Adrien (2022). “Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion”. *American Economic Journal: Economic Policy* 14.1, pp. 81–110.
- FAO (2021). *FAOSTAT*. URL: <http://www.fao.org/faostat/en/#home>.
- Ferman, Bruno and Pinto, Cristine (2021). “Synthetic controls with imperfect pretreatment fit”. *Quantitative Economics* 12.4, pp. 1197–1221.
- Ferraro, Paul J and Pattanayak, Subhrendu K (2006). “Money for Nothing? A Call for Empirical Evaluation of Biodiversity Conservation Investments”. *PLOS Biology* 4.4.
- Ferraro, Paul J., Sanchirico, James N., and Smith, Martin D. (2019). “Causal inference in coupled human and natural systems”. *Proceedings of the National Academy of Sciences* 116.12, pp. 5311–5318.
- Fowlie, Meredith, Rubin, Edward, and Walker, Reed (2019). “Bringing Satellite-Based Air Quality Estimates Down to Earth”. *AEA Papers and Proceedings* 109, pp. 283–88.
- Friberg, Richard, Steen, Frode, and Ulsaker, Simen A. (2022). “Hump-Shaped Cross-Price Effects and the Extensive Margin in Cross-Border Shopping”. *American Economic Journal: Microeconomics* 14.2, pp. 408–38.
- Government of Indonesia (2018). *Indonesia Second Biennial Update Report Under the United Nations Framework Convention on Climate Change*. Directorate General of Climate Change, Ministry of Environment and Forestry, Government of Indonesia report.
- Government of Norway (2010). *Letter of Intent on Cooperation on Reducing Greenhouse Gas Emissions from Deforestation and Forest Degradation*. Ministry of the Environment, Oslo, Norway.
- Gruber, Jonathan (1994). “The Incidence of Mandated Maternity Benefits”. *The American Economic Review* 84.3, pp. 622–641.

- Heckman, James J., Ichimura, Hidehiko, and Todd, Petra (1998). “Matching As An Econometric Evaluation Estimator”. *Review of Economic Studies* 65.2, pp. 261–294.
- Heckman, James J., Ichimura, Hidehiko, and Todd, Petra (1997). “Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme”. *Review of Economic Studies* 64, pp. 605–654.
- High-Level Commission on Carbon Prices (2017). *Report of the High-Level Commission on Carbon Prices*. World Bank, Washington, DC.
- Holland, Paul W. (1986). “Statistics and Causal Inference”. *Journal of the American Statistical Association* 81.396, pp. 945–960.
- Hotelling, Harold (1929). “Stability in Competition”. *The Economic Journal* 39.153, pp. 41–57.
- Jbaily, Abdulrahman, Zhou, Xiaodan, Liu, Jie, Lee, Ting-Hwan, Kamareddine, Leila, Verguet, Stéphane, and Dominici, Francesca (2022). “Air pollution exposure disparities across US population and income groups”. *Nature* 601.7892, pp. 228–233.
- Joppa, Lucas N and Pfaff, Alexander (2010). “Global protected area impacts”. *Proceedings of the Royal Society B* 278, 1633–1638.
- Krewski, D. *et al.* (2009). “Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality”. *Research Reports, Health Effects Institute*. Discussion 115-36.140, pp. 5–114.
- Lawley, Chad and Thivierge, Vincent (2018). “Refining the evidence: British Columbia’s carbon tax and household gasoline consumption”. *The Energy Journal* 39.2.
- Lee, David S. and Lemieux, Thomas (2010). “Regression discontinuity designs in economics”. *Journal of Economic Literature* 48.2, pp. 281–355.
- Lepeule, J., Laden, F., Dockery, D., and Schwartz, J. (2012). “Chronic exposure to fine particles and mortality: an extended follow-up of the Harvard Six Cities study from 1974 to 2009”. *Environmental Health Perspectives* 120.7, pp. 965–970.
- Li, Mingwei, Zhang, Da, Li, Chiao-Ting, Mulvaney, Kathleen M., Selin, Noelle E., and Karplus, Valerie J. (2018). “Air quality co-benefits of carbon pricing in China”. *Nature Climate Change* 8.5, pp. 398–403.
- Mildenberger, Matto, Lachapelle, Erick, Harrison, Kathryn, and Stadelmann-Steffen, Isabelle (2022). “Limited impacts of carbon tax rebate programmes on public support for carbon pricing”. *Nature Climate Change* 12.2, pp. 141–147.
- Murray, Brian and Rivers, Nicholas (2015). “British Columbia’s revenue-neutral carbon tax : A review of the latest ‘grand experiment’ in environmental policy”. *Energy Policy* 86, pp. 674–683.
- Nordhaus, William D (1977). “Economic Growth and Climate: The Carbon Dioxide Problem”. *American Economic Review* 67.1, pp. 341–346.
- Nordhaus, William (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Pigou, Arthur (1920). *The Economics of Welfare*. MacMillan and Co., London.
- Rivers, Nicholas and Schaufele, Brandon (2015). “Salience of carbon taxes in the gasoline market”. *Journal of Environmental Economics and Management* 74, pp. 23–36.

- Rubin, Donald B (2005). “Causal Inference Using Potential Outcomes”. *Journal of the American Statistical Association* 100.469, pp. 322–331.
- Sager, Lutz and Singer, Gregor (2022). “Clean Identification? The Effects of the Clean Air Act on Air Pollution, Exposure Disparities and House Prices”. *Grantham Research Institute on Climate Change and the Environment Working Paper No. 376*, p. 57.
- Schleicher, Judith, Eklund, Johanna, D. Barnes, Megan, Geldmann, Jonas, Oldekop, Johan A., and Jones, Julia P. G. (2020). “Statistical matching for conservation science”. *Conservation Biology* 34.3, pp. 538–549.
- Timilsina, Govinda R. (2022). “Carbon Taxes”. *Journal of Economic Literature* 60.4, pp. 1456–1502.
- UNFCCC (2008). *UNITED Report of the Conference of the Parties on its thirteenth session, held in Bali from 3 to 15 December 2007 Part Two : Action taken by the Conference of the Parties at its thirteenth session Decisions adopted by the Conference of the Parties*. Tech. rep., pp. 1–60.
- Vandyck, Toon, Keramidas, Kimon, Kitous, Alban, Spadaro, Joseph, Van Dingenen, Rita, Holland, Mike, and Saveyn, Bert (2018). “Air quality co-benefits for human health and agriculture counterbalance costs to meet Paris Agreement pledges”. *Nature Communications* 9.
- Weitzman, Martin L. (2015). “Internationally-Tradable Permits Can Be Riskier for a Country than an Internally-Imposed Carbon Price”. *SSRN Electronic Journal*.
- Weitzman, Martin (1974). “Prices vs. Quantities”. *Review of Economic Studies* 41.4, pp. 477–491.
- Woollacott, Jared (2018). “The Economic Costs and Co-benefits of Carbon Taxation: a General Equilibrium Assessment”. *Climate Change Economics* 9.1, pp. 1–22.
- World Bank (2022). *State and Trends of Carbon Pricing 2022*. World Bank, Washington, DC.
- Zhang, Wen-Wen, Zhao, Bin, Ding, Dian, Sharp, Basil, Gu, Yu, Xu, Shi-Chun, Xing, Jia, Wang, Shu-Xiao, Liou, Kuo-Nan, and Rao, Lan-Lan (2021). “Co-benefits of subnationally differentiated carbon pricing policies in China: Alleviation of heavy PM2.5 pollution and improvement in environmental equity”. *Energy Policy* 149, p. 112060.

## Part I

# The Impact of Carbon Pricing on Emissions

# Chapter 1

## Carbon Pricing with Permeable Borders: The Impact of British Columbia's Carbon Tax on Road Transportation CO<sub>2</sub> Emissions

### Abstract

This paper undertakes a quasi-experimental evaluation of British Columbia's 2008 carbon tax, with a focus on road transportation emissions. A naive application of the synthetic control method identifies emissions reductions of 0.47 metric tons per capita in an average year, or 13.8%, compared to a counterfactual constructed using all other Canadian provinces. A rigorous application of the synthetic control method finds instead a reduction in transportation emissions of 0.22 metric tons per capita in an average year, or 6.8%; however this result is not robust to traditional falsification procedures. Recent extensions of the synthetic control, such as synthetic difference-in-differences, de-meaned and ridge regularised synthetic controls, and matrix completion methods, all highlight a negative effect of the tax on road transportation emissions, albeit with large standard errors, thus not being able to identify a statistically significant estimate. The study also provides the first analysis of the impact of carbon pricing on cross-border fuel shopping in the transportation sector, estimating the impact of the 2008 carbon tax on retail sales in Whatcom county, Washington State, USA. A synthetic difference-in-differences analysis is unable to identify a positive shock on gasoline stations retail sales, decreasing the confidence on the hypothesis that the tax has resulted in significant carbon leakage. Rigorous applications of the synthetic control methodology and its extensions are essential in order to correctly inform the policymaker about the effectiveness of carbon pricing initiatives.

---

For helpful discussions and comments we thank, in alphabetical order: Mook Bangalore, Andrea Ciaccio, Eugenie Dugoua, Ben Groom, David Hendry, Charles Palmer, Elena Perra, Sefi Roth, Filippo Santi and all participants to research seminars hosted by the Department of Geography and Environment at the LSE in 2019, 2020 and 2021, and to the SASCA PhD Conference at Ca'Foscari University of Venice, 2022. The author acknowledges funding from the UK Economic and Social Research Council (ESRC). All remaining errors are my own.



## 1.1 Introduction

In 2019, the US Economists' Statement on Carbon Dividends and the European Economists' Statement on Carbon Pricing have once again highlighted the necessity to urgently correct the global market failure arising from carbon emissions, by either introducing carbon pricing outright (in the US), or coupling existing emission trading schemes (the EU ETS) with carbon taxes, targeted specifically towards the transportation and housing sectors<sup>1</sup>. Yet, empirical evidence on the effectiveness of carbon pricing is scarce, also due to their relative infrequency with respect to ETSs. Carbon taxes have indeed proven difficult to implement, *in primis* due to a lack of public support, which economists have proposed to overcome by redistributing prospective revenues to citizens (Carattini *et al.*, 2019). While this strategy has been found to attain limited effectiveness in practice (Mildenberger *et al.*, 2022), revenue neutrality has been a distinguishing feature of one of the most publicised carbon pricing schemes, Canadian province British Columbia's 2008 carbon tax (Murray and Rivers, 2015).

While recent research by Pretis (2022) has analysed the impact of the British Columbian carbon tax on aggregate emissions, finding negligible effects of the scheme in reducing CO<sub>2</sub> levels, this paper is concerned with estimating the impact of the tax on *road transportation* emissions. Notwithstanding the relatively broad coverage of the 2008 carbon tax, amounting approximately to 70% of provincial emissions, the most affected sector is indeed certainly transportation, which contributed to 43.9% of the province's total CO<sub>2</sub> levels in 2007, the year preceding the tax roll-out, making it an ideal candidate for an empirical evaluation. Road transportation has been deemed one of the most challenging sectors to decarbonise, due to the slow turnover of vehicle fleets and infrastructural and behavioural path dependency (IPCC, 2018). Moreover, other segments of the transportation sector in British Columbia (maritime shipping, air travel, locomotive fuel and fuel used in agricultural operations) were exempted from the 2008 carbon tax upon its implementation.

Previous studies have attempted to estimate the reduction in emissions registered in the BC transportation sector *indirectly*, by simulating emission reductions corresponding to estimated gasoline demand reductions (Rivers and Schaufele, 2015). We instead rely on the synthetic control method for comparative case studies (Abadie and Gardeazabal, 2003; Abadie *et al.*, 2010; Abadie *et al.*, 2015; Abadie, 2021) and on a set of recent extensions which enhance its robustness, such as synthetic difference-in-differences (Arkhangelsky *et al.*, 2021), demeaned synthetic controls (Doudchenko and Imbens, 2016; Ferman and Pinto, 2021) and matrix completion methods (Athey *et al.*, 2021), which allow us to estimate the *direct* effect of the 2008 carbon tax on road transportation CO<sub>2</sub> emissions, as pioneered in a recent paper by Andersson (2019). By avoiding simulations, we are able to account for several factors which are assumed away in estimates based on gasoline demand elasticities: fuel substitution between gasoline and diesel, suggestive evidence of which is found in Saberian (2017); long-term changes in fuel efficiency driven by vehicle fleet turnover

---

<sup>1</sup>These sectors are uncovered by the EU ETS.

(Antweiler and Gulati, 2016); and behavioural changes such as the modal switch to public transport, biking and walking. Moreover, with respect to the short-term analysis of Rivers and Schaufele (2015), with an endpoint in 2011, we augment the length of the post-intervention series to 2016, therefore shedding light on the evolution of the behavioural response to the carbon tax with a medium-term 8-year horizon. The synthetic control approach allows us to overcome issues found in previous analyses based on panel data and difference-in-differences (Metcalf, 2019), which are very likely to be biased due to violations of the foundational parallel trends assumption: by constructing a synthetic unit based on a weighted combination of untreated Canadian provinces, the method identifies a counterfactual British Columbia which reproduces the potential outcome that the actual unit would have experienced, had the carbon tax not been implemented in 2008. With respect to Pretis (2022), who also inspects the effect of the tax on transportation emissions in an ancillary analysis, we improve by restricting the focus to road transportation emissions only, thereby avoiding to include treated and untreated segments of the transportation sector in the evaluation and respecting the assumptions underlying thorough causal inference (Andersson, 2019); moreover, we provide the full set of placebo tests for synthetic controls, as prescribed by Abadie (2021), showing that a “naive implementation” of the synthetic control method does not survive the standard falsification routines, raising concerns about the lack of statistical power in this particular application. Through the introduction of synthetic difference-in-differences, de-measured synthetic controls and matrix completion methods for the first time in the literature on programme evaluation of carbon pricing, we confirm this hypothesis: placebo-based standard errors around the point estimate of the emissions reductions are indeed large, and suggest that while it is likely that transportation emissions reductions have arisen due to the 2008 carbon tax, neither the synthetic control method nor its extensions are able to clearly identify them when using the remaining Canadian provinces as the control pool. All methods produce negative estimates of transportation emissions effects following the carbon tax: the most conservative is identified via synthetic difference-in-differences and corresponds to an annual 6.1% reduction, or 0.21 tCO<sub>2</sub> per capita, in an average year. Notably, due to the large uncertainties connected with the lack of power, these estimates ought not to be considered true causal impacts of the policy.

Aware of the concerns due to violations of the stable unit treatment value assumption (SUTVA), arising due to potential carbon leakage towards the US state of Washington, as mentioned in Rivers and Schaufele (2015), Antweiler and Gulati (2016), and Andersson (2019), in the second part of the paper we calibrate a variant of the Friberg *et al.* (2022) model of cross-border shopping with parameters reflecting the British Columbia-Washington State, USA geographical region and the price gradient observed in their motor fuel markets. Theoretical predictions from the model assign to population distributions, price (and tax) differentials and to idiosyncratic consumers’ fixed costs a relevant role in shaping cross-border trade, on one hand confirming the possibility that cross-border fuelling may have arisen in response to the carbon tax, but dispelling it on the other due to minimal differences in real gasoline prices and on the high fixed costs of travelling across the border for refuelling. Empirically, we first estimate the influence of gasoline prices and exchange rates in determining border

crossings from British Columbia to Washington following [Chandra \*et al.\* \(2014\)](#), and confirm that the real exchange rate is the principal determinant of cross-border travel; secondly, we leverage taxable retail sales data from Washington’s Department of Revenue to perform a second synthetic difference-in-differences analysis at the Washington county level, in order to inspect whether retail sales in Whatcom county, bordering the most densely populated region of British Columbia, have increased following the 2008 carbon tax. With the exception of the food and beverage retail sector, we find negligible effects on retail trade trends; most importantly, retail sales within NAICS sector 447, corresponding to gasoline stations, do not seem affected by cross-border shopping behaviour by British Columbians.

The present study thus contributes to the literature along several dimensions. First, it provides a methodologically rigorous attempt at estimating the effect of the 2008 British Columbian carbon tax on road transportation emissions, expanding the evidence base from European economies to North American jurisdictions, which exhibit peculiarities in terms of higher average distances and heavier reliance on gasoline vis-à-vis diesel fuel but also lower prices and fuel taxes. Secondly, it highlights the shortcomings in which researchers may incur when applying the synthetic control method and its extensions without deploying the standard set of caveats and falsification exercises: previous analyses of the British Columbian carbon tax have indeed claimed to find a causal effect with respect to transportation emissions where instead a rigorous empirical analysis cannot establish causality for the registered 6.1% decrease. Finally, this paper presents the first empirical analysis of international carbon leakage in the transportation sector following the implementation of a carbon pricing scheme: due to the high incidence of fixed costs for private consumers and to the relatively low wedge in fuel prices introduced by the carbon tax in British Columbia, significant carbon leakage has not arisen in response to the tax. This is a reassuring result in light of the necessity of introducing carbon pricing in jurisdictions where it is politically feasible, which due to their geographic location may share land borders with countries whose fuel markets present an attractive price gradient.

The remainder of the paper proceeds as follows: [Section 1.2](#) provides background on the the 2008 British Columbian carbon tax, and reviews the burgeoning literature on the “poster child” of carbon pricing applications; [Section 1.3](#) describes the data used in the empirical analyses; [Section 1.4](#) introduces and explains the empirical methodology, describing the synthetic control method and its recent extensions; [Section 1.5](#) presents the results from the synthetic control method and its extensions; [Section 1.6](#) models cross-border shopping behaviour between BC and Washington, estimates the elasticity of cross-border travel to fuel prices and exchange rates and analyses carbon leakage in the context of the 2008 carbon tax; results are discussed and compared to the literature in [Section 1.7](#), and [Section 1.8](#) concludes the paper.

## 1.2 The 2008 British Columbian carbon tax

### 1.2.1 Background

The Canadian province of British Columbia first announced the introduction of its revenue-neutral carbon tax in February 2008, before eventually implementing it on July 1st, 2008 (Rivers and Schaufele, 2015). The tax was initially introduced at CAD 5/tCO<sub>2</sub> and projected to rise by CAD 5 each year; in 2012, however, the tax was frozen at CAD 30/tCO<sub>2</sub> until 2018 upon a change in local government.

A prominent feature of BC's carbon pricing scheme, at its onset, was revenue-neutrality: during the first four years of its implementation, close to 100% of the tax revenues were redistributed via personal income tax cuts for households in lower income brackets, low income tax credits, and corporate income tax cuts. After 2012, some of the revenues generated by the carbon price started to be earmarked for corporate tax cuts and credits in specific sectors (Murray and Rivers, 2015), thus departing from perfect neutrality and giving rise to a mixed system of redistribution. Revenue neutrality was a central feature of the tax proposal, as a mechanism to obtain public support and reduce hostility towards additional taxation, one of the main obstacles to the implementation of carbon pricing schemes (Carattini *et al.*, 2017; Carattini *et al.*, 2019); indeed, after the initial "Axe the tax" campaigns calling for the abrogation of carbon pricing ahead of the 2009 provincial elections, polling results showed a sustained increase in public approval of the tax until 2015 (Murray and Rivers, 2015). A recently published study exploiting data on the only two countries with ongoing tax and rebate schemes (Canada and Switzerland), however, pins the rates of approval and disapproval of carbon pricing to partisan identities rather than updated information about eventual rebates (Mildenberger *et al.*, 2022).

British Columbia, however, was not the only Canadian province to have introduced a tax on fossil fuel consumption at the time; indeed, in 2007 the province of Québec imposed an annual duty payable to the Green Fund, at the minimal rate of CAD 3 per tonne of CO<sub>2</sub> (Houle, 2013), or 0.8 cents per litre of gasoline. This levy fell on energy producers, with no intended pass-through to consumers at the pump, and received substantially less attention than its BC counterpart: it is not included as a carbon pricing scheme in the World Bank's State and Trends of Carbon Pricing annual report (World Bank, 2022). The duty was later subsumed in Québec's ETS system, which was approved in 2012 but became effective from 2013 onward, and only covered transportation emissions from 2015 (World Bank, 2018). The BC carbon tax had a substantially higher profile and received far more attention in the media, thereby often obtaining recognition as the first scheme of its kind in North America (Rivers and Schaufele, 2015).

### 1.2.2 Review of the literature

A handful of studies has attempted to estimate the impact of the BC carbon tax in reducing gasoline demand. Using a 2000-2011 panel dataset at the provincial

level, [Rivers and Schaufele \(2015\)](#) estimate that a five-cent increase in the carbon price results in a 8.4% decrease in gasoline consumption, which roughly translates to a 2.37 MtCO<sub>2</sub> reduction in emissions from road transportation over the first four years of the tax implementation. A similar aggregate analysis is performed by [Antweiler and Gulati \(2016\)](#), who instead find gasoline demand reductions in the 1-7% range between 2001 and 2014. [Lawley and Thivierge \(2018\)](#) employ a 2001-2012 panel of household-level data, finding 5-8% reductions in gasoline demand stemming from geographically heterogeneous responses to the carbon tax: residents of large cities exhibit a strong tax elasticity (around 10-12%), while inhabitants of rural and northern areas do not significantly respond to the tax. A common potential shortcoming across these studies is the shortness of the post-intervention time series employed in the empirical analyses. While it is important to evaluate the short-term behavioural response to carbon pricing, which translates into a reduction in gasoline demand, consumers adjust to the introduction of a carbon tax on both the intensive and extensive margin, as [Antweiler and Gulati \(2016\)](#) highlight; a longer post-treatment time frame could then shed light on the effect of pricing on the renewal of the car fleet and on fuel substitution strategies. Moreover, the impacts found in the three aforementioned studies, with the possible exception of [Antweiler and Gulati \(2016\)](#), who indeed find smaller and less precisely estimated effects, could be contaminated by the aftermath of the 2007-2009 global financial crisis, which may have had asymmetric repercussions across Canadian provinces. Another inflection point is represented by the decrease in the global crude oil price after 2014, which has entailed a consequent decrease in local gasoline prices in BC, and a rebound effect in terms of total gasoline consumption ([Arcila Vasquez and Baker, 2022](#)).

Two recently published studies relate to this investigation more directly, [Pretis \(2022\)](#) and [Arcila Vasquez and Baker \(2022\)](#). [Pretis \(2022\)](#) primarily analyses the BC carbon tax's impact on the province's aggregate emission levels, using traditional difference-in-differences and synthetic control methods and a novel break-detection approach, and failing to find an effect on BC's total emissions. He then proceeds to disaggregate the overall effect into its sectoral components, finding an average 5% significant reduction in emissions in the transportation sector alone. The absence of any effects in remaining sectors seems to be in contrast with [Xiang and Lawley \(2019\)](#) who, employing synthetic controls, identify a significant reduction in natural gas demand in BC following the introduction of the tax. However, given the imperfect fit between BC and its synthetic control in both studies, and especially in the sectoral analysis in [Pretis \(2022\)](#), it is difficult to establish a preferential order between the two contrasting results. For what concerns transportation, given the primary importance of the sector in BC's emissions mix<sup>2</sup>, it is surprising that a significant reduction in road CO<sub>2</sub> emissions would not result in reductions in the province's aggregate trajectory; however it must be noted that the estimates contained in [Pretis \(2022\)](#) are not robust to traditional inference for synthetic controls, including those obtained for the transportation sector<sup>3</sup>. This result is partially confirmed in the

---

<sup>2</sup>In 2007, the year before the implementation of the carbon tax, transportation emission accounted for 43.9% of the total CO<sub>2</sub> emissions for the province.

<sup>3</sup>Limited robustness and significance are probably also an artifact of the small size of the control pool,

recent paper by [Arcila Vasquez and Baker \(2022\)](#), which consists in a re-evaluation of the effectiveness of the tax on aggregate CO<sub>2</sub> emissions and gasoline consumption employing the synthetic control method, finding effects which “stand out” with respect to the existing literature, namely an increase in aggregate emissions and gasoline consumption. While they also leverage the similarities between Canadian provinces and US states by including the latter in their control pool, their estimations are very likely flawed by the less-than-ideal fit between their synthetic controls and BC, therefore failing to lend utmost credibility to their findings.

## 1.3 Data

### 1.3.1 Effect of the 2008 carbon tax on CO<sub>2</sub> emissions

In order to analyse the effect of the implementation of the 2008 British Columbian carbon tax on emissions, we construct an annual panel dataset spanning the years 1990-2019. we obtain road transportation, total transportation, and overall CO<sub>2</sub> emissions data for all Canadian provinces from [Statistics Canada \(2021a\)](#). We then calculate per capita emissions profiles by dividing gross emissions data by provincial and state population, obtained from [Statistics Canada \(2021d\)](#), thereby making emissions paths comparable across observational units.

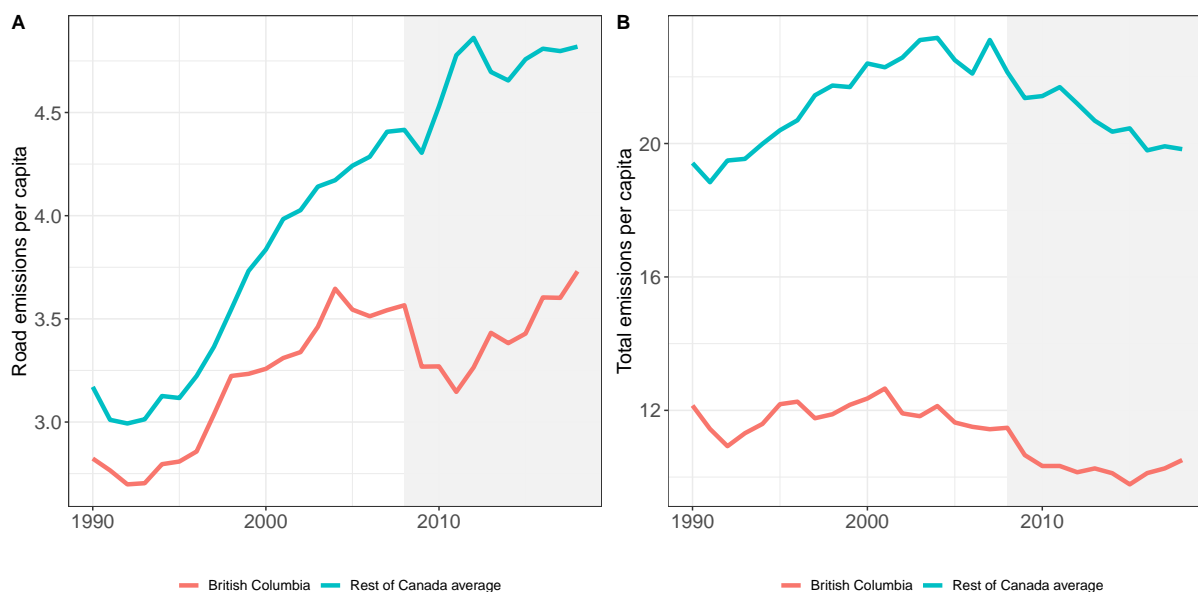
Trends of per capita road transportation emissions (panel A) and total emissions (panel B) for British Columbia and the rest of Canadian provinces are reported in [Figure 1.1](#), where the period of the implementation of the British Columbian carbon tax is shaded in grey.

Noticeably, while BC’s road transport emissions follow a similar trajectory to the average of other Canadian provinces up to 1998, it is evident how the trends start to diverge thereafter. BC’s road transport CO<sub>2</sub> experiences a sharp drop from 2008-2011, coinciding with the first three years of the implementation of the carbon tax, a decline which is not observed in other Canadian provinces. It is also noteworthy to observe how, after 2011, the series of road transportation CO<sub>2</sub> emissions for BC starts growing again at a rate comparable to the pace observed between 1991 and 2005, perhaps owing to the post-financial crisis economic recovery.

The post-2005 divergence in trends between BC and the rest of Canadian provinces is also observed in the paths of total CO<sub>2</sub> emissions, although less markedly; moreover, the drop in emissions registered after the implementation of the 2008 carbon tax is less pronounced and seemingly parallel to that observed in the rest of Canada, suggesting a modest incidence of the reductions in transportation emissions in driving the overall emission profile for the province, notwithstanding the large share of total emissions owed to the transportation sector in BC.

---

even though [Pretis \(2022\)](#) maximises it by including the provinces of Québec and Saskatchewan, whose transportation emissions trends are likely affected by another carbon tax in the first case, and an anomalous increase in high-emitting vehicles in the second - see [Section 1.3](#).



**Figure 1.1:** (A) Trends in road transportation and (B) total CO<sub>2</sub> emissions for British Columbia, an average of all other Canadian provinces and an average of all US States between 1990 and 2017.

In order to aid the performance of the synthetic control estimator, we follow [Andersson \(2019\)](#) in obtaining data for a range of predictors that the literature deems determinant for predicting emission trends of the transportation sector: (1) Annual sales of transportation fuels, namely motor gasoline and diesel, from [Statistics Canada \(2021e\)](#); (2) Size of the vehicle fleet, both in aggregate and by vehicle weight class<sup>4</sup> from [Statistics Canada \(2021g\)](#); (3) Gross Domestic Product and GDP sectoral splits at the NAICS-2 digit level, which allow me to account for the economic sectors more likely to contribute to transportation and overall emissions (e.g. agriculture, mining, construction, and manufacturing) from [Statistics Canada \(2021b\)](#); (4) Urban population shares and population density from [Statistics Canada \(2021d\)](#). All variables are aggregate at the annual and provincial/state level for the maximum possible length of the available time series.

In [Section 1.5](#), following [Abadie and Gardeazabal \(2003\)](#), [Abadie et al. \(2010\)](#), [Abadie et al. \(2011\)](#), and [Abadie et al. \(2015\)](#), we exclude from the panel those observational units which have experienced shocks to their emissions profiles, such as the introduction of gasoline taxes, the participation in Emission Trading Schemes or large idiosyncratic movements of their emission paths. We therefore exclude from the Canadian province panel: (1) Québec, due to the introduction, in 2007, of the “annual duty payable to the Green Fund”, albeit small in magnitude at CAD 3/tCO<sub>2</sub> (approximately 0.8 cents per litre of fuel, see [Section 1.2](#)) and probably irrelevant in driving the province’s emission reductions ([Houle, 2013](#)); (2) Saskatchewan, which experienced a sharp increase in road transportation emissions starting in 2006,

<sup>4</sup>Light vehicles < 4500 Kg, Medium vehicles > 4500 Kg and < 15000 Kg, and Heavy vehicles > 15000 Kg.

possibly explained by a sudden spike in heavy vehicle registrations (Dolter, 2016); (3) Yukon, Nunavut and the Northwest Territories, due to incomparably high and unstable emissions profiles, given the relatively much lower population base and density. Moreover, data series for Nunavut and the Northwest Territories are only distinct from 1999 onward, curtailing the pre-treatment period to only 10 years. Given Abadie and Gardeazabal (2003), Abadie *et al.* (2010), Abadie *et al.* (2011), and Abadie *et al.* (2015)'s recommendations about the necessity of maximising the length of the pre-intervention data series in order to augment the resemblance of the synthetic control unit to its treated counterpart, we opt to lose in granularity what we gain in terms of goodness of fit.

### 1.3.2 Carbon leakage to Washington State, USA

The determinants of eventual carbon leakage to the US state of Washington are analysed using monthly data at the British Columbian province level. We define border crossings as the monthly number of Canadian vehicles returning to Canada from the US<sup>5</sup>, and we classify the data into four categories: total crossings, total crossings made by automobiles, same-day trips and trips spanning two or more days. Further, we obtain gasoline price data from Kalibrate (formerly Kent Group Ltd.) at the monthly level for the city of Vancouver, which we consider representative of the entire province; exchange rate data is retrieved from the Pacific Exchange Rate Service at University of British Columbia's Sauder School of Business. Monthly income and unemployment rate data are extracted from Statistics Canada (2021f). All data is collected for the 1990-2019 period. We also construct a yearly panel dataset for the US state of Washington, employing data webscraped from the Washington Department of Revenue. We obtain county-level information on taxable retail sales between 1994 and 2019 for multiple industries, at the three-digit level defined by the NAICS classification system. In particular, we retrieve data for NAICS sectors: 441 (gasoline stations and convenience stores with pumps); 447 (automobile dealers); 443 (food and beverage), and for the entire NAICS retail trade sector (codes 44-45 at the two-digit level).

---

<sup>5</sup>Since Statistics Canada (Statistics Canada, 2021c) only records data for vehicles crossing the Canada-US border in the inward direction.



## 1.4 Empirical Strategy

### 1.4.1 Synthetic Control Method

#### Description

The synthetic control method (SCM) is a procedure that synthesises a control group based on a weighted convex combination of outcomes and explanatory variables of the “donor regions” which reproduces the variables of interest of the treated region in the pre-treatment period. By construction, the comparison units’ outcome closely match the treated unit’s outcome trend until the treatment, and if the intervention is successful the two paths diverge in the expected direction from the moment the treatment is implemented. In the following, we describe the SCM, drawing from [Abadie \*et al.\* \(2010\)](#).

Let units  $j = \{1, \dots, J + 1\}$  be the Canadian provinces. These units are observed for a range of time that goes from  $t = 1$  to  $t = T$ . To facilitate the exposition, we consider the first region as the British Columbia, the one where the intervention is implemented. Hence,  $j = 1$  if  $j$  denotes British Columbia. All the other Canadian provinces,  $j = \{2, \dots, J + 1\}$  will be our control units, which contribute to construct the synthetic control. Henceforth, we deem the set of control units “donor pool”. [Abadie \*et al.\* \(2010\)](#) assume the treated unit is continuously affected by the intervention after its implementation.

Denote the intervention period  $T_0 + 1$ . Pre-intervention periods are then  $t = \{1, 2, \dots, T_0\}$ , while post-intervention periods are  $t = \{T_0 + 1, \dots, T\}$ . In our model, which spans from 1990 to 2017,  $t = 1990, \dots, 2017$ , and  $T_0 + 1 = 2008$ .

Let us define two outcomes of interest: (1)  $Y_{jt}^N$  is the outcome that would be observed for unit  $j$  at time  $t$  if unit  $j$  is not exposed to the intervention; (2)  $Y_{jt}^I$  is the outcome that would be observed if unit  $j$  is exposed to the treatment.

The model assumes that the treatment has no effect on the outcome variable of interest before of the implementation of the intervention, meaning that the equality  $Y_{jt}^N = Y_{jt}^I$  must hold for all  $t = \{1, \dots, T_0\}$  and for all the units  $j = \{1, \dots, N\}$ . In other words, we set two strong assumptions: (1) No temporal spillovers (e.g. anticipation effect); (2) No spatial spillovers (i.e. no interference and indirect influences between units).

Our objective is to estimate the effect of the intervention on the outcome variable for the treated unit in the post-intervention period, comparing this result with the outcome of the synthetic control. The magnitude of the effect is defined as the estimated difference between the two potential outcomes  $\alpha_{1t} = Y_{jt}^I - Y_{jt}^N$  for all the post-intervention periods,  $T_0 + 1, \dots, T$ , where  $1 \leq T_0 \leq T$ . For construction,  $Y_{jt}^N$  is unobservable for the treated unit in the post-intervention period, as the treated unit has certainly received the treatment. The SCM aims to construct a synthetic counterfactual which estimates this unobservable outcome for the period following

the intervention (Abadie *et al.*, 2011).

### Implementation of the Model

The SCM constructs a synthetic British Columbia as a weighted combination of donor Canadian provinces: for  $j = \{2, \dots, J + 1\}$ , there exists a positive  $(J \times 1)$  vector  $W$  of weights which sum to one,  $W = (w_2, \dots, w_{J+1})^T$  with  $0 \leq w_j \leq 1$  and  $w_2 + \dots + w_{J+1} = 1$ . When the vector  $W$  varies, the weighted combination of control countries varies, creating an unique synthetic control.

The SCM aims to find the optimal vector of weights,  $W^*$  such that the synthetic control closely matches the treated units on several pre-treatment predictors of the outcome variable. If  $Z_j$  is the vector of predictors for each observational unit in our sample, the SCM algorithm calculates  $W^*$  such that  $W^* = (w_2^*, w_3^*, \dots, w_{J+1}^*)$  equalises the following sums for  $t = 1, 2, \dots, T_0$ :

$$\sum_{j=2}^{J+1} w_j^* Y_{j1} = Y_{11}, \sum_{j=2}^{J+1} w_j^* Y_{j2} = Y_{12}, \dots, \sum_{j=2}^{J+1} w_j^* Y_{jT_0} = Y_{1T_0}, \quad \text{and} \quad \sum_{j=2}^{J+1} w_j^* Z_j = Z_1$$

i.e. the convex combination of the outcome variable for the control units (the Synthetic Control) matches the outcome variable for the treated unit in every period up to the intervention period, and that the convex combination of the predictors for the control units matches the predictors for the treated units (Abadie *et al.*, 2010; Andersson, 2019). However, this procedure rests on a strong assumption: the constraint on weights to be non-negative implies that exact balance may be achieved if and only if the treated unit's outcomes and predictors lie within the convex hull of those for the control units (Ben-Michael *et al.*, 2021; Abadie *et al.*, 2010).

If this condition holds, Abadie *et al.* (2010) prove that  $\alpha_{1t}$  can be estimated for the post-treatment period using the unbiased estimator:

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}$$

Where the first term on the right hand side is the outcome variable of the treated region and the second term is the outcome variable of the synthetic control, expressed as a convex combination of the outcome variables of the donor units.

To select the optimal vector of weights  $W^*$ , we need to minimise the ‘‘pseudo-distance’’ (Abadie *et al.*, 2011) between British Columbia and the Synthetic British Columbia. Let  $X_1$  be a  $(k \times 1)$  vector of pre-intervention predictors and outcomes for British Columbia, defined as  $X_1 = (Z_1^T, Y_{11}, \dots, Y_{1T_0})^T$ . Here,  $Z_1^T$  is a transposed vector of predictors, as defined above, and  $Y_{1i}$ , for  $i = 1, \dots, T_0$  is the sequence of outcome variables for British Columbia.

Define  $X_0$  as a  $(k \times J)$  matrix which replicates  $X_1$  for all of the control regions. The SCM selects  $W^*$  to minimise the distance  $\|X_1 - X_0W\|$  for the pre-intervention period, subject to the constraints defined for the individual weights  $w_2, \dots, w_{J+1}$ .

Following [Abadie \*et al.\* \(2010\)](#); [Abadie \*et al.\* \(2011\)](#), let  $V$  be a  $(k \times k)$  symmetric and positive semidefinite matrix which measures the distance between  $X_1$  and  $X_0W$ , defined by the following expression:

$$\|X_1 - X_0W\|_v = \sqrt{(X_1 - X_0W)^T V (X_1 - X_0W)}$$

A matrix  $V$  defined as such<sup>6</sup> minimises the mean square prediction error (MSPE) ([Abadie \*et al.\*, 2015](#)) of  $Y_j$  over  $t = \{1, \dots, T_0\}$ , i.e.

$$MSPE = \frac{1}{T_0} \sum_{t=1}^{T_0} \left( Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt} \right)^2$$

### 1.4.2 Extensions of the Synthetic Control Method

We first introduce a recently developed estimation method for comparative case studies where a single unit is exposed to treatment: the synthetic difference-in-differences (hereafter, SDID) estimator proposed by [Arkhangelsky \*et al.\* \(2021\)](#). In the BC context, the traditional synthetic control method (SCM), of which we have discussed case-specific pitfalls and limitations, likely dominates traditional difference-in-differences (DID) applications, due to the fact that a single observational unit is “administered” the treatment, and that CO<sub>2</sub> emissions trends prior to treatment are not parallel among BC and control Canadian provinces (see [Section 1.3](#)). As it is argued in [Arkhangelsky \*et al.\* \(2021\)](#), SDID also dominates DID in applications of this kind, due to its reliance on the inclusion of unit and time weights in the regression function, which effectively “localise” the two-way fixed effects (TWFE) regression, by giving more importance to units whose pre-treatment outcomes are more similar to the treated unit, and to time periods which are on average comparable to the treated periods. SDID is also likely to dominate SCM in applications where SCM has so far been considered the preferred strategy in recent empirical practice ([Athey and Imbens, 2017](#)): by introducing time weights in the TWFE regression, SDID is able to modulate the effect of pre-treatment time periods which are distant from post-treatment time periods, thereby increasing precision and removing bias; moreover, by accounting for unit fixed effects as in DID, the novel estimator can remove some of the bias due to unexplained variation in outcomes unaccounted for by the SCM<sup>7</sup>.

<sup>6</sup>The statistical package `synth`, developed by [Abadie \*et al.\* \(2011\)](#) estimates  $V$  through a data-driven procedure, which allows to place non-identical weights on the predictors  $Z_j$  in order to assess their relative contribution to the evolution of the outcome variable  $Y_{jt}$  over time.

<sup>7</sup>Notably, when the pre-treatment weighted average of the outcome for the control units is identical to the outcome path of the treated unit, the SCM is able to account for the role of unit fixed

More specifically, [Arkhangelsky et al. \(2021\)](#) consider a balanced panel of  $N$  units and  $T$  time periods, where  $Y_{it}$  is the outcome for unit  $i$  in time period  $t$ , and  $N$  is partitioned such that  $(N_1, \dots, N_{co})$  are units which never experience treatment during all time periods  $(T_0, \dots, T)$ , and  $(N_{tr}, \dots, N)$  are exposed to treatment after time  $T_{pre}$ , with  $T_0 < T_{pre} < T$ .

In the BC case, the SDID estimator,  $\hat{\tau}^{SDID}$  is constructed by finding unit weights  $\hat{\omega}^{SDID}$  such that pre-2008 outcome trends for BC's transportation CO<sub>2</sub> emissions are aligned with those of control units:

$$\sum_{i=1}^{N_{co}} \hat{\omega}^{SDID} Y_{it} \approx \frac{1}{N_{tr}} \sum_{i=N_{tr}}^N Y_{it} \quad \forall t = 1, \dots, T_{pre} \quad (1.1)$$

And time weights,  $\hat{\lambda}_t^{SDID}$ , which align pre-2008 and post-2008 time periods. The average causal effect of the 2008 carbon tax on transportation CO<sub>2</sub> emissions,  $\tau_{tax}$ , is then estimated via the following TWFE regression:

$$(\hat{\tau}_{tax}^{SDID}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \underset{\tau, \mu, \alpha, \beta}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau)^2 \hat{\omega}_i^{SDID} \hat{\lambda}_t^{SDID} \right\} \quad (1.2)$$

Whereby, in comparison with standard TWFE-DID, the unit and time weights  $\hat{\omega}_i^{SDID}$  and  $\hat{\lambda}_t^{SDID}$  are added as a ‘localisation’ procedure for the regression, and in comparison with SCM, the unit fixed effects  $\alpha_i$  and the time weight  $\hat{\lambda}_t^{SDID}$  are added in order to remove some of the bias remaining from the SCM procedure, as discussed above. It is notable that the procedure in which the unit weights  $\hat{\omega}_i^{SDID}$  are chosen<sup>8</sup> differs from the one employed in [Abadie et al. \(2010\)](#) due to the presence of an intercept term,  $\omega_0$ , and a regularisation parameter following [Doudchenko and Imbens \(2016\)](#). In particular, the introduction of the intercept term allows for unit weights to be less stringent than unit weights employed in traditional SCM applications, since the SCM necessary condition of closely matched pre-treatment trends relaxes to the SDID sufficient condition of parallel pre-treatment outcomes. This additional flexibility is enabled by the introduction in the TWFE regressions of the unit fixed effects  $\alpha_i$ , which sweep constant differences between units. In the BC context, where the province has experienced a particularly idiosyncratic trend which is unlikely to be perfectly reproduced by a combination of control units, also due to the relative lack of available controls<sup>9</sup>, this improvement is likely to play a substantial role in improving upon SCM; moreover, the use of relatively more disperse weights avoids reliance on a very concentrated number of units chosen from a relatively small control pool.

---

effects; however, this is rarely the case in practical applications, in which the pre-treatment paths of the treated and control units are closely matched but not exactly equal, and specifically in the BC case, where we will show (see [Section 1.5](#)) how the fit between the treated and synthetic unit is probably plagued by problems of interpolation bias.

<sup>8</sup>In order not to be excessively redundant with [Arkhangelsky et al. \(2021\)](#), we omit the procedure used to estimate unit and time weights for SDID.

<sup>9</sup>See the discussion of inclusion/exclusion of control Canadian provinces in [Section 1.3](#).

Furthermore, the implementation of SCM and SDID in [Arkhangelsky \*et al.\* \(2021\)](#) allows to calculate confidence intervals around our estimates, by leveraging the “placebo variance estimator” ([Arkhangelsky \*et al.\*, 2021, p. 4110](#)). In each iteration  $b = 1, \dots, B$ , each control unit is sampled without replacement to receive the placebo treatment, and the SCM and SDID estimator  $\hat{\tau}^{(b)}$  is calculated. The variance is then defined as:

$$\hat{V}_{\tau}^{placebo} = \frac{1}{B} \sum_{b=1}^B \left( \hat{\tau}^{(b)} - \frac{1}{B} \sum_{b=1}^B \hat{\tau}^{(b)} \right)^2 \quad (1.3)$$

With a limited number of control units, and especially when the fit between the placebo synthetic units and the placebo treated units is less than perfect,  $\hat{V}_{\tau}^{placebo}$  will be large, due to an underpowered analysis.

In addition to SDID, we rely on other extensions of the synthetic control method mentioned in [Abadie \(2021\)](#), such as the demeaned synthetic control by [Doudchenko and Imbens \(2016\)](#) and [Ferman and Pinto \(2021\)](#), the matrix completion method of [Athey \*et al.\* \(2021\)](#), and ridge-regularised version of the synthetic control and demeaned synthetic control estimator, in order to provide a wider base on which to evaluate the effect of the 2008 carbon tax on road transportation emissions.

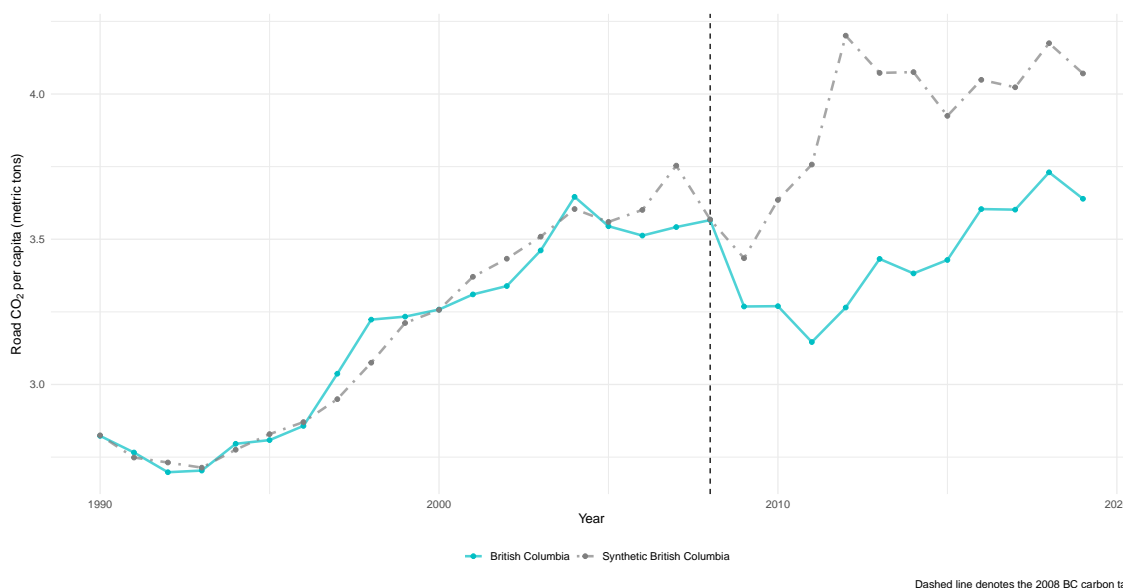
## 1.5 Results

### 1.5.1 Synthetic control method: naive implementation

#### Assessment of Fit

The first requirement for correct identification of the synthetic control method is that the outcome path of the synthetic control unit must be able to adhere to the outcome path of the treated unit. In [Figure 1.2](#), this condition seems to be respected: road transportation CO<sub>2</sub> emissions from the treated unit and the synthetic control differ on average by 0.052 metric tons per capita in absolute value during 1990-2008, a 1.6% average divergence. Nonetheless, due to the relative shortness of the pre-intervention time series, which probably fails to smooth some of the outcome variable’s volatility, the fit is less than ideal from 1996 onwards, an issue also observed, but not discussed, in [Pretis \(2022\)](#). It is important to note, however, that Synthetic BC approximates the treated unit well with respect to key predictors of road transportation CO<sub>2</sub> emissions, as evidenced in [Table 1.1](#). The predictors’ values, averaged in all cases across all available years except for outcome variable lags<sup>10</sup>, all move closer to British Columbia’s, underpinning a vastly ameliorated fit with respect to the standard average of donor pool provinces. The fit is slightly less ideal, albeit improved, for what concerns gasoline consumption per capita, and the share of light and heavy vehicles per thousand people. These predictors indeed receive less weight in the data-driven

<sup>10</sup>1990-2008 for gasoline sales per capita; 1997-2008 for GDP; 1999-2008 for light and heavy vehicles per thousand people; and 1990-2008 for population density.



**Figure 1.2:** Path plot of road transportation CO<sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia

estimation of the  $V$  matrix of covariates' weights, which identifies GDP per capita and CO<sub>2</sub> emissions lags as the main determinants of British Columbia's outcome path.

**Table 1.1:** Covariate balance between British Columbia and synthetic British Columbia

Variable	British Columbia	Synth BC	Donor Pool	$V$ weight
Gasoline per capita	1081.72	1154.97	1321.39	0.00
GDP per capita	40302.13	40293.20	45736.34	0.20
Heavy vehicles (x1000)	23.85	19.50	33.41	0.00
Light vehicles (x1000)	563.47	531.85	587.40	0.00
Population density	4.15	4.17	7.58	0.04
Road CO <sub>2</sub> per capita 1990	2.82	2.82	3.62	0.30
Road CO <sub>2</sub> per capita 2000	3.26	3.26	4.01	0.21
Road CO <sub>2</sub> per capita 2008	3.57	3.57	4.85	0.25


Table 1.2 reports instead the weights computed for the  $W$  matrix, i.e. pertaining to donor provinces: BC is reproduced via a weighted combination of Manitoba, Ontario, New Brunswick and Newfoundland and Labrador, which is strikingly different than the one reported by Pretis (2022) for aggregate CO<sub>2</sub> emissions. It is reassuring to see that Québec, the other Canadian province with a carbon tax scheme in place since 2007, obtains near-zero weight; however, it must be noted that its inclusion could still give rise to interpolation biases (see Section 1.4). Reporting predictor and unit weights is one of the characteristics of the traditional SCM which makes it attractive for comparable case studies, since it allows for a transparent interpretation of the counterfactual unit (Abadie, 2021); in this case, assessing the similarity of the

**Table 1.2:** Composition of synthetic British Columbia

Province	$W$ Weight
Alberta	0.00
Manitoba	0.69
New Brunswick	0.06
Newfoundland and Labrador	0.04
Nova Scotia	0.00
Ontario	0.21
Prince Edward Island	0.00
Quebec	0.00
Saskatchewan	0.00
Yukon	0.00

positive-weighted control units to the treated unit is arguably less of a concern, since the analysis is performed at a sub-national scale rather than at a cross-country level, thus restricting inherent socio-economic and cultural differences to a minimum. A potential remaining concern, if any, is the high weight assigned to Manitoba, which constitutes almost 70% of the synthetic control (90% with Ontario): such a heavy reliance on two donor provinces could be problematic if either of them is singularly driving the results.

## Results



**Figure 1.3:** Gap plot of road transportation CO<sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia

Figure 1.3 highlights the divergence between synthetic BC and the treated unit

reported in [Figure 1.2](#). In 2019, the last year in the sample<sup>11</sup>, the measured reduction in road transportation emissions per capita amounts to 0.43 metric tons per capita, 11.8% lower than if the carbon tax had not been introduced. The end-of-sample effect is smaller than the 0.47 metric tons per capita average effect (13.8%), probably due to a post-2011 rebound in the trend of road transportation emissions, coinciding with the period in which the tax was frozen at a rate of CAD 30/tCO<sub>2</sub> due to political reasons. Another possible reason for the levelling of emissions reductions is due to a sharp decrease in crude oil prices starting in 2014, which more than compensated the level of the carbon tax; notwithstanding the much higher tax salience vis-à-vis prices ([Rivers and Schaufele, 2015](#)), crude oil price volatility might have dampened the effect of the tax for consumers. It is also noteworthy that the measured effect is larger than the results found by [Andersson \(2019\)](#) for Sweden with a much larger carbon price; even though there exist large differences in gasoline and diesel prices between North America and Europe, thereby making a lower tax in Canada similar in incidence with respect to a higher levy in Europe, the magnitude of the result is striking.

### Standard Placebo Tests


Rigorous implementations of the synthetic control method are often accompanied by a suite of placebo tests, which assess the robustness of the results with respect to relevant characteristics (e.g. the composition of the actual donor pool or the set of predictors). Several recent studies fail to fully address the issue of the stability of their results to these types of perturbations ([Pretis, 2022](#)) or violate the standard requirements for causal effect identification using synthetic controls ([Arcila Vasquez and Baker, 2022](#)), while others identify a credible estimate via a careful inspection of potential biases in the study design ([Andersson, 2019](#)). Adhering to this latter tradition, and aware of the methodological suggestions brought forward by synthetic control proponents ([Abadie, 2021](#)), we discuss the robustness of the naive SCM implementation via four standard placebo tests: (i) an “in-time placebo”, where the carbon tax is backdated to 2002; (ii) an “in-space placebo”, also described in the literature as a uniform permutation test; (iii) the mean squared prediction error (MSPE) ratio test; (iv) a “leave-one-out” test.

[Figure 1.4](#) provides the first: the carbon tax is assigned to British Columbia in 2002, and the data-driven optimisation procedure to find  $V$  and  $W$  is also curtailed in time. Notably, no divergence in trends between the synthetic and treated unit arises right after the placebo treatment, but the two series only diverge post-2008 as in the original implementation, consistently with [Abadie \(2021\)](#). The emergence of a large placebo effect after 2002 would have cast doubts over the ability of the synthetic control estimator to provide a robust counterfactual to the actual treated BC.

The uniform permutation “in-space” placebo, shown in [Figure 1.5](#), consists in a reassignment of the treatment to each unit in the donor pool (without replacement),

<sup>11</sup>As explained in [Section 1.4](#), the sample must be restricted to 2016 if Alberta is kept in the control pool, since the province implemented its CAD 20/tCO<sub>2</sub> carbon tax in January 2017.



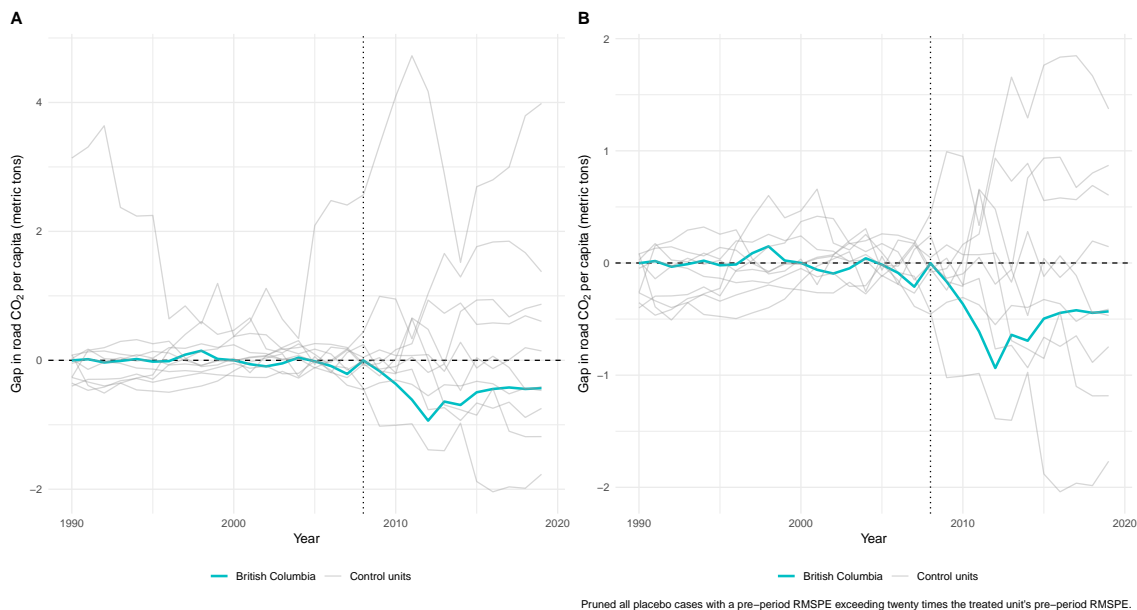


**Figure 1.4:** Path plot of road transportation CO<sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia, in-time placebo with treatment in 2002

and an estimation of the same synthetic control model for each control unit as if they were treated. This procedure<sup>12</sup> allows to evaluate whether the effect estimated by the synthetic control method for the treated unit lies in a tail of the distribution of the placebo effects estimated for control units, and to calculate  $p$ -values; however, it relies on the ability of the synthetic control method to produce credible counterfactuals for untreated as well as the treated units. In panel A, the gap between the treated and synthetic unit is reported for all provinces in the panel; panel B prunes observation whose pre-treatment RMSPE is 20 times higher than British Columbia's. Nonetheless, the placebo synthetic controls almost never attain a similar pre-treatment fit to that of BC: only the pre-treatment RMSPE for Manitoba and Nova Scotia have the same order of magnitude, thereby highlighting potential issues in terms of specific predictor or outcome lag choices (the  $V$  matrix) singularly determining the choice of the synthetic control.


Moreover, the magnitude of the effect identified for BC does not appear to be a particular outlier with respect to the other units, which also show explosive trends (either positive or negative) after 2008, often higher in absolute terms. The MSPE ratio test, visualised in Figure 1.6 identifies BC as the province with the highest post-treatment to pre-treatment MSPE ratio, and hence a Fisher's exact  $p$ -value of 0.0909 (1/11, where 1 is BC's MSPE ratio rank and 11 is the number of units in the panel) associated with the probability of estimating a gap of this magnitude between the treated and control unit. These seemingly contradictory results are reconciliated by the fact that the MSPE ratio test inherently depends on the ability of the synthetic control methodology to calculate a plausible placebo for the control

<sup>12</sup>Pretis (2022) only performs a uniform permutation test among the suite of available robustness checks for synthetic controls.



**Figure 1.5:** In-space placebo


provinces: indeed, a high pre-treatment MSPE, corresponding to an imperfect fit, yields a low rank in the MSPE ratio test even when the estimated placebo effects are large. Read in conjunction, these two robustness exercises cast some uncertainty over the ability of the synthetic control method, as specified in [Section 1.4](#), to carefully identify the effect of the 2008 BC carbon tax.



**Figure 1.6:** MSPE ratio test

The uncertainty is reinforced by the results of the leave-one-out test, which iteratively removes one of the donor units with positive weight from the control pool and

re-estimates the synthetic control in its absence. The test is reported in [Figure 1.7](#), where it is immediate to assess how the fit is worsened by the removal of critical control units; only the synthetic unit excluding New Brunswick is almost identical to the main synthetic control. In particular, the removal of Manitoba pushes the average pre-treatment difference between treated and synthetic unit to 0.089 metric tons per capita, while removing Ontario brings it up to 0.155: the procedure is not stable with respect to the exclusion of these two provinces, as hypothesised in [Section 1.4](#). The end-of-sample effects also vary considerably, ranging from a 5.5 to a 18.9% reduction in emissions, while the average effect ranges between 11.3 and 20%. The sensitivity of the procedure to the exclusion of either Manitoba or Ontario is of particular concern, given that no such synthetic unit using the same  $V$  weights is capable of reproducing the treated unit’s outcome path in their absence, thereby pinning the identified effect to a particular model specification.



**Figure 1.7:** Leave-one-out test

### 1.5.2 Implementation fulfilling contextual requirements


As outlined in detail in [Abadie \(2021\)](#), the traditional synthetic control method is an appropriate tool for policy evaluation of a single (aggregate) treated unit if and only if a set of contextual and data requirements is satisfied. In the context of the analysis of the impact of carbon pricing on transportation emissions, the seminal paper by [Andersson \(2019\)](#), analysing the case of Sweden, discusses data and contextual requirements in detail, applying the relevant data modifications and inspecting the robustness of the results to the full set of falsification exercises performed in the previous paragraph. On the other hand, the case of [Pretis \(2022\)](#) is much more akin to the “naive implementation” case described in [Section 1.4](#), omitting a rigorous analysis of the instances in which synthetic controls may fail to produce an adequate

counterfactual for the treated unit. Accordingly, we proceed to examine a set of data constraints and conditions which may invalidate the precision of the synthetic control estimator in the case of British Columbia’s 2008 carbon tax and in the context of transportation emissions.

First and foremost, any implementation of the traditional synthetic control method should start with a careful inspection of the characteristics of the proposed donor pool. While it is certainly true that, by restricting the analysis to the remaining Canadian provinces, the risk of insurgence of interpolation bias is minimised (as the analysis is confined to a set of observational units pertaining to the same federal nation), such instances are not necessarily eliminated if large discrepancies between the treated and synthetic units persist. Another concern is that of eliminating from the analysis donors which have suffered from idiosyncratic shocks to their outcomes during the period under investigation. Even though Québec’s 2007 duty has received far less attention than British Columbia’s carbon tax and has since been absorbed in its cap-and-trade scheme (see [Section 1.2](#)), both its implementation and the subsequent implementation of the ETS are sufficient grounds for exclusion of the province from BC’s control pool<sup>13</sup>. Before the eventual roll-out of the federal carbon tax in 2019, another province had introduced a carbon pricing scheme: Alberta, whose carbon tax started at CAD 20/tCO<sub>2</sub> on January 1st, 2017. The model in [Section 1.4](#) uses data up to 2019, hence estimates for the last three years would overlap with the Albertan carbon tax: the estimation period thus needs to be curtailed to 2016, as rightfully done in [Pretis \(2022\)](#). Moreover, two other provinces need to be removed from the donor pool due to concerns about their outcome paths: Yukon, whose road emissions per capita are much higher than those of other provinces and suffer from a spike from 2004 onwards, and Saskatchewan, due to a large positive shock from 2005 onwards, possibly due to a large increase in new heavy vehicle registrations ([Dolter, 2016](#)).

Removal of these donor provinces is sufficient to fundamentally worsen the fit between the synthetic and treated control unit when using the same set of predictors  $V$ , as shown in [Figure 1.8](#); the average absolute difference between the synthetic and treated unit prior to treatment jumps to 0.083 metric tons per capita, a 37.7% increase from the naive model. Even though the excluded provinces had received a near-zero  $W$ -weight, their presence in the control pool was probably essential in order to attain the fit shown in [Figure 1.2](#): indeed, BC is at the left tail of the distribution for what concerns road transportation emissions per capita during the whole duration of the study sample, hence its outcome path is difficult to reproduce when relying only on a limited number of control Canadian provinces. Moreover, the weighting of the donor provinces changes substantially upon the exclusion of Québec, Yukon and Saskatchewan; synthetic BC is indeed a combination of Ontario ( $W = 0.628$ ), Nova Scotia (0.300) and Manitoba (0.072). The effect of the carbon tax is also largely diminished: it is negligible at the end of the sample (2016) and amounts to 0.22 metric tons per capita, or a 6.8% reduction, in an average post-treatment year.

<sup>13</sup>Even if Québec receives negligible weighting in the “naive model” of [Section 1.4](#), its presence in the donor pool could still be the cause of interpolation bias.



**Figure 1.8:** Path plot of road transportation CO<sub>2</sub> emissions per capita in British Columbia and synthetic British Columbia, fulfilling contextual requirements

### 1.5.3 Extensions of the synthetic control method


In [Figure 1.9](#), we report the results of traditional TWFE DID, SCM and SDID estimations of the average treatment effect of BC’s 2008 carbon tax on road transportation emissions using the pool of Canadian provinces, excluding Québec, Yukon and Saskatchewan as mentioned in [Section 1.4](#), as controls<sup>14</sup>. It is immediate to notice how the DID estimator may give rise to biased estimates: the outcome trends of BC and control provinces prior to 2008 are only approximately parallel prior to 1998, diverging significantly thereafter; in contrast, both the SCM and the SDID estimator appear to achieve significantly better balance in pre-treatment trends. In this particular case, with the control pool limited to 7 Canadian provinces, it is not straightforward to ascertain the eventual dominance of SDID on SCM; indeed, the SDID procedure fails to impose perfectly parallel trends during the whole pre-treatment period, while trends for SCM treated and control units appear instead closely matched, as in [Section 1.4](#). However, as is evident from the time weights graph reported together with the outcome plots, SDID assigns a significant weight to the very last periods preceding the carbon tax implementation in 2008, during which the trends appear indeed to be parallel. This particular structure of time weights is nonetheless potentially problematic: analyses putting 100% of the weight on the period preceding treatment risk incurring into biases due to anticipatory behaviour, as reported e.g. by [Heckman and Smith \(1999\)](#). As reported in [Table 1.3](#), all three methods identify a negative impact of the carbon tax on transportation CO<sub>2</sub> emissions; however, both the 90% and the 95% confidence intervals always contain zero, thereby failing to identify a statistically significant impact of the carbon tax on transport emissions.

As a further check on the stability and significance of our estimates, we add a suite of regressions performed using related methods, such as the de-meaned synthetic control (DIFP) proposed by [Doudchenko and Imbens \(2016\)](#) and [Ferman and Pinto \(2021\)](#), the matrix completion (MC) method introduced by [Athey \*et al.\* \(2021\)](#), and penalised version of the traditional SCM and of the DIFP method, where a ridge regularisation is added to the estimation of the synthetic control weights. Importantly, coefficients are consistently negative and similar in magnitude across methods, suggesting that BC has indeed experienced a negative transportation emissions shock after the 2008 carbon tax; however, no point estimate is statistically significant at conventional levels, thereby failing to identify a causal effect of the tax in driving down transportation emissions stably in the long run. [Figure 1.10](#) reports the unit weights assigned to each Canadian control province by DID, SCM and SDID: it is straightforward to determine how SDID weights are more balanced on average than DID weights and employ the full set of control provinces with respect to the 3 control units receiving weights in the SCM estimation.

<sup>14</sup>In [Section 1.A](#), we report all results with an extended control sample, including Québec and Saskatchewan. The results are qualitatively unchanged.

**Table 1.3:** Summary of  $\hat{\tau}$  point estimates and relative standard errors from seven different estimation methods, control pool restricted to Canadian provinces.

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>	<i>DIFP</i>	<i>MC</i>	<i>SCM<sub>ridge</sub></i>	<i>DIFP<sub>ridge</sub></i>
$\hat{\tau}$	-0.59	-0.32	-0.21	-0.34	-0.34	-0.42	-0.30
S.E.	0.63	1.01	0.33	0.51	0.55	1.00	0.55



**Figure 1.9:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on road transportation CO<sub>2</sub> emissions in British Columbia.




Figure 1.10: DID, SCM and SDID units weights.

## 1.6 Potential confounder: carbon leakage from cross-border travel

In this Section, we examine a potential violation of the Stable Unit Treatment Value Assumption (SUTVA) in the context of synthetic controls and synthetic differences-in-differences, consisting in carbon leakage from cross-border shopping for gasoline and diesel fuel. The determinants of cross-border travel have been extensively studied in the Canada-USA context by [Chandra \*et al.\* \(2014\)](#), who find that fluctuations in the CAD-USD exchange rate are the dominant factor in driving consumers to cross the border for shopping. Given that the distribution of US and Canadian populations is unequally concentrated in close proximity of the Canada-USA border, with Canadians living on average closer to the border than Americans, cross-border shopping by Canadian residents is more responsive to swings in the exchange rate than similar behaviour by American citizens. As discussed in [Rivers and Schaufele \(2015\)](#), [Antweiler and Gulati \(2016\)](#), and [Andersson \(2019\)](#) and in [Section 1.3](#), British Columbia shares a land border with the Canadian provinces of Alberta, Yukon and the Northwest Territories, and with the US states of Alaska, Washington, Idaho and Montana. While the majority of these border regions are fundamentally depopulated, there is substantial two-way traffic between BC and Washington State (WA), since a large majority of BC's population lives close to the WA border: there exists then the possibility that, in response to the carbon tax, residents of BC have engaged in cross-border fuel shopping, which would have reduced the registered levels of gasoline sales in BC and therefore biased its road transportation emissions totals downwards.

[Abadie \(2021\)](#) warns about the possibility of violations of the SUTVA stemming



from interference, i.e. spillover effects among treatment and donor units; however, the possibility of spillovers outside the control pool is not discussed extensively. Given the set of results reported in [Section 1.5](#), an eventual downward bias due to cross-border shopping would not change the essence of our insights: unlike [Pretis \(2022\)](#), we do not claim that transportation emissions reductions registered in BC after 2008 are significantly different from zero, even if there is suggestive evidence for such an interpretation due to a reduction in per capita gasoline consumption ([Rivers and Schaufele, 2015](#))<sup>15</sup>. In order to analyse the determinants of cross-border travel in the specific case of BC, we follow [Chandra \*et al.\* \(2014\)](#) in setting up a reduced-form estimation of the joint effect of exchange rate swings and gasoline price fluctuations on the monthly number of border crossings. In traditional models of cross-border shopping, the domestic price of gasoline enters the estimation as a proxy for travel costs, and the sign of its coefficient is expected to be negative, as higher travel costs reduce consumers' propensity to cross the border; however, if consumers cross the border to buy a bundle of goods which includes foreign gasoline, domestic prices play a more ambiguous role, since an increase in the domestic gasoline price vis-à-vis the foreign one could also induce the marginal consumer to become more inclined to travel across the border.

It is straightforward to show this in a simple extension of the [Hotelling \(1929\)](#) model to continuous demand, as proposed by [Friberg \*et al.\* \(2022\)](#). Let  $x_{i1}$  and  $x_{i2}$  be consumer  $i$ 's demand for product 1 (motor fuel) and product 2 (a bundle of all other goods). Let individual  $i$ 's utility function be a Cobb-Douglas of the form  $U_i = x_{i1}^\alpha x_{i2}^{1-\alpha}$ , and her income  $m$ . Product 1 can either be purchased in British Columbia ( $B$ ) or in the US state of Washington ( $W$ ), across the border; product 2 is always purchased domestically<sup>16</sup>. Assume that British Columbian consumers are located at distance  $d_i$  from Washington state, and incur in travel costs  $t = \kappa p_{1B}$  for each kilometre of distance, where  $p_{1B}$  is the price of fuel in BC and  $\kappa$  is fuel consumption per kilometre. Moreover, let  $F_i$  be the idiosyncratic fixed cost incurred by BC consumers in travelling to Washington state, which may reflect opportunity costs, or different preferences regarding waiting times at the BC-WA border stations.  $p_{1W}$  denotes the cost of fuel in Washington state, while  $p_2$  denotes the cost of product 2. The utility maximisation problem of consumer  $i$  can be expressed as:

$$U = \begin{cases} \max_{x_1, x_2} & x_1^\alpha x_2^\beta & \text{s.t.} & p_{1B}x_1 + p_2x_2 + d\kappa(p_{1B}) + F \leq m & \text{when buying fuel in B.C.} \\ \max_{x_1, x_2} & x_1^\alpha x_2^\beta & \text{s.t.} & p_{1W}x_1 + p_2x_2 \leq m & \text{when buying fuel in WA} \end{cases} \quad (1.4)$$

Solving the utility maximisation problem yields two sets of demand functions, one for the case in which the consumer shops for fuel locally and one in which she crosses the border to buy fuel:

<sup>15</sup>Who, however, base their estimation on a short post-intervention period, spanning only 2008-2011.

<sup>16</sup>As shown in [Chandra \*et al.\* \(2014\)](#), this is unlikely to be the case, as cross-border travel between Canada and the US involves shopping for multiple consumption goods and is thus heavily influenced by the CAD-USD exchange rate

$$(x_{1B}^*, x_2^*) = \begin{cases} x_{1B} = \frac{\alpha m}{p_{1B}} \\ x_2 = \frac{(1-\alpha)m}{p_2} \end{cases}$$

$$(x_{1W}^*, x_2^*) = \begin{cases} x_{1W} = \frac{\alpha(m-d\kappa p_{1B}-F)}{p_{1W}} \\ x_2 = \frac{(1-\alpha)(m-d\kappa p_{1B}-F)}{p_2} \end{cases}$$

The role of domestic fuel prices  $p_{1B}$  is thus ambiguous, as the consumer's decision to shop for foreign or local motor fuel depends on swings in the relative prices of domestic and foreign fuel, on travel costs  $d\kappa p_{1B}$  (a combination of fuel prices, fuel efficiency and distance from the border) and on idiosyncratic fixed costs  $F$ . In this context, the BC carbon tax acts as a further wedge in the per litre price of domestic fuel, thereby appearing in both the domestic and in the foreign demand functions. Let  $\tau$  be the per litre level of the carbon tax; demand functions stemming from the same utility maximisation problem can now be expressed as:

$$(x_{1B}^*, x_2^*) = \begin{cases} x_{1B} = \frac{\alpha m}{p_{1B}(1+\tau)} \\ x_2 = \frac{(1-\alpha)m}{p_2} \end{cases}$$

$$(x_{1W}^*, x_2^*) = \begin{cases} x_{1W} = \frac{\alpha(m-d\kappa p_{1B}(1+\tau)-F)}{p_{1W}} \\ x_2 = \frac{(1-\alpha)(m-d\kappa p_{1B}(1+\tau)-F)}{p_2} \end{cases}$$

We parameterise these results in a similar fashion to that of [Friberg \*et al.\* \(2022\)](#), that is,  $m = 1000$ ,  $\alpha = 0.05$ ,  $t = d * \kappa = 0.1$ ,  $F \sim \mathcal{N}(9, 5)$ . We slightly modify the price ratio to reflect that observed between BC and Washington between both 2000 and 2007 and 2008 and 2016, where it's 0.77. We keep  $p_{1B} = 30$ , and therefore,  $p_{1W} = 23.33$ . We also assume that distance increases in increments of  $1/6$  and that 900 consumers inhabit each location on the unit line<sup>17</sup>. We consider two exercises in comparative statics: (1) In a regime of low fixed costs  $F \sim \mathcal{N}(9, 5)$  (where the mean of the individual-level fixed cost is much lower than the price of fuel), we let the price of fuel in BC suffer a positive shock due to a carbon tax. The shock is either  $\tau_{low} = \text{USD } 1.53/\text{L}$  (proportional to the maximum level of the ratio between the BC carbon tax and BC's gasoline price, registered in 2012-2013), or  $\tau_{high} = \text{USD } 6/\text{L}$  (equal to 20% of the prior gasoline price and similar to the Swedish carbon tax studied in [Andersson \(2019\)](#)); (2) we repeat exercise (1) in a regime of higher and less dispersed fixed costs,  $F \sim \mathcal{N}(18, 5)$ . Indeed, imposing a mean of the fixed cost distribution at 9 USD would imply individual fixed costs lower than both the domestic and foreign fuel prices; in practice, as e.g. shown by [Chandra \*et al.\* \(2014\)](#), travellers face a mean 26-minute wait to cross the Canada-US border in each direction, which implies an op-

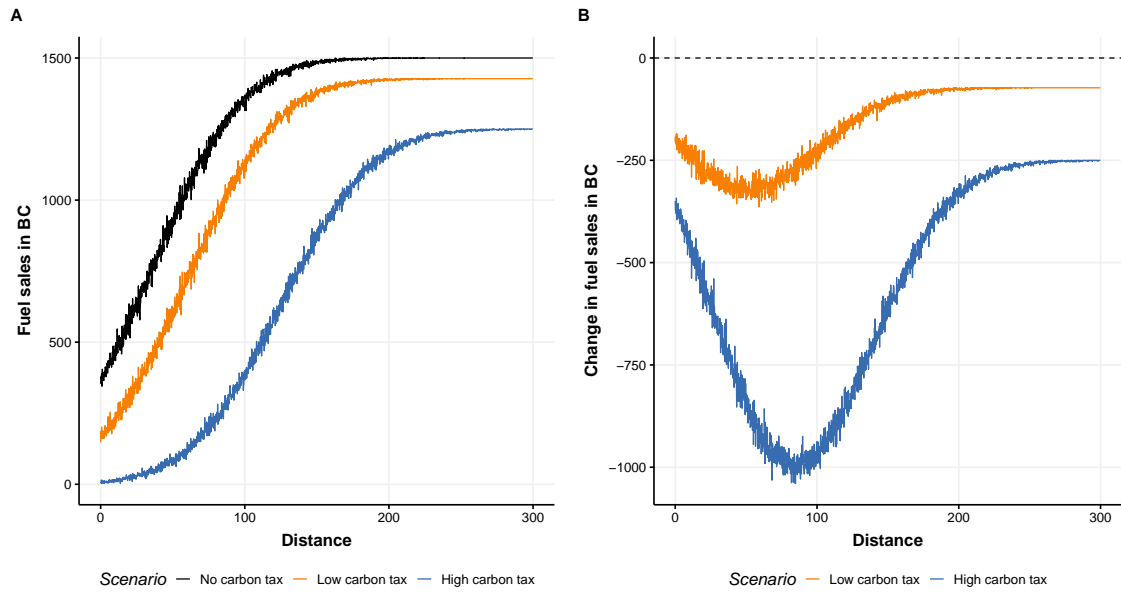
<sup>17</sup>This latter assumption is important in order to determine the shape of the demand response to price and tax increases, a fact which is unaddressed in [Friberg \*et al.\* \(2022\)](#). A uniform population distribution is unlikely in practice: for this reason, in [Figure 1.B.1](#) and [Figure 1.B.2](#) in the Appendix, we replicate the same exercise but positing that consumers are distributed over the  $d \in [0, 300]$  distance continuum according to a truncated Gaussian with  $\mu = 50$  and  $\sigma = 50$ , which approximately reflects the distribution of Canadian residents in British Columbia. Results are qualitatively unchanged, but the hump shape arises also in a high fixed costs regime.

portunity cost of time that is likely to be much higher than the per litre price of fuel<sup>18</sup>.

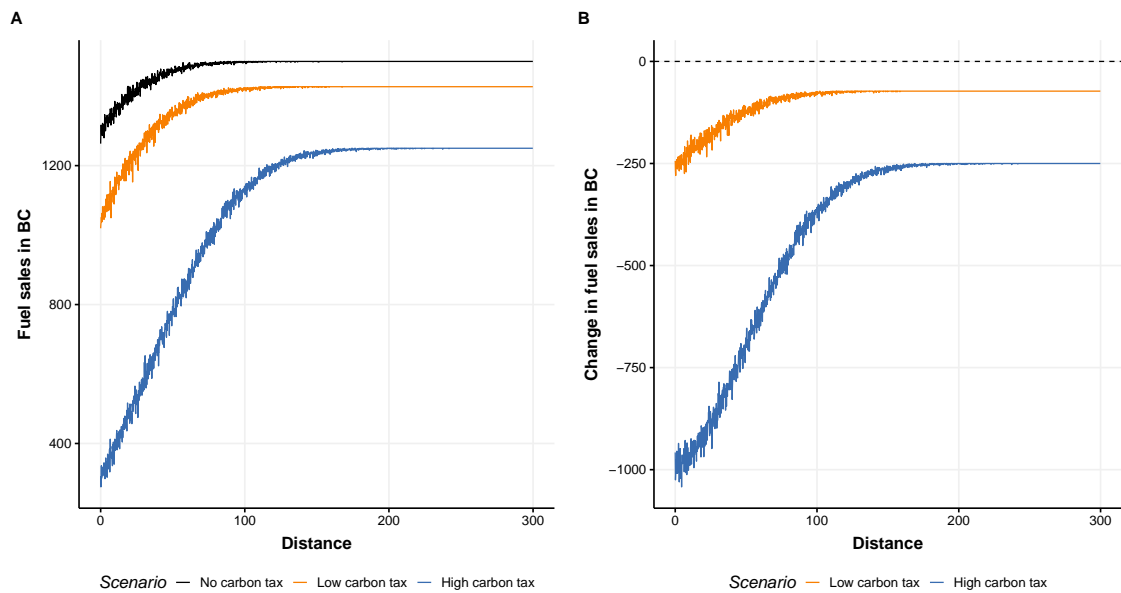
Figure 1.11 and Figure 1.12 illustrate the results of exercise (1) and (2), respectively. First of all, the level of fuel sales after the introduction of any carbon tax is lower at any distance from the international border, as expected. Further, it is immediate to notice how in a regime of relatively low fixed costs, the hump-shaped curved elasticity of fuel sales to shocks in fuel prices appears even at low levels of carbon taxation. This behaviour could give rise to substantial cross-border shopping, as the most elastic locations are relatively inland with respect to the BC-Washington border. Equally important is the role of the distribution of individual level fixed costs in giving rise to hump-shaped elasticities. Indeed, in Figure 1.12, the bell shaped part of the curve vanishes, and the change in domestic fuel sales consequent to an impulse to the carbon price takes a more conventional upward-sloping form, with higher elasticities found at lower distances from the international border. Moreover (panel A), the *level* of domestic sales is much less dependent on distance than with low fixed costs, as in Figure 1.11, and only consumers located in close proximity to the border exhibit cross-border shopping behaviour, while for inland locations fuel demand is practically inelastic with respect to distance. While in principle swings in the fuel price difference between BC and Washington could give rise to extensive cross-border carbon leakage due to hump-shaped elasticity responses, in practice the average real price gap has been equal to 0.21 USD between 2000 and 2007 and to 0.27 USD between 2008 and 2016; up to a certain distance, cross-border shopping for fuel has been rational for consumers before the introduction of the BC carbon tax, and its low level is unlikely to have affected consumers located further inland in presence of high fixed costs of travelling to the border for fuel shopping.

---

<sup>18</sup>The estimated compensation per hour of travel found by Chandra *et al.* (2014) is 29.69 USD.



**Figure 1.11:** Model parameterisation with low individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price.



**Figure 1.12:** Model parameterisation with high individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price.

### 1.6.1 Empirics

#### Determinants of border crossings

In order to understand the relationship between fuel prices, exchange rates and border crossings, and to uncover the determinants of cross-border travel between BC and Washington, we set up a simple model in logarithms following [Chandra \*et al.\* \(2014\)](#) and [Friberg \*et al.\* \(2022\)](#). In particular, we estimate the following regression:

$$\begin{aligned} \ln(\text{Crossings}_t) = & \alpha + \beta_1 \ln(\phi_t) + \beta_2 \ln(e_t) + \beta_3 \ln(e_t) \times [e > 1.45] + \\ & + \beta_4 \ln(e_t) \times [e < 1.11] + \beta_5 \ln(X_t) + \lambda_t + \delta_s + \epsilon_t \end{aligned} \quad (1.5)$$

Where  $\text{Crossings}_t$  is the monthly number of vehicles returning to Canada from the US, analysed separately as the total number of vehicles returning from any trip (Total), the total number of automobiles returning from any trip (Cars), the number of automobiles returning from same day trips (Daytrips) and the number of automobiles returning from trips spanning multiple days (Overnight).  $\phi_t$  is the monthly price of gasoline in BC;  $e_t$  is the monthly real CAD-USD exchange rate. We add interactions of  $e_t$  with its highest and lowest quartiles as in [Chandra \*et al.\* \(2014\)](#), and control for monthly after tax income and unemployment rate (contained in the vector of controls  $X_t$ ).  $\lambda_t$  are month fixed effects and  $\delta_s$  are year fixed effects, while  $\epsilon_t$  is the usual idiosyncratic error term. Since border crossings are serially autocorrelated, we compute standard errors using the [Newey and West \(1987\)](#) procedure, with robustness up to 60 lags. The potential for cross-border trips to influence the price of gasoline in BC and viceversa, may give rise to endogeneity concerns; we therefore estimate the same model with two stage least squares (TSLS), instrumenting the price of gasoline in BC with the crude oil price.

[Table 1.4](#) reports the results from this exercise. While high gasoline prices seem to have a significant influence on cross-border travel in the OLS specification, thereby suggesting consumers behavioural choices compatible with carbon leakage, significance vanishes once the gasoline price is instrumented with the crude oil price. On the other hand, the magnitude and significance of the  $\beta_2$  coefficient measuring the role of the real exchange rate are always consistent and fundamentally unaltered. A weak Canadian dollar is associated with less frequent trips across the border; moreover, this effect is stronger for same day trips, suggesting that purchasing power plays a fundamental role in stimulating crossing behaviour for shopping motives by BC residents.

#### Carbon tax impact on retail sales in Washington state

After having individuated the real exchange rate as the main factor pushing BC residents to cross the US border, the effect of the 2008 BC carbon tax on eventual carbon leakage remains unidentified. In this section, we aim to circumvent the lack of data on BC fuel sales at the station level by exploiting tax revenue data from Washington state's Department of Revenue. The theoretical predictions from the model described in [Section 1.6](#) can indeed be interpreted diametrically from the other

**Table 1.4:** Determinants of border crossings: log-log

	Total (OLS)	Cars (OLS)	Daytrips (OLS)	Overnight (OLS)	Total (IV)	Cars (IV)	Daytrips (IV)	Overnight (IV)
ln Gas price, $\phi$	0.3316*** (0.0561)	0.3503*** (0.0664)	0.3985*** (0.0703)	0.0564 (0.0904)	0.2172 (0.1891)	0.1966 (0.2016)	0.2491 (0.2244)	-0.0402 (0.1537)
ln $e$	-0.8153*** (0.1506)	-0.9296*** (0.1658)	-0.9251*** (0.1781)	-0.8756*** (0.1164)	-0.8636*** (0.2034)	-0.9944*** (0.2196)	-0.9881*** (0.2326)	-0.9163*** (0.1408)
ln $e \times [e > 1.45]$	0.0078 (0.0396)	0.0105 (0.0446)	0.0012 (0.0460)	0.1588*** (0.0601)	0.0148 (0.0421)	0.0199 (0.0462)	0.0104 (0.0491)	0.1647*** (0.0574)
ln $e \times [e < 1.11]$	0.2430 (0.2508)	0.2260 (0.2911)	0.2626 (0.3552)	-0.6484*** (0.2174)	0.2727 (0.2841)	0.2659 (0.3293)	0.3014 (0.3932)	-0.6233*** (0.2363)
ln Income	1.459 (1.641)	0.9444 (1.864)	1.089 (1.945)	4.430*** (1.654)	0.6136 (1.861)	-0.1909 (2.057)	-0.0145 (2.147)	3.716** (1.449)
ln Unemployment	-0.1297 (0.1289)	-0.1194 (0.1500)	-0.1476 (0.1588)	-0.0159 (0.0794)	-0.1344 (0.1302)	-0.1258 (0.1515)	-0.1538 (0.1590)	-0.0199 (0.0817)
Month FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	312	312	312	312	312	312	312	312
R <sup>2</sup>	0.97596	0.97573	0.97564	0.97235	0.97573	0.97541	0.97537	0.97222
Within R <sup>2</sup>	0.30683	0.29827	0.28787	0.18661	0.30028	0.28885	0.28020	0.18291
F-test					149.50	149.50	149.50	149.50
p-value					0.000	0.000	0.000	0.000

*Notes:* Newey-West (1987) SEs robust to serial autocorrelation with 60 lags in parentheses. The Stock-Yogo (2005) test for weak IV is reported together with its p-value. \*\*\*:  $p < 0.01$ , \*\*:  $p < 0.05$ , \*:  $p < 0.1$ .

side of the border: rational behaviour by cross-border shoppers entails that, during periods of high cross-border travel and high fuel price differentials, locations immediately beyond the BC-WA border should see their fuel sales increase by an amount proportional to the domestic decrease in sales. It is plausible that fuel retailers in the state of Washington would respond to the increase in foreign fuel demand by adjusting their margins upwards by a measure low enough to keep attracting foreign consumers and high enough to profit with respect to regular pricing. It is therefore to be expected that the gradient of Washington fuel prices decreases with respect to the inland distance to the Canadian border. On the extensive margin, Washington fuel retailers could respond by increasing the number of service stations or fuelling positions in order to accommodate the relative increase in demand. Nonetheless, these responses are likely to be relatively localised with respect of the size of the entire state of Washington. If the motive pushing BC consumers to cross the border is purely related to shopping, it is reasonable to expect any increase in retail sales to be concentrated in Washington cities located near the BC border, as any increase in inland distance travelled increases the travel costs and the opportunity cost of time, therefore lowering demand and thus utility from cross-border shopping.

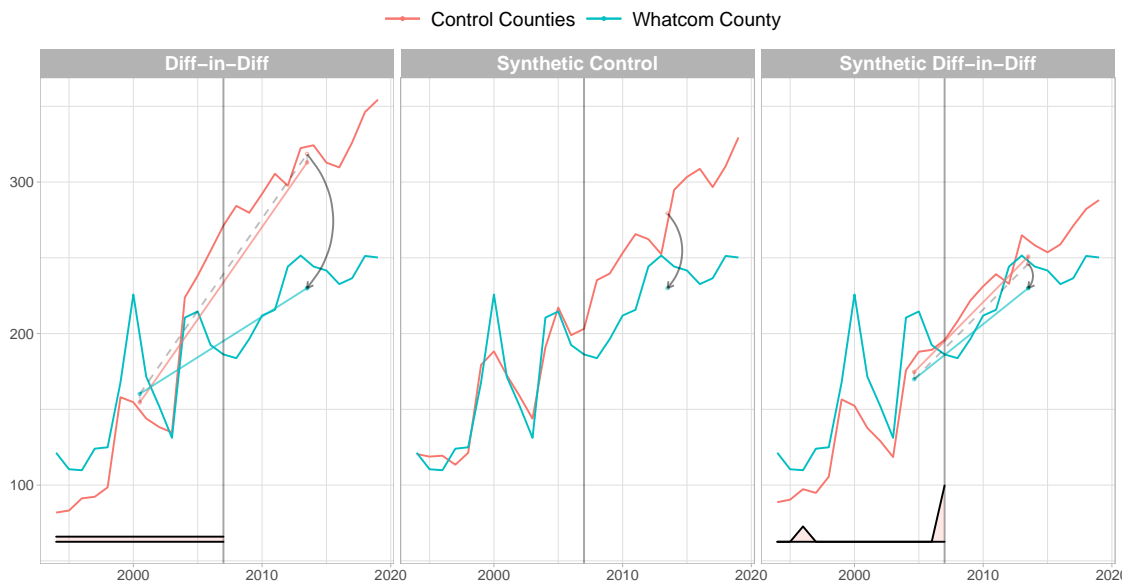
These stylised facts motivate the analysis of cross-border retail trade that we report in what follows. In order to identify the eventual causal effect of the 2008 BC carbon tax on cross-border retail sales, we consider Whatcom county, which hosts the two most trafficked border crossing stations between BC and Washington state, the Blaine-Peace Arch and Blaine-Pacific Highway crossings, as the county “treated” by

the tax-mediated border shopping shock, and estimate the effect of the carbon tax on its retail sales revenues using the other 38 Washington state counties as the control pool. It is plausible that eventual cross-border shopping behaviour by Canadians would be concentrated in the proximity of the BC-WA border, especially considering same day cross-border trips as the mechanism through which cross-border shopping manifests. We consider retail sales at the NAICS 447 industrial classification (Gasoline stations and convenience stores with pumps) in order to directly account for carbon leakage; we further examine NAICS 441 (Automobile dealers) grouped with 447, then NAICS 443 (Food and beverages) and all NAICS 44-45 codes (all retail trade sectors) as a robustness check and a further inspection of general responses in cross-border shopping arising from the carbon tax.


## Results

In [Figure 1.13](#), [1.14](#), [1.15](#), and [1.16](#) we report the results from TWFE, SCM and SDID estimation on the effect of the 2008 BC carbon tax on per capita taxable retail sales from gasoline stations (NAICS 447), gasoline stations coupled with auto dealers (441+447), food and beverage retailers (443) and all retail trade sectors (44-45), respectively. No effect is found for the first two categories, which is suggestive of the lack of significance of cross-border shopping behaviour in driving carbon leakage from BC into Whatcom County, WA. In particular, in [Figure 1.13](#), it is immediate to notice how there is indeed a hump in gasoline stations TRS between 2010/11 and 2014/15, the period coinciding with a strong Canadian dollar; however, the registered increase is not picked up by the models as a statistically significant anomaly with respect to gasoline stations TRS in control counties. The entire automotive and fuel industry shows a similar absence of effects due to the carbon tax, as does the retail trade sector when considered in its entirety. On the contrary, a positive effect on TRS is found in the retail food and beverage sector (which includes supermarkets and large food retailers); however, as evidenced in [Figure 1.15](#), the effect arises from 2010 onwards, and it is thus more aptly attributed to the incidence of exchange rate swings and increases in Canadian consumers' purchasing power rather than to the 2008 carbon tax.

Overall, there does not appear to be an impact of the 2008 BC carbon tax in increasing taxable retail sales in Whatcom county. While there is suggestive evidence that the food and beverage retail sectors have indeed experienced an uptick in sales in the post-intervention period, the increase seems related to exchange rate dynamics rather than to a cross-border travelling incentive induced by higher fuel prices. Consistently, no effects are found for the NAICS category pertaining to gasoline stations, hence reducing confidence in the causal claim that the tax has resulted in significant levels of carbon leakage.




**Figure 1.13:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on gasoline stations (NAICS 447) taxable retail sales in Whatcom County, WA.

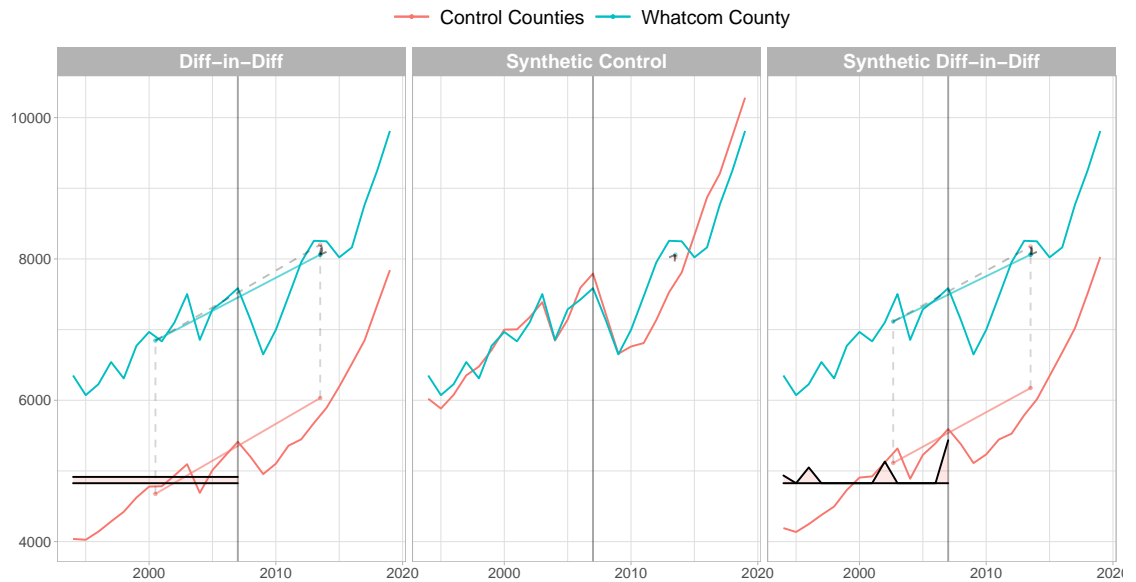


**Figure 1.14:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on auto dealers and gasoline stations (NAICS 441-447) taxable retail sales in Whatcom County, WA.





**Figure 1.15:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on the food sector (NAICS 445) taxable retail sales in Whatcom County, WA.



**Figure 1.16:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on all NAICS retail trade codes (44-45) taxable retail sales in Whatcom County, WA.

## 1.7 Discussion

A highly publicised, globally relevant policy implemented at a relatively aggregate administrative level such as the 2008 British Columbian carbon tax is a natural candidate for a quasi-experimental evaluation using the synthetic control method and its recent extensions. As noted in [Abadie \(2021\)](#), however, “mechanical applications” that overlook the importance of the context in which the analysis is performed or do not carefully assess the robustness of the methodology vis-à-vis a suite of falsification exercises, are prone to yield misleading estimates, which could imply important consequences when evaluation informs the policymaker. In this paper, we provide evidence of the potential pitfalls of naively implementing the synthetic control method in the context of the 2008 British Columbian carbon tax. A recent paper by [Arcila Vasquez and Baker \(2022\)](#) has analysed, among other outcomes, the effect of the 2008 carbon tax on BC’s gasoline consumption, finding a contradictory long-run null/positive effect of the tax on gasoline demand. However, their synthetic control units’ outcome path is far from identical to the treated unit’s, and they fail to provide any of the traditional instruments for inference using synthetic controls besides a leave-one-out test (which confirms that BC’s outcome path is difficult to reproduce with their data)<sup>19</sup>. Much more thorough is the study performed in [Pretis \(2022\)](#), who analyses the whole range of aggregate and sectoral emissions in BC and finds a negative effect on transportation emissions ranging from 3 to 15%. Nonetheless, the study violates several of the contextual requirements listed in [Abadie \(2021\)](#), by including provinces in the donor pool which are treated or experience large shocks to the outcome variable, grouping emissions from treated and untreated segments of the transportation sector, and perhaps most importantly, claiming causality where a less-than-perfect fit between the treated and synthetic unit is observed. By omitting robustness checks and predictor weights, it is also impossible to verify whether the SCM analysis stands up to scientific scrutiny: in [Section 1.5](#), we have illustrated how it is sufficient to exclude unsuitable donor provinces from the control pool to make a credible synthetic unit unattainable in practice. Furthermore, none of these studies examine an oft-cited “elephant in the room”: the possibility that BC’s transportation emissions figures are biased in the first place due to carbon leakage arising from cross-border fuel shopping, mentioned in [Rivers and Schaufele \(2015\)](#), [Antweiler and Gulati \(2016\)](#), and [Andersson \(2019\)](#). Even though the analysis in [Section 1.6](#) has shown that this problem is likely to be overstated, its empirical relevance is undeniable and a careful practitioner needs to take it into account when providing recommendations to prospective policymakers. By avoiding to claim causality for SCM estimates, given the instability of the synthetic control unit to standard falsification exercises, and by extending the analysis via recent advances in the synthetic control literature, in this paper we provide a set of estimates which confirm the negative direction of the impact of the carbon tax on road transportation emissions. Results from the

---

<sup>19</sup>A synthetic control and synthetic difference-in-differences estimation for gasoline consumption was also separately performed for this paper and is available on request. While SCM fails to identify a credible control unit for BC, SDID does a slightly better job, even if the method fails to impose perfectly parallel trends for 2006-2008. The results contradict [Arcila Vasquez and Baker \(2022\)](#), pointing again towards a reduction of gasoline consumption in BC which is consistent with the remaining literature.

synthetic difference-in-differences method are at the lower bound of all estimators and identify a 6.1% reductions, or 0.21 in terms of annual metric tons per capita, which is almost identical to the 6% reduction found in [Andersson \(2019\)](#) for Sweden. While the Swedish tax is much higher than BC's, fuel prices and taxes in Canada are much lower than in Europe, and thus a smaller tax represents a higher relative increase in the price of gasoline: a smaller tax can then lead to similar emissions reductions given the context ([Andersson, 2019](#)). However, it must be noted that the placebo-based standard errors around the point estimates for emissions reductions are large for every estimator considered in [Section 1.4](#), and that standard confidence intervals always include zero, thereby failing to identify a statistically significant result. Given the relative unavailability of donor pool provinces and the large heterogeneities in outcome paths among observational units, however, these concerns seem related to a lack of statistical power more than to a lack of significant effects, also considering that all of the related literature analysing the 2008 carbon tax has found reductions in fuel consumption ([Rivers and Schaufele, 2015](#); [Antweiler and Gulati, 2016](#); [Bernard and Kichian, 2019](#); [Lawley and Thivierge, 2018](#)), or in emissions ([Xiang and Lawley, 2019](#); [Pretis, 2022](#); [Ahmadi \*et al.\*, 2022](#)). Nonetheless, these same caveats highlight a general naivety in synthetic control applications in this field which perhaps veil attitudes related to publication bias. Given that no US state has enacted a carbon tax so far, that sectoral emissions data for US at the state level are available from 1980 onwards, and that carbon taxation *will* be required in the US if the 2015 Paris Agreement targets are to be fully or partially met, future research will undoubtedly leverage the methodology employed in this paper in order to analyse the impact of not-yet-implemented US carbon taxes: a great dose of scientific rigour will be needed in order to avoid the proliferation of misleading estimates of their impact based on “perfunctory applications” of comparative case study methodology ([Abadie, 2021](#)).

## 1.8 Conclusion

As highlighted by the 2019 US and EU economists' statement on carbon pricing, schemes such as the 2008 British Columbian carbon tax will need to be widely applied in the coming years in order to abide to the targets established in the 2015 Paris Agreement. The Canadian federal government itself has imposed a federal backstop mechanism which prescribes that each province which had not implemented its own carbon pricing scheme as of 2019 would be subject to a CAD 20/tCO<sub>2</sub> federal carbon tax. The widespread application of carbon taxes in different jurisdictions calls for effective and rigorous evaluation of their performance, in order to inform policymakers about the efficacy of such measures and to correctly calibrate tax rates to the context in which they are rolled out. The British Columbian carbon tax, hailed as the first North American "grand experiment" in climate policy, has gained praise for its features of gradual ramp-up and revenue neutrality, which have enabled it to overcome well-known obstacles due to public opposition and to persist unrepealed for 14 years. In this paper, we first show that there is a good degree of confidence about its effectiveness, with the most conservative estimate identifying a 6.1% reduction in road transportation emissions in an average year, comparable with higher taxes implemented in the 1990s in Sweden. Concerns about carbon leakage to the neighbouring US State of Washington are likely overstated, having taxable retail sales in the gasoline sector and in the whole retail trade economy of Whatcom county not increased as a consequence of tax-induced cross-border travel. Lastly, we show that methodological concerns are of primary importance when analysing climate policies implemented at aggregate scales: our estimates of emissions reductions are indeed not statistically significant, due primarily to a lack of power and to BC being at the left tail of the emissions distribution, but also possibly to a weak behavioural response to the tax signal. While it is encouraging that British Columbia's road transportation emissions have trended downwards with respect to its synthetic counterfactuals, the *level* of emissions reductions is still not nearly sufficient for the decarbonisation of the province's most polluting sector; moreover, these reductions seem to not have translated into "aggregate" emissions reductions, calling for a stronger ramp-up of the tax rate and a coupling with additional policy tools. Future analyses of carbon pricing schemes at the sub-national level ought to adhere to the strictest empirical rigour in order to correctly inform the policymaker about the advantages and shortcoming of carbon taxes, conscious that they represent just one of the multiple tools at our avail in order to attain decarbonisation rather than a "climate silver bullet".

# References

- Abadie, Alberto (2021). “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects”. *Journal of Economic Literature* 59.2, pp. 391–425.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2011). “Synth: An R package for synthetic control methods in comparative case studies”. *Journal of Statistical Software* 42.13, pp. 1–17.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2010). “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program”. *Journal of the American Statistical Association* 105.490, pp. 493–505.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2015). “Comparative Politics and the Synthetic Control Method”. *American Journal of Political Science* 59.2, pp. 495–510.
- Abadie, Alberto and Gardeazabal, Javier (2003). “The Economic Costs of Conflict : A Case Study of the Basque Country”. *American Economic Review* 93.1, pp. 113–132.
- Ahmadi, Younes, Yamazaki, Akio, and Kabore, Philippe (2022). “How Do Carbon Taxes Affect Emissions? Plant-Level Evidence from Manufacturing”. *Environmental Resource Economics* 82.2, pp. 285–325.
- Andersson, Julius J. (2019). “Carbon Taxes and CO2 Emissions: Sweden as a Case Study”. *American Economic Journal: Economic Policy* 11.4, pp. 1–30.
- Antweiler, Werner and Gulati, Sumeet (2016). “Frugal Cars or Frugal Drivers? How Carbon and Fuel Taxes Influence the Choice and Use of Cars”. *SSRN Electronic Journal*.
- Arcila Vasquez, Andrés and Baker, John D. (2022). “Evaluating carbon tax policy: A methodological reassessment of a natural experiment”. *Energy Economics* 111.C.
- Arkhangelsky, Dmitry, Athey, Susan, Hirshberg, David A., Imbens, Guido W., and Wager, Stefan (2021). “Synthetic Difference-in-Differences”. *American Economic Review* 111.12, pp. 4088–4118.
- Athey, Susan, Bayati, Mohsen, Doudchenko, Nikolay, Imbens, Guido, and Khosravi, Khashayar (2021). “Matrix Completion Methods for Causal Panel Data Models”. *Journal of the American Statistical Association* 116.536, pp. 1716–1730.
- Athey, Susan and Imbens, Guido W. (2017). “The State of Applied Econometrics: Causality and Policy Evaluation”. *Journal of Economic Perspectives* 31.2, pp. 3–32.
- Ben-Michael, Eli, Feller, Avi, and Rothstein, Jesse (2021). “The Augmented Synthetic Control Method”. *Journal of the American Statistical Association* 116.536, pp. 1789–1803.


- Bernard, Jean-Thomas and Kichian, Maral (2019). “The long and short run effects of British Columbia’s carbon tax on diesel demand”. *Energy Policy* 131.C, pp. 380–389.
- Carattini, Stefano, Baranzini, Andrea, Thalmann, Philippe, Varone, Frédéric, and Vöhringer, Frank (2017). “Green Taxes in a Post-Paris World: Are Millions of Nays Inevitable?” *Environmental and Resource Economics* 68.1, pp. 97–128.
- Carattini, Stefano, Kallbekken, Steffen, and Orlov, Anton (2019). “How to win public support for a global carbon tax”. *Nature* 565.
- Chandra, Ambarish, Head, Keith, and Tappata, Mariano (2014). “The Economics of Cross-border Travel”. *The Review of Economics and Statistics* 96.4, pp. 648–661.
- Dolter, Brett (2016). *A Response to Saskatchewan CLIMATE CHANGE White Paper*. Tech. rep. Canadian Centre for Policy Alternatives.
- Doudchenko, Nikolay and Imbens, Guido W. (2016). *Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis*. Working Paper 22791. National Bureau of Economic Research.
- Ferman, Bruno and Pinto, Cristine (2021). “Synthetic controls with imperfect pretreatment fit”. *Quantitative Economics* 12.4, pp. 1197–1221.
- Friberg, Richard, Steen, Frode, and Ulsaker, Simen A. (2022). “Hump-Shaped Cross-Price Effects and the Extensive Margin in Cross-Border Shopping”. *American Economic Journal: Microeconomics* 14.2, pp. 408–38.
- Heckman, James J. and Smith, Jeffrey A. (1999). “The Pre-Programme Earnings Dip and the Determinants of Participation in a Social Programme. Implications for Simple Programme Evaluation Strategies”. *The Economic Journal* 109.457, pp. 313–348.
- Hotelling, Harold (1929). “Stability in Competition”. *The Economic Journal* 39.153, pp. 41–57.
- Houle, David (2013). *Obstacles to Carbon Pricing In Canadian Provinces*. Tech. rep. Sustainable Prosperity.
- IPCC (2018). *Global Warming of 1.5°C, Summary for Policymakers*. IPCC, Switzerland.
- Lawley, Chad and Thivierge, Vincent (2018). “Refining the evidence: British Columbia’s carbon tax and household gasoline consumption”. *The Energy Journal* 39.2.
- Metcalf, Gilbert E (2019). *On the Economics of a Carbon Tax for the United States*. Tech. rep. Brooking Papers on Economic Activity, BPEA Conference Drafts, March 7-8, 2019.
- Mildenberger, Matto, Lachapelle, Erick, Harrison, Kathryn, and Stadelmann-Steffen, Isabelle (2022). “Limited impacts of carbon tax rebate programmes on public support for carbon pricing”. *Nature Climate Change* 12.2, pp. 141–147.
- Murray, Brian and Rivers, Nicholas (2015). “British Columbia’s revenue-neutral carbon tax : A review of the latest ‘grand experiment’ in environmental policy”. *Energy Policy* 86, pp. 674–683.
- Newey, Whitney K. and West, Kenneth D. (1987). “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix”. *Econometrica* 55.3, pp. 703–708.

- Pretis, Felix (2022). “Does a Carbon Tax Reduce CO2 Emissions? Evidence from British Columbia”. *Environmental and Resource Economics* 83, pp. 115–144.
- Rivers, Nicholas and Schaufele, Brandon (2015). “Salience of carbon taxes in the gasoline market”. *Journal of Environmental Economics and Management* 74, pp. 23–36.
- Saberian, Soodeh (2017). “The Negative Effect of the BC Carbon Tax on Vancouver Air Quality: A Good Climate for Bad Air?” *Working Paper*.
- Statistics Canada (2021a). *Canada’s official greenhouse gas inventory*. URL: <https://donnees.ec.gc.ca/data/substances/monitor/canada-s-official-greenhouse-gas-inventory/>.
- Statistics Canada (2021b). *Gross domestic product (GDP) at basic prices, by industry, provinces and territories (x 1,000,000)*. Table 36-10-0402-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610040201>.
- Statistics Canada (2021c). *Number of vehicles travelling between Canada and the United States*. Table 24-10-0002-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=2410000201>.
- Statistics Canada (2021d). *Population estimates on July 1st, by age and sex*. Table 17-10-0005-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1710000501>.
- Statistics Canada (2021e). *Sales of fuel used for road motor vehicles, annual (x 1000)*. Table 23-10-0066-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=2310006601>.
- Statistics Canada (2021f). *Upper income limit, income share and average of adjusted market, total and after-tax income by income decile*. Table 11-10-0193-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1110019301>.
- Statistics Canada (2021g). *Vehicle registrations, by type of vehicle*. Table 23-10-0067-01. URL: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=2310006701>.
- World Bank (2018). *State and Trends of Carbon Pricing 2018*. World Bank.
- World Bank (2022). *State and Trends of Carbon Pricing 2022*. World Bank, Washington, DC.
- Xiang, Di and Lawley, Chad (2019). “The impact of British Columbia’s carbon tax on residential natural gas consumption”. *Energy Economics* 80.C, pp. 206–218.

## 1.A Full Sample Results

**Table 1.A.1:** Summary of  $\hat{\tau}$  point estimates and relative standard errors from seven different estimation methods, full control pool.

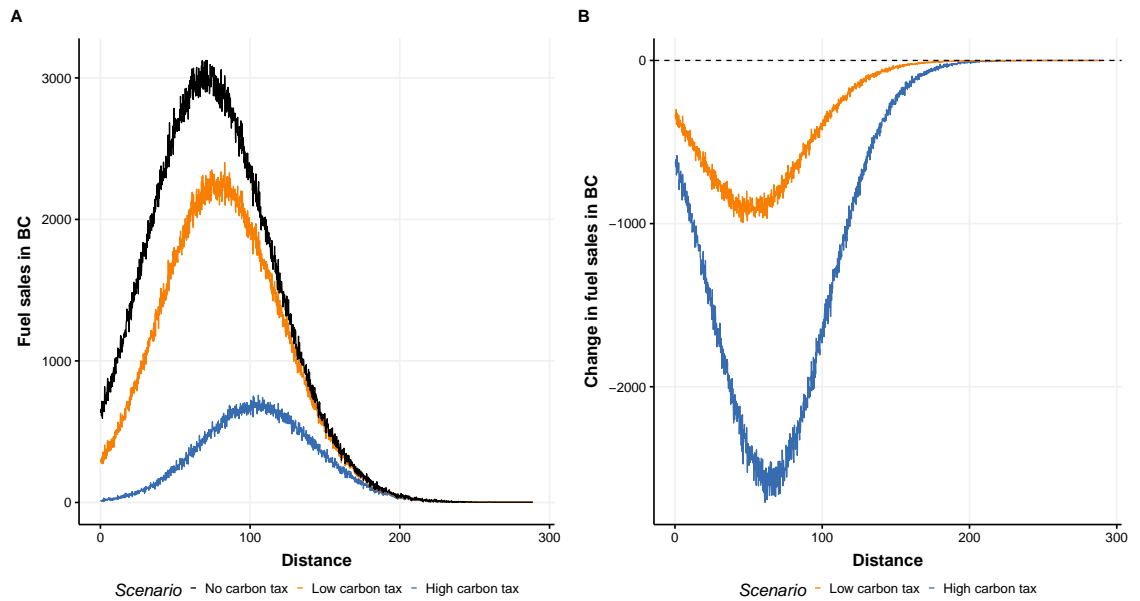
	<i>DID</i>	<i>SCM</i>	<i>SDID</i>	<i>DIFP</i>	<i>MC</i>	<i>SCM<sub>ridge</sub></i>	<i>DIFP<sub>ridge</sub></i>
$\hat{\tau}$	-0.77	-0.33	-0.17	-0.34	-0.47	-0.28	-0.27
S.E.	0.96	0.78	0.48	0.44	0.77	0.88	0.63



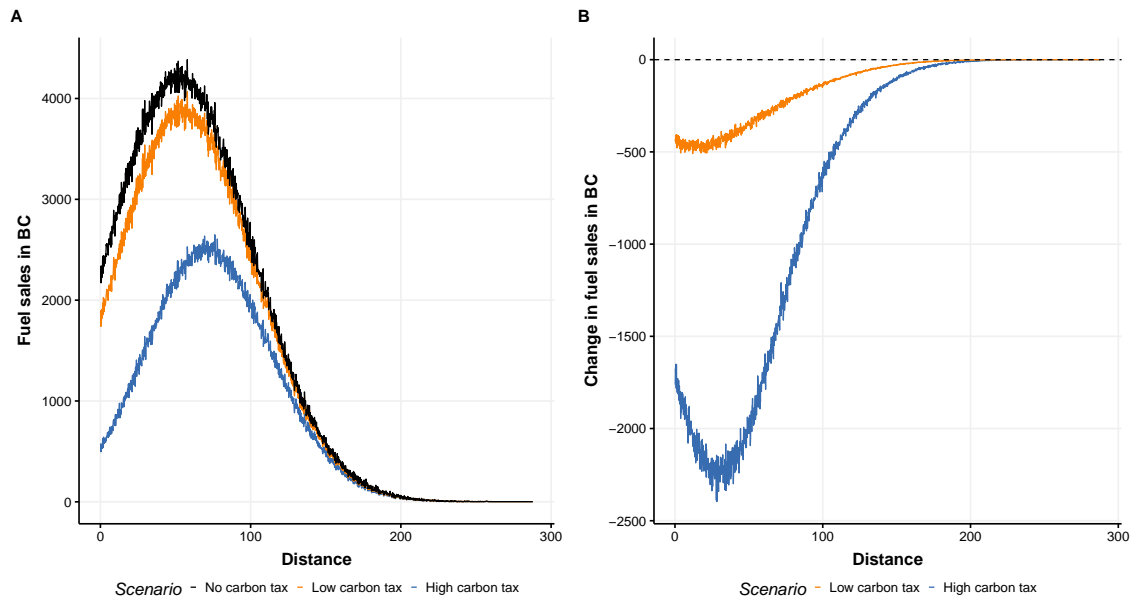
**Figure 1.A.1:** DID, SCM and SDID estimates of the effect of the 2008 carbon tax on road transportation CO<sub>2</sub> emissions in British Columbia, full control pool.



## 1.B Alternative specification of the demand model



**Figure 1.B.1:** Model parameterisation with low individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price.



**Figure 1.B.2:** Model parameterisation with high individual fixed costs. Panel (A): Domestic fuel sales with no carbon tax, low carbon price or high carbon price; Panel (B): Change in domestic fuel sales with low carbon price and high carbon price.

# Chapter 2

## Carbon Pricing with Regressive Co-benefits: Air Quality and Health Effects from British Columbia's Carbon Tax are Positive but Unequally Distributed

### Abstract

Empirical estimates of realised co-benefits arising from carbon pricing, such as improvements in local air pollution, are practically inexistent. In this paper, the 2008 carbon tax implemented in British Columbia, Canada is exploited as a source of exogenous variation in order to evaluate the effect of a carbon pricing scheme on air quality. Combining a granular dataset at the census dissemination area level (corresponding to US census tracts), satellite observations of local pollutants, and a synthetic difference-in-differences methodology, we find that the carbon tax has reduced PM<sub>2.5</sub> emissions by 0.36-0.89  $\mu\text{g}/\text{m}^3$  in British Columbian metropolitan areas compared to the rest of Canada, or 5.2-10.9% compared to pre-tax levels. The result is heterogeneously distributed across census dissemination units: areas which present lower levels of baseline pollution, less dense and in higher income brackets experience greater reductions, suggesting a spatial dimension of the regressive nature of climate policy. A mechanism underlying the results is the substitution of transport mode from driving to public transit, cycling, and walking. Despite the regressive nature of the carbon tax, its per capita co-benefits arising from the conversion of air pollution hazard rates into monetary values are large, at \$198 per capita.

---

For helpful discussions and comments we thank, in alphabetical order: Antonio Avila-Uribe, Eugenie Dugoua, Glen Gostlow, Ben Groom, Beatriz Jambrina-Canseco, Charles Palmer, Alberto Parmigiani, Elena Perra, Julien Picard, Capucine Riom, Sefi Roth, and all participants at research seminars hosted by the Department of Geography and Environment at the LSE in 2022. The author acknowledges funding from the UK Economic and Social Research Council (ESRC). All remaining errors are my own.

## 2.1 Introduction

The major sources of CO<sub>2</sub> emissions are the fossil fuel combustion processes which also give rise to emissions of air pollutants. Climate change and air pollution can then be categorised as complementary global and local externalities from fossil fuel use. Therefore, efforts to control CO<sub>2</sub> emissions by internalising the social cost of carbon are bound to give rise to significant “co-benefits” in terms of air quality improvements.

The preferred economic policy instrument to achieve CO<sub>2</sub> reductions is a carbon tax (Nordhaus, 2008; Weitzman, 2015; Weitzman, 2016); however, the worldwide rate of carbon tax adoption has been historically low, with sparse jurisdictions enacting this form of regulation (World Bank, 2022). Given the relative scarcity of long-tenured carbon pricing schemes, it is perhaps unsurprising that empirical evidence of their causal impact on global pollutants is sporadic at best (Andersson, 2019), while attempts to quantify the local co-benefits from carbon taxation from observational data are practically inexistent. On the contrary, there is a large and growing literature which, using theoretical insights (Parry *et al.*, 2015) and simulation models (Zhang *et al.*, 2021; Knittel and Sandler, 2011), has attempted to calculate the monetary value of air pollution improvements due to carbon taxation and compare them with the cost of mitigation policies.

In particular, net health co-benefits arising from carbon taxation can reach a high enough magnitude to partially or fully offset the mitigation costs for households at a national (Li *et al.*, 2018; Shindell *et al.*, 2016) and global (West *et al.*, 2013; Vandyck *et al.*, 2018) level, and may provide strong additional incentives for a swift transition to a low-carbon economy. Moreover, reductions in morbidity and mortality due to improvements in air quality are not likely to capture the full extent of the local pollution externality: a large body of research has linked air pollution to negative educational outcomes (Ebenstein *et al.*, 2016; Wen and Burke, 2022), increase in crime rates (Bondy *et al.*, 2020), reductions in labour productivity (Graff Zivin and Neidell, 2012), housing prices (Sager and Singer, 2022; Freeman *et al.*, 2019) and other non-health outcomes (Aguilar-Gomez *et al.*, 2022), suggesting that any attempt at quantifying the monetary impact of co-benefits based on health outcomes only would, at best, provide a lower bound of the beneficial consequences of air quality improvements.

It is therefore of fundamental importance to analyse, *ex post*, the eventual causal effect of carbon pricing instruments in reducing local pollution, in order to estimate the magnitude of co-benefits from carbon taxation; incorporating these estimates in cost-benefit analyses of carbon pricing is indeed susceptible of raising the optimal carbon price to reflect not only the social cost of carbon, but also the social cost of complementary local pollutants (Parry *et al.*, 2015). Furthermore, identifying significant co-benefits from carbon pricing could incentivise policymakers to enact mixed instruments for abatement, which appear to be the optimal choice when pollutants are complements (Ambec and Coria, 2013), and to couple carbon pricing with layered, spatially heterogeneous interventions specific to air quality (Zhang

*et al.*, 2021), in order to reap the benefits from policy additionality.

In this paper, we exploit the ideal policy setting of British Columbia, Canada, which introduced a revenue-neutral carbon tax on July 1st, 2008, in order to empirically estimate the impact of carbon taxation on PM<sub>2.5</sub> concentrations in the Province’s Census Metropolitan Areas. The tax, covering approximately 70% of British Columbian CO<sub>2</sub> emissions, was initially introduced at a rate of \$10/tCO<sub>2</sub>, and sequentially ramped up by \$5 per year until 2012, when it was frozen at \$30/tCO<sub>2</sub> until 2018. Importantly, no other Canadian Province introduced relevant<sup>1</sup> carbon pricing schemes between 2008 and 2018, when the tax was rolled out on a federal basis, which allows us to rely on a control pool comprised of Census Metropolitan Areas in other Canadian Provinces. We leverage high-resolution data on PM<sub>2.5</sub>, based on a combination of satellite observations, geo-chemical models and ground-based monitoring stations, from [Meng \*et al.\* \(2019\)](#) and [van Donkelaar \*et al.\* \(2019\)](#), and combine them with highly disaggregated socio-economic data at the Dissemination Area level<sup>2</sup> retrieved from the Canadian Census at 5-year intervals between 2001 and 2016, in order to assess whether the tax has resulted in significant reductions in air pollution concentrations within British Columbian cities. We contribute to the nascent literature on the econometric evaluation of the effect of carbon pricing on atmospheric emissions ([Andersson, 2019](#)), and, to the best of our knowledge, provide the first empirical analysis and quantification of the magnitude of air pollution co-benefits arising from a carbon tax<sup>3</sup>.

The central result of the paper is that the 2008 British Columbian carbon tax has resulted in statistically significant reductions in PM<sub>2.5</sub> concentrations, with a lower bound average estimate of  $-0.36 \mu\text{g}/\text{m}^3$  (using the [Meng \*et al.\* \(2019\)](#) dataset over 2000-2016) and an upper bound average estimate of  $-0.89 \mu\text{g}/\text{m}^3$  (using the [van Donkelaar \*et al.\* \(2019\)](#) dataset over 2000-2018), corresponding to a 5.2-10.9% reduction in particulate matter concentrations with respect to pre-treatment average levels. Importantly, as in e.g. [Andersson \(2019\)](#) and [Sager and Singer \(2022\)](#), this result is obtained by moving away from traditional two-way fixed effects difference-in-differences (TWFE-DID) estimation, in light of a violation of the foundational parallel trends assumption: particulate matter trends between British Columbian and control Dissemination Areas diverge prior to the implementation of the carbon tax, thereby biasing DID estimates. As in [Andersson \(2019\)](#), we rely on a family of estimators related to the synthetic control method (SCM) for comparative case studies ([Abadie and Gardeazabal, 2003](#); [Abadie, 2021](#)), employing in particular the synthetic difference-in-differences (SDID) estimator by [Arkhangelsky \*et al.\* \(2021\)](#) as our preferred methodology. The estimator, which to the best of our knowledge

<sup>1</sup>As discussed in Chapter 1, Québec implemented a minimal fuel duty in 2007, at a rate of 0.8 cents per litre of gasoline. Since this intervention is not regarded as a carbon pricing scheme ([World Bank, 2022](#)), it amounts to a simple change in gasoline excise taxation.

<sup>2</sup>Corresponding roughly to US Census tracts.

<sup>3</sup>A precedent attempt, confined in geographical scope to the city of Vancouver and in the temporal dimension to the year 2013, is the unpublished work by [Saberian \(2017\)](#). Our contribution goes beyond the scope of that paper by extending the series to 2016-2018, exploiting high-resolution rasters in order to avoid incurring in measurement error, inspecting the heterogeneity and mechanisms underlying our results, and quantifying the impacts in terms of health gains.

has not been utilised in any applied study in environmental economics to date, has attractive features which combine the use of unit and time fixed effects from TWFE-DID, unit weights from SCM, and novel time-specific weights in order to construct a set of synthetic counterfactuals whose outcome path is approximately parallel to that of treated units prior to the intervention. In our setting, with multiple treated units and a large number of control units to draw synthetic counterfactuals from, both the SCM and SDID perform well in addressing concerns about diverging pre-treatment trends and appear to identify unbiased and robust estimates of the impact of the carbon tax on  $PM_{2.5}$  levels.

We subsequently inspect whether these reductions arise heterogeneously within British Columbian metropolitan areas. A substantial body of research has indeed documented geographic and socio-economic disparities in pollution levels within cities (e.g. [Jbaily \*et al.\*, 2022](#); [Currie \*et al.\*, 2020](#)) and it is paramount to inspect whether carbon taxation, which has been shown to be regressive over income levels ([Douenne, 2020](#)), presents similar characteristics with respect to other socio-demographic components. We split the pool of treated units in quintiles of pre-existing pollution, population and road density, night-time luminosity, median income levels and dwelling values and estimate the impact of the tax on  $PM_{2.5}$  reductions for each quintile of each baseline characteristics. The carbon tax appears to be doubly regressive, with an additional spatial dimension: reductions are higher in previously less polluted, better off areas, and lower in denser urban conglomerations. Pricing carbon, while giving rise to co-benefits across the entirety of urban areas, may thus exacerbate the pollution-income and the pollution-density gaps (see e.g. [Carozzi and Roth, 2022](#)), highlighting the need for spatially differentiated climate interventions ([Nehiba, 2022](#)) or for additional layered instruments aimed at internalising the congestion externality in urban centres and reducing local pollution (e.g. [Pestel and Wozny, 2021](#); [Sarmiento \*et al.\*, 2022](#); [Gehrsitz, 2017](#)). Other examples of incremental policies to aid carbon pricing in providing co-benefits are incentives for alternative transport modes: by exploiting information on commute modes contained in the Canadian census data, we show that one underlying feature driving the reductions in  $PM_{2.5}$  in British Columbia is the switch from high to low emission transport modes, principally public transport.

Finally, drawing from the insights of [Fowlie \*et al.\* \(2019\)](#) and [Carozzi and Roth \(2022\)](#), we convert our estimates of particle pollution reductions into mortality reductions, based on the environmental health and epidemiology literature ([Lepeule \*et al.\*, 2012](#); [Krewski \*et al.\*, 2009](#)) and associated monetary gains, relying on the concept of the Value of a Statistical Life. The median monetary health gains appear to be large, in the order of \$88-402 per capita depending on the specification; our preferred estimate of \$198 is 66% of the Low Income Climate Action Tax Credit, the carbon tax governmental rebate accruing to low-income families to mitigate the cost of carbon pricing. Health gains stemming from  $PM_{2.5}$  are similarly heterogeneous over space, with greater benefits manifesting in peri-urban areas rather than in city centres, and exhibit an inverse correlation with income within metropolitan areas, corroborating the “spatial regressiveness” hypothesis.

The remainder of the paper begins with a detailed description and graphical analysis of the data sources in [Section 2.2](#). In [Section 2.3](#), we describe the empirical analysis, highlighting the shortcomings of TWFE-DID in the context, outlining the SCM and SDID methodologies, the heterogeneity analysis, and the inspection of underlying mechanisms; we present all main and heterogeneous results in [Section 2.4](#), while the commute mode analysis is reported in [Section 2.5](#). In [Section 2.6](#), we describe how we obtained estimates of mortality reductions and monetary health gains, and visualise the results spatially. [Section 2.7](#) concludes, and additional information is reported in the Appendices.

## 2.2 Data and Descriptive Statistics

In order to analyse the effect of British Columbia’s 2008 carbon tax on air quality, we assemble and process information on local pollutants’ concentrations, geographic characteristics, and socio-economic dynamics from multiple sources. The observational units which we employ in the analysis are Dissemination Areas (DAs), the smallest standard geographic areas for which Canadian census data are disseminated. Since the paper is concerned with analysing the effect of carbon pricing on air quality in cities, we restrict the geographic scope of the dataset to 26 Canadian Census Metropolitan Areas (CMAs), thereby excluding rural areas and smaller towns<sup>4</sup>. Canadian census data is obtained from [von Bergmann \*et al.\* \(2022\)](#), while DA census boundaries are converted to common geographies based on [von Bergmann \(2021\)](#), and using DA administrative boundaries from the 2016 Canadian census as the target geography. Our final dataset is thus comprised of 25,479 DAs observed over 19 years, from 2000 to 2018, across 26 CMAs.

The dependent variable employed in the main part of the paper is yearly average  $PM_{2.5}$  concentration, for which we rely on the [Meng \*et al.\* \(2019\)](#) dataset, which combines information from satellite-retrieved Aerosol Optical Depth with simulations and ground-based observations obtained from monitoring stations readings. We extract the mean value of yearly  $PM_{2.5}$ , weighted by grid-cell level population counts obtained from [Rose \*et al.\* \(2020\)](#), onto the 25,479 DAs which constitute our dataset for every year between 2000 and 2018. The resolution of the  $PM_{2.5}$  raster data is  $0.01^\circ \times 0.01^\circ$ , while population data is available for grid cells of dimension  $0.0083^\circ \times 0.0083^\circ$ , hence implying that the population raster had to be resampled at the resolution of the  $PM_{2.5}$  raster in order to be viable for use in the weighted mean calculation. Hence, for each DA, the dependent variable takes the form:

---

<sup>4</sup>The CMAs in the dataset are: St. John’s, Halifax, Saint John, Quebec, Trois Rivieres, Sherbrooke, Montreal, Ottawa, Saguenay, Kingston, Toronto, Hamilton, St. Catharine’s, Kitchener, London, Windsor, Sudbury, Thunder Bay, Winnipeg, Regina, Saskatoon, Calgary, Edmonton, Abbotsford, Vancouver, and Victoria. While the number of Canadian CMAs is 35 in the latest available census wave (2016), we only keep in the dataset those CMAs which were designated as such in the 2001 Census, in order to ensure compatibility across all waves.

$$PM_{2.5it} = \frac{1}{N_j} \sum_{j=1}^N Pop_{jt} * PM_{2.5jt} \quad (2.1)$$

Where  $j = 1, \dots, N$  is the number of raster grid cells in a DA  $i$ ,  $Pop_{jt}$  is the population count in grid cell  $j$  at time  $t$ , and  $PM_{2.5jt}$  is the value of the particulate matter raster in grid cell  $j$  at time  $t$ .

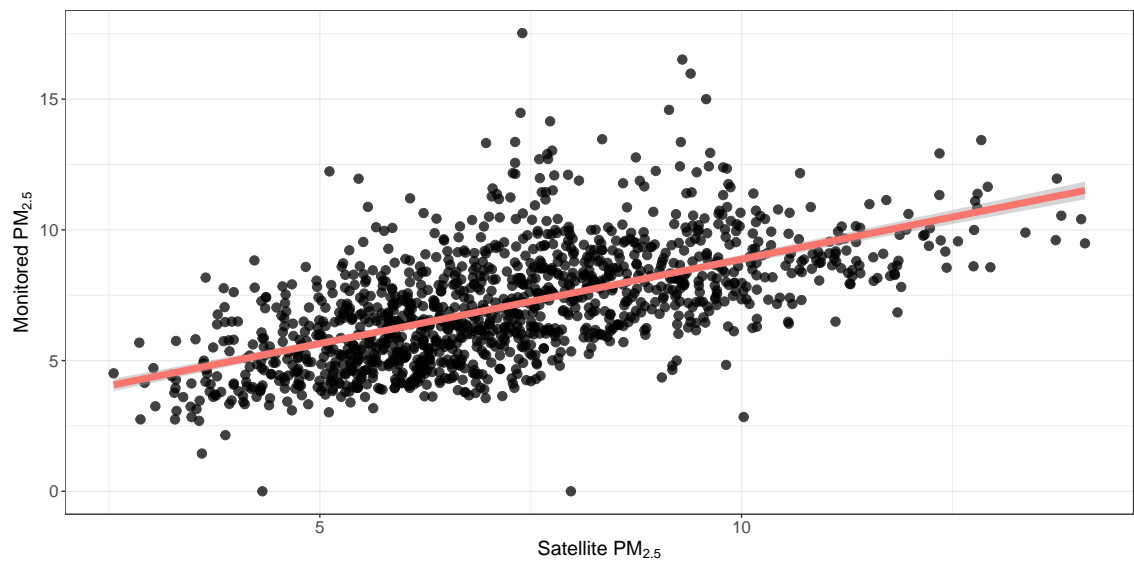
The main advantage of this source compared to data obtained from monitoring stations only (Saberian, 2017), is their much wider spatial and temporal coverage, which also allows us to overcome the selection problem mentioned in Carozzi and Roth (2022) relative to the endogenous location of monitoring stations within urban areas<sup>5</sup>. The entity of data loss when using ground-based data is considerable: PM<sub>2.5</sub> data from the National Atmospheric Surveillance Program (NAPS) is only available for 61 DAs in 2000, growing to 230 in 2018 as new monitoring stations get added every year (see Figure 2.A.1). Nonetheless, the satellite-retrieved measurements from Meng *et al.* (2019), when restricted to the DAs with at least one PM<sub>2.5</sub> ground monitoring station, correlate well with the NAPS readings, as shown in Figure 2.2.1.

We rely on the Meng *et al.* (2019) PM<sub>2.5</sub> estimates in order to produce our main results. However, we also run the main analysis using PM<sub>2.5</sub> concentration data from van Donkelaar *et al.* (2019), as done e.g. in Sager and Singer (2022). While the two estimates are highly related, with a Pearson correlation coefficient of 0.729 (see Figure 2.2.2), the concentrations from Meng *et al.* (2019) are generally lower throughout the sample. Moreover, a closer inspection of the van Donkelaar *et al.* (2019) rasters reveals that, beginning with the year 2004, much of the variability of PM<sub>2.5</sub> pixel values over Canadian CMAs is swept out, resulting in unrealistic estimates of pollution concentrations, especially with respect to their distribution over densely populated DAs. The choice of employing data from Meng *et al.* (2019) is therefore conservative, as results using the van Donkelaar *et al.* (2019) dataset are generally higher in magnitude, as Section 2.4 shows.


Aware of a burgeoning literature relating population density and air pollution (Carozzi and Roth, 2022; Borck and Schrauth, 2021), we obtain population counts at the DA level from Rose *et al.* (2020), which are available for all years between 2000-2018<sup>6</sup>. Road density is also likely to be a highly influential predictor of particulate pollution from traffic congestion: therefore, we obtain data on the road network predating the carbon tax from Statistics Canada (2006 Census Road Network) and calculate baseline road density by dividing the total length of the network in each DA by its surface. We complement these information with data on night-time luminosity, often considered as a valuable proxy for economic activity (e.g. Henderson *et al.*, 2012), by extracting the harmonised version of the nighttime lights dataset from Xuecao *et al.* (2020), from which we calculate the mean level of night lights (NTLs) at the

<sup>5</sup>Monitoring stations are likely to be located where air pollution is higher, thereby introducing measurement error in an eventual empirical analysis.

<sup>6</sup>The dataset also contains population counts for all DAs extrapolated from Canadian censuses; however, this data is only available in 5-years intervals between 2001 and 2016.



**Figure 2.2.1:** Satellite PM<sub>2.5</sub> (Meng *et al.*, 2019) vs PM<sub>2.5</sub> from NAPS monitoring stations. Both measures are in  $\mu\text{g}/\text{m}^3$ . The correlation coefficient is 0.597.

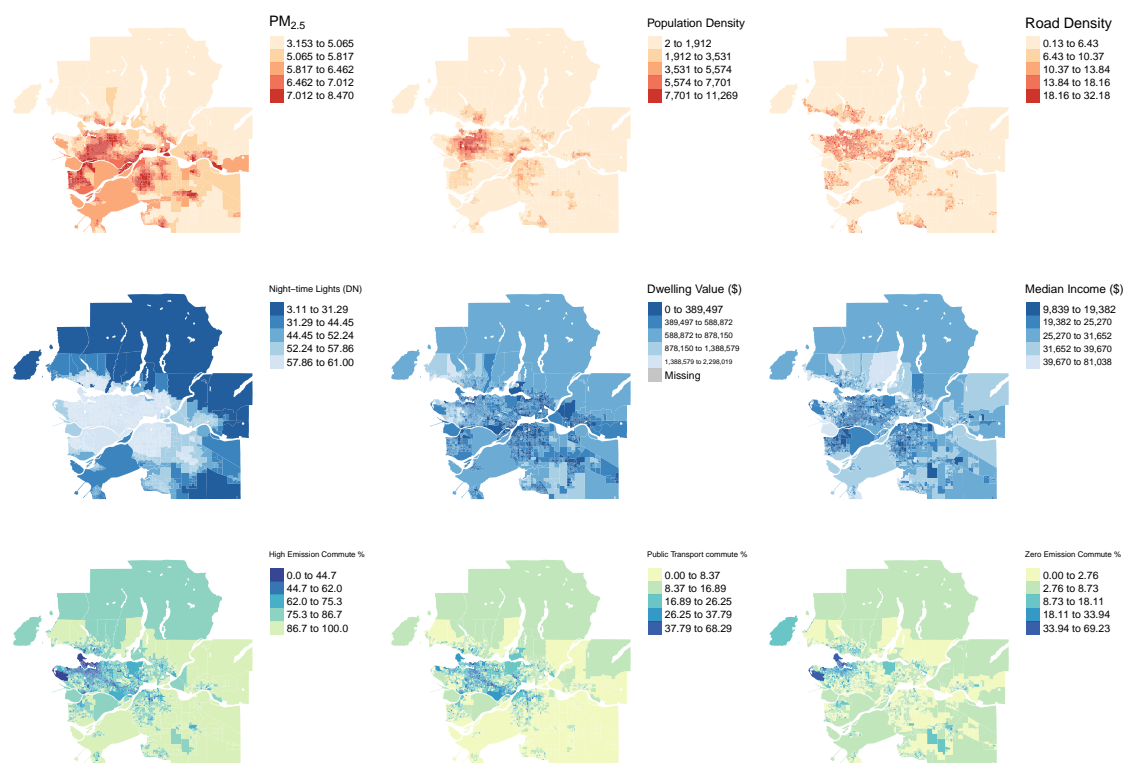


**Figure 2.2.2:** Satellite PM<sub>2.5</sub> (Meng *et al.*, 2019) vs Satellite PM<sub>2.5</sub> (van Donkelaar *et al.*, 2019). Both measures are in  $\mu\text{g}/\text{m}^3$ . The correlation coefficient is 0.729.



DA level for 2000-2018. Moreover, we extract information about weather controls, namely precipitation, maximum and minimum temperature, and wind speed, from [Abatzoglou \*et al.\* \(2018\)](#). We convert monthly to annual averages and compute the mean of each variable at the DA level between 2000 and 2018. If the carbon tax was successful in producing a behavioural adjustment in BC residents, we are likely to observe higher take up of alternative means of transport, especially within metropolitan areas, as an economic coping mechanism. Therefore, we leverage the detailed information contained in the four waves of Canadian census data between 2001-2016 to retrieve DA-level data on commute modes. We divide commute modes in two different categories: (1) High emission (cars, taxis, and motorcycles); (2) Low emission (public transport, bicycles, and walking). Furthermore, we divide up the low emission category in (3) Public transport only and (4) Zero emission (cycling and walking). Finally, we also employ the Canadian censuses to retrieve information on average and median income, and average dwelling values at the DA level, in order to inspect whether the tax has produced heterogeneous impacts along these dimensions.

[Figure 2.2.3](#) plots the baseline spatial distribution of the dependent variable and the main covariates over the Vancouver CMA, the most populated metropolitan area in the treated province of British Columbia. Time-varying variables are averaged over 2005-2007, the three years preceding the implementation of the carbon tax, while road density is time-invariant and all variables retrieved from the Canadian Census are taken at their 2006 values, the last observation before the tax was instituted. The spatial distribution of  $PM_{2.5}$  concentrations is as expected: values are indeed higher in central areas rather than in the periphery, qualitatively adhering to the traditional association with population density found e.g. in the US ([Carozzi and Roth, 2022](#)) or Germany ([Borck and Schrauth, 2021](#)); moreover, road density and night-time lights seem to be highly spatially correlated with air pollution at the baseline, while the pattern is less clear with respect to dwelling values (taken as a proxy of local wealth) and median income. On the contrary, baseline commute mode seems to be inversely related with the spatial distribution of  $PM_{2.5}$ : areas whose inhabitants are less reliant on cars, taxis and motorbikes seem to be more polluted on average, a result probably due to their centrality with respect to the road networks and urban form. Summary statistics for the whole sample, split across treatment and control CMAs, are presented in [Table 2.A.1](#) and [Table 2.A.2](#) for the pre-treatment and post-treatment periods, respectively.



**Figure 2.2.3:** Spatial distribution of PM<sub>2.5</sub> and relevant covariates within the Vancouver CMA. Top row: PM<sub>2.5</sub>, population density and road density; Middle row: NTLs, dwelling value and median income; Bottom row: high emission commute mode %, public transport % and cycling/walking %.

## 2.3 Empirical Strategy

### 2.3.1 Two-way fixed effects difference-in-differences (TWFE-DID)

The core aim of our empirical strategy is to estimate the treatment effect of the 2008 British Columbian carbon tax on local air pollution, measured in terms of  $PM_{2.5}$  concentrations at the DA level. A first attempt at this exercise can be operationalised in terms of a traditional two-way fixed effects difference-in-differences (TWFE-DID) strategy. The estimating equation takes the form:

$$PM_{2.5it} = \tau TAX_{it} + \theta_t + \eta_i + \epsilon_{it} \quad (2.2)$$


Where  $TAX_{it}$  is the DID binary indicator, taking value 1 for all treated units after the implementation of the carbon tax in 2008, and 0 for all other observations;  $\theta_t$  and  $\eta_i$  are respectively time and unit specific fixed effects,  $\epsilon_{it}$  is a time-varying idiosyncratic error term, and  $\beta$  is the coefficient of interest, i.e. the average effect of being exposed to the carbon tax.

In order for  $\beta$  to be equal to the average treatment effect on the treated cohort (ATT), the identifying assumption is that parallel outcome trends between the treated and the control units hold, i.e. if the 2008 carbon tax had not been implemented in British Columbia,  $PM_{2.5}$  levels in British Columbian DAs would have followed the same trajectory as  $PM_{2.5}$  levels in DAs located in other Canadian provinces. The parallel trends assumption is not testable; however, in empirical work, it is standard to plot the outcome path for the variables of interest in order to assess whether treatment and control units exhibit similar or diverging trends. [Figure 2.3.1](#) and [Figure 2.3.2](#) report the average  $PM_{2.5}$  trends for 2000-2016 and 2000-2018, respectively, for British Columbian and control DAs; when using the [van Donkelaar \*et al.\* \(2019\)](#) data, it is immediate to dispel the possibility that outcome paths are parallel prior to treatment, while trends using the [Meng \*et al.\* \(2019\)](#) dataset are slightly less divergent, but noisy, especially for what concerns the control observations. Nonetheless, in both cases there is reason to suspect that a TWFE-DID regression would fail to identify the correct ATT. By giving equal weight to all control observations, TWFE-DID indeed will include units whose pre and post-treatment outcome paths fundamentally differ from those of DAs in British Columbia, and with greater potential for abatement.


It is also worth noting that in both cases, in addition to diverging trends, the level of  $PM_{2.5}$  pollution is almost always<sup>7</sup> lower for British Columbian vis-à-vis control DAs. Province-specific factors such as city morphology, more progressive environmental attitudes, different car fleet compositions and heterogeneous availability of alternative means of transportation could be the reason why trends and levels diverge across British Columbian CMAs and control Provinces. Treatment status in this instance is place-based and dependent on the political choice of an individual province – although sufficiently exogenous in timing ([Rivers and Schaufele, 2015](#)); nonetheless, as in [Sager and Singer \(2022\)](#), bias in the TWFE-DID estimator introduced by the

<sup>7</sup>Except for the [van Donkelaar \*et al.\* \(2019\)](#) dataset in the very first year of the panel, 2000.

failure of the parallel trends assumption needs to be acknowledged and a different estimation strategy can give rise to more precise estimates. In particular, our two alternative empirical methodologies are the synthetic control method for comparative case studies (SCM) (Abadie and Gardeazabal, 2003; Abadie, 2021) and the newly introduced synthetic difference-in-differences (SDID) methodology (Arkhangelsky *et al.*, 2021).



**Figure 2.3.1:** Trends in satellite PM<sub>2.5</sub> (Meng *et al.*, 2019), British Columbia vs control provinces, between 2000 and 2016. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line.



**Figure 2.3.2:** Trends in satellite PM<sub>2.5</sub> (van Donkelaar *et al.*, 2019), British Columbia vs control provinces, between 2000 and 2018. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line.

### 2.3.2 Synthetic control method and synthetic difference-in-differences

A traditional solution to diverging pre-treatment trends in empirical applications (usually with a unique treated unit, but extensible to the case of multiple treated units) is the SCM (Abadie and Gardeazabal, 2003; Abadie, 2021). In the BC carbon tax case, the SCM constructs a set of synthetic DAs as a weighted combination of control DAs by finding, for each treated unit  $i$ , a non-negative vector of weights  $\omega_i^{sc}$  summing to one, which ensures that each convex combination of the outcome variable for control units matches each outcome variable for the treated units for all periods up to the intervention date. Through this procedure, reliance on the parallel trends assumption, which is violated in our setting, is fundamentally weakened.

In order to combine the attractive features of both TWFE-DID (the inclusion of additive unit-specific and time-specific fixed effects), and SCM (reducing the reliance on the parallel trends assumption by weighting observations in order to ensure closely matched pre-intervention trends), Arkhangelsky *et al.* (2021)<sup>8</sup> have introduced a new method, synthetic difference-in-differences (SDID), which employs time and unit (two-way) fixed effects in the regression function (as in TWFE-DID), together with unit-specific weights (as in SCM) and time-specific weights which lessen the role of time periods that are largely divergent from post-treatment time periods. In a nutshell, for each treated unit SDID estimates: (1) unit weights  $\omega_i^{sdid}$  which underpin a synthetic control whose outcome is approximately parallel to the outcome for the treated unit; (2) time weights  $\lambda_t^{sdid}$  which ensure that the average post-treatment outcome for control units only differs by a constant from the weighted average of pre-treatment outcome for the control units – a synthetic pre-treatment period using controls. Once unit and time weights are calculated, SDID estimates a TWFE regression on the resulting panel, identifying the SDID ATT.

In order to formally explain how SDID combines features from TWFE-DID and SCM, let us consider a balanced panel with  $N$  observations and  $T$  time periods. In the British Columbian case, the outcome variable is  $PM_{2.5it}$ , and the binary treatment is  $TAX_{it}$ . Let  $i = 1, \dots, N_{tr}$  be the treated DAs in BC, and  $i = N_{tr} + 1, \dots, N_{co}$  be the DAs in control provinces. The baseline TWFE-DID regression problem can be expressed as:

$$(\hat{\tau}^{did}, \hat{\mu}, \hat{\eta}, \hat{\theta}) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \sum_{t=1}^T (PM_{2.5it} - \mu - \eta_i - \theta_t - \tau TAX_{it})^2 \right\} \quad (2.3)$$

Which is solved without the use of unit or time-specific weights, but with the inclusion of unit and time-specific fixed effects  $\eta_i$  and  $\theta_t$  as also illustrated in Equation 2.2. The SCM estimator, instead, does not employ unit fixed effects, but includes time fixed effects and unit-specific weights  $\omega_i^{sc}$ :

$$(\hat{\tau}^{sc}, \hat{\mu}, \hat{\eta}, \hat{\theta}) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \sum_{t=1}^T (PM_{2.5it} - \mu - \theta_t - \tau TAX_{it}) \hat{\omega}_i^{sc} \right\} \quad (2.4)$$

<sup>8</sup>This Section draws heavily on the seminal work of Arkhangelsky *et al.* (2021). Omissions are in the interest of avoiding redundancy.

Finally, the SDID estimator combines features from [Equation 2.3](#) and [Equation 2.4](#). Unit weights  $\hat{\omega}_i^{sdid}$  are chosen such that the pre-treatment outcome path of control DAs are parallel to those of the treated units<sup>9</sup>:

$$\omega_0 + \sum_{i=N_{tr}+1}^{N_{co}} \hat{\omega}_i^{sdid} PM_{2.5it} \approx \frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} PM_{2.5it} \quad (2.5)$$

Moreover, time weights  $\hat{\lambda}_t^{sdid}$  need to ensure that the pre-treatment levels for the control units differs from the post-treatment levels for the same units only by a constant. If we let  $t = 1, \dots, T$  be the total length of the panel,  $T_{pre}$  be the number of pre-intervention periods, and  $T_{post}$  be the number of post-intervention periods, the condition can be expressed as:

$$\lambda_0 + \sum_{t=1}^{T_{pre}} \hat{\lambda}_t^{sdid} PM_{2.5it} \approx \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^T PM_{2.5it} \quad (2.6)$$

Thus, the regression problem for the SDID estimator can be expressed as a weighted TWFE-DID problem which incorporates unit and time-specific fixed effects  $\eta_i$  and  $\theta_t$ , plus unit and time-specific weights  $\omega_i$  and  $\lambda_t$ , as illustrated in [Equation 2.7](#):

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\eta}, \hat{\theta}) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \mu - \eta_i - \theta_t - \tau TAX_{it})^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \right\} \quad (2.7)$$

In the remainder of the paper, we regard SDID as our preferred method in order to estimate the effect of the 2008 BC carbon tax on air pollution co-benefits, as the methodology allows us to overcome the apparent violation of the parallel trends assumption in conventional TWFE-DID; nonetheless, we estimate our main regression and robustness checks using all three of TWFE-DID, SCM and SDID, in order to inspect the direction of the eventual bias.

We calculate standard errors for all methods using the bootstrap variance estimation algorithm described in [Arkhangelsky \*et al.\* \(2021, p. 4109\)](#), with 200 replications. The procedure constructs a bootstrap dataset by sampling a portion of the original dataset with replacement, and computes the SDID estimator  $\tau^{(b)}$  on this subset for each iteration  $b$ . The variance is then defined as:

$$\hat{V}_\tau^b = \frac{1}{B} \sum_{b=1}^B \left( \hat{\tau}^{(b)} - \frac{1}{B} \sum_{b=1}^B \hat{\tau}^{(b)} \right)^2 \quad (2.8)$$

<sup>9</sup>Unit-specific weights are found using a regularisation parameter  $\zeta$ , as in [Doudchenko and Imbens \(2016\)](#), which aids the estimation strategy by increasing the dispersion of the weights and ensuring their uniqueness. When the intercept  $\omega_0$  and the regularisation parameter are set to 0, the unit weights  $\omega_i$  correspond to the SCM weights in [Abadie \*et al.\* \(2010\)](#). For further details on the procedure used to estimate  $\zeta$ , please refer to [Arkhangelsky \*et al.\* \(2021\)](#).

### 2.3.3 Heterogeneous Treatment Effects

In [Section 2.3.1](#) and [Section 2.3.2](#), the parameter identifying the effect of the 2008 BC carbon tax on  $PM_{2.5}$  emissions has always been assumed as constant across treated units. In essence, the methodology explained above computes a homogeneous ATT across DAs. Nonetheless, when dealing with disaggregated data within Census Metropolitan Areas, a homogeneously estimated ATT is likely to mask substantial heterogeneities across DAs which could be highly informative about the performance of different locations within metropolitan areas.

A first channel to examine is certainly that of pre-existing pollution levels: standard economic theory would in fact predict that emission abatement would happen first where the marginal cost of reducing emissions is lower, i.e. where pre-existing pollution is higher (that is, lower-hanging fruits would be picked earlier). This avenue is explored by [Sager and Singer \(2022\)](#) and [Auffhammer \*et al.\* \(2009\)](#), who find substantially higher reductions in  $PM_{2.5}$  and  $PM_{10}$  due to the Clean Air Act in nonattainment US census tracts that are more polluted in the three years preceding the implementation of the policy. In light of the results of [Carozzi and Roth \(2022\)](#) and [Borck and Schrauth \(2021\)](#), it is also worth exploring whether heterogeneity in air pollution reductions arises at different levels of the population density distribution: indeed, while densely populated areas have been shown to experience higher concentrations of  $PM_{2.5}$  particulate, usually population density is higher in city centres, where greater opportunities for substitution away from cars may arise. Following the same logic, we also test whether results are heterogeneously distributed according to road density and night-time lights. Lastly, an unexplored channel in carbon pricing is that of “spatial regressiveness”. A large body of research has shown that carbon pricing is regressive along income and wealth dimensions, but the relationship between the geographic distribution of income and wealth and the burden of carbon taxation is relatively underinvestigated. Therefore, we also inspect the heterogeneity of  $PM_{2.5}$  reductions with respect to the geographic distribution of mean dwelling value (as a proxy of wealth) and median income.

As the SDID methodology does not allow the inclusion of interactions in the estimation procedure, we split the treatment sample into quintiles of baseline<sup>10</sup> variables:  $PM_{2.5}$ , population density, road density, NTLs, average dwelling value and median income. We then run SDID separately for each quintile and, in [Section 2.4.3](#), we summarise the results graphically. Moreover, we also inspect a slightly more extreme trimming of the treated sample by dividing it in deciles of each baseline covariate and comparing the 90th percentile with the 10th percentile, in order to analyse the ratio of the effect at the top and bottom of the covariate distribution.

<sup>10</sup>For time-varying covariates we use the average of the three years prior to treatment as the baseline value; for variables retrieved from the Canadian census, we use their 2006 values, i.e. the last observation prior to the implementation of the carbon tax.

### 2.3.4 Mechanisms

It is paramount to be able to understand the potential mechanisms that determine the relationship between the carbon tax and  $PM_{2.5}$  concentrations. Ideally, when concerned with the estimation of  $PM_{2.5}$  reductions arising from the implementation of carbon pricing, we would look at DA-level reductions in motor fuel sales or in the quantity of vehicle kilometres travelled; however, these data are not available at the desired level of granularity for Canada between 2000 and 2018. The only precedent of a paper studying the relationship between the 2008 BC carbon tax and air quality in Canadian cities is the working paper of [Saberian \(2017\)](#), who restricts the analysis to Vancouver and uses monitoring stations data in order to infer her result – pointing to a worsening in air pollution following the carbon tax. The analysis of mechanisms leading to the result in [Saberian \(2017\)](#) highlights gasoline-to-diesel fuel switching as the potential causal driver of increased air pollution. However, the evidence is only anecdotal, as no evidence supporting the claim is presented in the study. Moreover, while Canadian province-level data on vehicle sales disaggregated by type of fuel is only available from 2011 onwards, the post-2011 trends in sales of diesel vehicles are relatively flat (See [Figure 2.A.3](#)), and the landscape seems to be dominated by gasoline cars (See [Figure 2.A.2](#)), suggesting that an eventual gas-to-diesel switch caused by the carbon tax incentive would have produced all of its results between July 2008 and January 2011 before bottoming out; the evidence for this conclusion is not very strong as a result. Another potential mechanism behind an increase in air pollution could derive from an exceptionally high rate of replacement in BC’s car fleet with respect to other Canadian provinces, caused by the willingness of BC’s residents to increase their cars’ fuel efficiency and realise savings at the pump. If the savings per each tank refuel were sufficient to offset the increase in gasoline prices due to the carbon tax, British Columbian residents could have potentially travelled more kilometres than prior to the tax, thereby increasing road congestion and hence pollution due to a rebound effect. As shown in [Figure 2.A.4](#) there has indeed been a rapid increase in truck and SUV sales in British Columbia after 2008; however, this increase is paralleled by similar jumps in truck sales in all large Canadian provinces<sup>11</sup>, and it thus seems implausible to attribute it to the marginal effect of the carbon tax in raising fuel prices.

In this paper, we instead exploit the information contained in the 2001, 2006, 2011 and 2016 waves of the Canadian census, which contains data on commute-to-work modes at the DA level for all Canadian CMAs. While the information on commute modes is not an exhaustive representation of all car trips made in each DA, the granularity of the data may shed light on whether residents of DAs located in British Columbia have adjusted their behaviour following the implementation of the carbon tax, substituting public transport or active commuting modes such as cycling and walking for car trips. Due to the structure of the data, collected at 5-year intervals, we are prevented from using the SCM and SDID methodology in this exercise; We thus resort to traditional TWFE-DID estimation of commute mode switching, analysing the data separately for each category of commute mode. In particular, we estimate

---

<sup>11</sup>Namely, Alberta, Ontario and Quebec.



the following equation:

$$Mode_{it} = \tau TAX_{it} + \theta_t + \eta_i + \epsilon_{it} \quad (2.9)$$

Where  $Mode_{it}$  is the share of each commute mode (high emission, low emission, public transport and zero emission, as described in [Section 2.2](#),  $Tax_{it}$  is the carbon tax DID binary variable also employed in [Equation 2.2](#),  $\theta_t$  and  $\eta_i$  are time and unit-specific fixed effects, and  $\epsilon_{it}$  is an idiosyncratic error term. In additional specifications, we also add a vector of controls  $X_{it}$  which account for population density, median income, and weather covariates (precipitation, maximum and minimum temperature, and wind speed), hence the estimating equation becomes:

$$Mode_{it} = \tau TAX_{it} + \beta X_{it} + \theta_t + \eta_i + \epsilon_{it} \quad (2.10)$$


We initially run the TWFE-DID regressions for the whole sample, without trimming the control pool. In further specifications, we restrict the control sample to the units which receive positive  $\omega_i$  weights in the SDID estimation of the main result, in order to ensure comparability across treatment and control cohorts and reduce the reliance on potentially violated parallel trends. Further, we retrieve the  $\omega_i$  weights from the SDID estimation and weigh our restricted TWFE-DID regressions with the SDID weights, assigning equal weights  $\frac{1}{N_{tr}}$  to the treatment cohort.

## 2.4 Results

### 2.4.1 Main Specification

#### With Meng *et al.* (2019) PM<sub>2.5</sub> data

In Figure 2.4.1 and Table 2.4.1, we report the results of the TWFE-DID, SCM and SDID regressions for the Meng *et al.* (2019) PM<sub>2.5</sub> dataset. Looking at the first panel of Figure 2.4.1, it is straightforward to infer how the baseline TWFE-DID strategy suffers from a clear violation of its foundational parallel trends assumption, as discussed in Section 2.3. The graphical representation of the regression analysis aids this line of interpretation: the DID ATT is indeed estimated by assuming that the outcome path of the treated units is parallel to the outcome path of the controls, thus the coefficient,  $\hat{\tau}^{did} = 0.393$  is biased. In the centre panel of Figure 2.4.1, we plot the average outcome path for the treated units and the traditional synthetic controls. The improvement in pre-treatment fit is dramatic, with minimal average deviation between British Columbian DAs and their controls, implying that the SCM performs well in giving positive weights to control units which best approximate treated DAs' outcome path and zero weight to control units which exhibit different trends. The direction of bias from the TWFE-DID regression is positive: SCM indeed identifies an effect of opposite sign to TWFE-DID,  $\hat{\tau}^{sc} = -0.142$ . In order to further augment the precision of our estimates, we also rely on the SDID estimator, graphically shown in the right-most panel of Figure 2.4.1. At the bottom of the panel, pre-treatment time-weights are represented in pink. The estimator gives positive temporal weights to periods for which the treated and control units exhibit similar trends.



**Figure 2.4.1:** Graphical results from DID, SCM and SDID for PM<sub>2.5</sub> concentrations, with Meng *et al.* (2019) data.

The SDID estimator does a particularly good job in imposing pre-treatment parallel trends in the years preceding the tax, even if weights  $\lambda_t$  are unevenly distributed over the pre-intervention period. However, negligible weights in 2007-2008 are reassuring,

given that a standard caveat in event-study methodologies is the excessive reliance on the single period immediately preceding the intervention (Heckman and Smith, 1999). The SDID procedure is able to select control units which exhibit pre-treatment trends that are almost perfectly parallel to BC’s outcome path, especially in the four year window preceding the intervention. The estimated ATT is  $\hat{\tau}^{sdid} = -0.363$ , therefore higher than the SCM ATT, and corresponding to a 5.2% reduction with respect to pre-intervention mean pollution levels. We regard SDID as the preferred methodology due to its greater flexibility; furthermore, due to the selection of a different, sparser set of control DAs. SDID selects indeed 6,258 control units among the untreated DAs and then performs DID on the matched sample with the inclusion of unit and time fixed effects to aid the estimation. On the contrary, SCM aggregates treated DAs into a single treated unit and assigns positive weight to just 10 units among 21,979 available DAs in the donor pool. Therefore, while SCM obtains a near-perfect fit pre-treatment, the outcome path of its synthetic unit heavily depends on the particular set of units receiving positive weights, which in our highly disaggregated setting is not ideal. In Figure B.1, we aggregate all 6,258 DAs which receive positive weights to the CMA level, in order to obtain the composition of synthetic BC in terms of percentages of other Canadian CMAs, in a similar vein to the traditional SCM methodology of Abadie (2021).


**Table 2.4.1:** Summary of  $\hat{\tau}$  point estimates and standard errors from all estimation methods, dependent variable from Meng *et al.* (2019).

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>
$\hat{\tau}$	0.3925	-0.1421	-0.3633
S.E.	0.0074	0.0809	0.0219
$N_{obs}$	432939	432939	432939

#### With van Donkelaar *et al.* (2019) PM<sub>2.5</sub> data

In order to gauge the robustness of this exercise to an alternative particulate emissions data source, we repeat the TWFE-DID, SCM and SDID estimation using the van Donkelaar *et al.* (2019) PM<sub>2.5</sub> dataset, which is available between 2000 and 2018. Notwithstanding the high correlation between the two outcome variables, as outlined in Figure 2.2.2, both the treatment and control pre-intervention trends exhibit some differences with respect to the Meng *et al.* (2019) dataset; however, the temporal location of peaks and troughs is generally respected, as is the relationship between the BC and control units outcome path. Indeed, DAs located in British Columbia always exhibit lower average annual concentrations of particulate pollution, and their PM<sub>2.5</sub> trend prior to 2008 appears to decline at an even faster pace than for control observations, barring some peaks in concentrations typical of the control provinces. The violation of the parallel trends assumption is once again highlighted in the graphical representation of the TWFE-DID regression in Figure 2.4.2, which, differently from the previous estimation, identifies a negative effect of the 2008 carbon

tax on emissions of  $\hat{\tau}^{did} = -0.495$  (see [Table 2.4.2](#)).



**Figure 2.4.2:** Graphical results from DID, SCM and SDID for  $PM_{2.5}$  concentrations, with [van Donkelaar \*et al.\* \(2019\)](#) data. Time weights  $\lambda_t$  are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.

The SCM, represented graphically in the middle panel of [Figure 2.4.2](#), again obtains a good pre-treatment fit in this instance, signalling that each British Columbian DA's outcome path is best approximated by a convex combination of control DAs rather than equally weighted control units. Furthermore, as evidenced in [Table 2.4.2](#), the direction of the TWFE-DID bias is confirmed: the SCM estimates a negative ATT of  $\hat{\tau}^{sc} = -0.709$ , therefore qualitatively reinforcing the SCM result of [Table 2.4.1](#). A similar conclusion can be drawn from the results of the SDID estimation, presented in the right-most panel of [Figure 2.4.2](#). The SDID procedure is able to select control units which exhibit pre-treatment trends that are almost perfectly parallel to BC's outcome path, with the exception of outlying time periods which receive zero-weights in the estimation. The estimated ATT is  $\hat{\tau}^{sdid} = -0.890$ , therefore slightly lower, but qualitatively similar to the SCM ATT. The two estimators are probably equivalent in this application. In terms of magnitude, both the SCM and SDID regressions identify a substantial drop in  $PM_{2.5}$  concentrations with respect to 2000-2007 levels, corresponding to a reduction of 10.9% from the pre-intervention  $PM_{2.5}$  mean for British Columbia.

## 2.4.2 Robustness Checks

We test the consistency of the main results by performing three additional analyses on sub-samples of the full dataset. First, we restrict the treated pool to DAs within the Vancouver metropolitan area, therefore excluding all DAs in the Abbotsford and Victoria CMAs. The resulting treatment cohort is comprised of 2874 DAs, vis-à-vis

**Table 2.4.2:** Summary of  $\hat{\tau}$  point estimates and standard errors from all estimation methods, dependent variable from van Donkelaar *et al.* (2019).

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>
$\hat{\tau}$	-0.4954	-0.7087	-0.8896
S.E.	0.0085	0.1540	0.0300
$N_{obs}$	483873	483873	483873

the 3490 DAs constituting the entire treatment unit pool; the control pool is kept the same, with 21989 control DAs. We then run TWFE-DID, SCM and SDID on the restricted sample, in order to be able to compare our results with those obtained by Saberian (2017). Perhaps unsurprisingly, given the relatively small number of DAs pertaining to the Abbotsford and Victoria CMAs, the results (reported in Figure B.3 and Table B.1) are qualitatively unchanged from the main regressions using the Meng *et al.* (2019) dataset.

We subsequently restrict the dataset to those DAs corresponding to the location of NAPS monitoring stations (see Figure 2.A.1), by spatially joining monitoring stations' locations to DAs<sup>12</sup>. Here, the size of the dataset is considerably restricted: the cross-section of DAs kept in the treated pool counts just 25 observations, while 106 DAs are kept in the control pool. This exercise allows us to infer whether our results also arise when considering just those locations in which pollution monitors have been established, thereby restricting the analysis to areas in which pollution is likely to be a greater concern. Once again, the results (presented in Figure B.4 and Table B.2), are qualitatively similar to the main specifications, a first signal that air pollution experiences greater reductions in places which exhibit greater levels of pre-intervention concerns about air quality<sup>13</sup>. Notably, the performance of the SDID estimator is not considerably worsened on this much smaller sample, with both estimators achieving a reasonable pre-treatment fit, and therefore identifying credible ATTs. On the contrary, the fit of the SCM seems to be substantially worse, and the method identifies a much higher ATT than in other specifications.

Lastly, we restrict the estimation window to 2000-2013, for two main reasons: (i) Checking whether the carbon tax ramp-up is the main mechanism behind the continuous reductions (the carbon tax was frozen at \$30/tCO<sub>2</sub> as mentioned in Section 2.1), and (ii) Comparing our results with Saberian (2017). The results, presented in Figure B.5 and Table B.3 identify a much higher ATT of  $\hat{\tau}^{sdid} = -0.67$ , which corroborates hypothesis (i) and is not comparable with the study by Saberian (2017), which identified an increase in particulate pollution over the same temporal window, possibly due to selection bias in the establishment of monitoring stations.

<sup>12</sup>We match DAs with all monitoring stations in the dataset, regardless of the date of establishment of each monitoring station, in order to maximise observations.

<sup>13</sup>Heterogeneity with respect to pre-intervention pollution levels is explored in greater depth in Section 2.4.3.

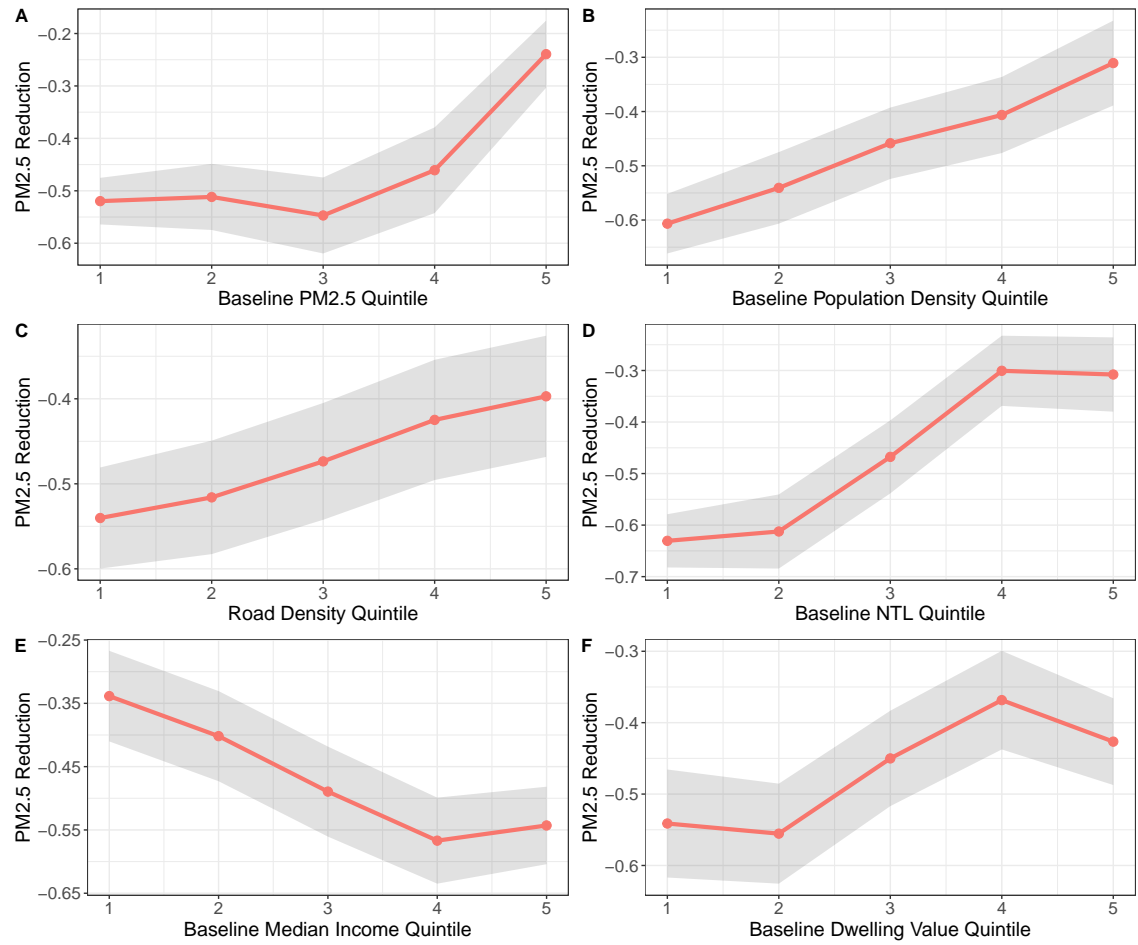
### 2.4.3 Heterogeneous Treatment Effects

After establishing robustness of our main result, the analysis turns to an in-depth evaluation of the heterogeneity of the findings with respect to a set of baseline DA-level characteristics. In light of the findings by [Sager and Singer \(2022\)](#), we first examine whether eventual reduction arise in locations which were more polluted at the baseline, i.e. where low-hanging benefits are more likely to be reaped. This is consistent with a long established tradition in the economics of pollution abatement, which points to the fact that abatement is likely to be cheaper where the initial level of pollution is higher. Moreover, in the context of air quality, and especially with respect to British Columbia, where average  $PM_{2.5}$  concentrations were low even before the introduction of carbon pricing, the left tail of the pollution distribution is likely to coincide with geographical areas in which road transportation is sparser and traffic levels are lower. Analysing whether the reductions arise at the right tail of the distribution is then useful in order to understand whether more problematic areas, with more potential for reaping higher benefits, see greater reductions.

In [Figure 2.4.3](#), we graphically<sup>14</sup> report the results from the quintile SDID regressions described in [Section 2.3](#), together with the 95% confidence interval around each estimate. We use quintiles for the distribution of each covariate of interest; further, we analyse the top and bottom of covariates' distribution via splitting the dataset in deciles, and calculating the 90th-10th decile ratio of ATTs (see [Table 2.4.3](#)). Quintile-SDID results for baseline  $PM_{2.5}$  concentrations are presented in panel A. It is immediate to infer that greater reductions arise in DAs with lower pollution levels between 2005 and 2007. Nonetheless, the ATT is in the  $[-0.2, -0.6]$  range for all quintiles; comparing the result with the baseline pollution levels reported in the first panel of [Figure 2.2.3](#), it is striking how most reductions appear to manifest in more peripheral areas, perhaps a sign of the necessity of driving through the most central locations within the urban network. The remaining panels analyse the results across population density (panel B), road density (C), NTL (D), median income (E) and dwelling value (F) quintiles. In panels B, C, and D, the ordering of ATT magnitudes is consistent with panel A, with greater reductions arising monotonically in less densely populated DAs, units with a lower concentration of road infrastructure, and DAs with lower emitted night-time luminosity. Taken in conjunction, these insights appear to confirm that the 2008 carbon tax was not effective in curtailing traffic in more central areas within British Columbian metropolitan areas, but rather had greater effect in peri-urban locations. More surprising is the result in panel E, which highlights the fact that richer DAs within metropolitan areas have experienced greater reductions, possibly reflecting an inverse relationship between density and income, but more importantly signalling that the pollution-income gap has increased as a result of the carbon tax. This result is a clear confirmation of the “spatial regressiveness” hypothesis, i.e. that a carbon tax is not only regressive on the vertical

---

<sup>14</sup>Graphical representation is necessary in order to ensure a parsimonious representation of these results. Tables are available on request.



**Figure 2.4.3:** Graphical SDID results: heterogeneity by quintiles of baseline characteristics for the treated sample. Panel A) Quintiles of baseline  $PM_{2.5}$ ; B) Quintiles of baseline population density; C) Quintiles of road density; D) Quintiles of baseline nightlights values. 95% confidence intervals are shaded in grey.

income dimension, but also geographically, with greater gains in better off areas. On the contrary, the relationship with dwelling value is consistent with those for density, possibly due to the association between centrality and land value.

In order to quantify the magnitude of the observed heterogeneity, in [Table 2.4.3](#), we report ATTs for the bottom 10% and the top 90% of baseline covariate distributions, together with the 90-10 ratio of ATTs. With respect to baseline pollution, the ATT for the 10th percentile is over 2 times greater than the ATT for the 90th percentile. Similar ratios are observed for all other characteristics, with the exception of income, for which reductions at the 90th percentile are 1.5 times stronger than at the 10th.

**Table 2.4.3:** SDID results for 10th and 90th percentiles, with 90-10 ratios.

	PM <sub>2.5</sub>	Pop.	Road	NTL	Income	Dwellings
<b>10th pct</b>	-0.5034	-0.6025	-0.5689	-0.6330	-0.3663	-0.5229
<b>90th pct</b>	-0.2081	-0.2541	-0.3967	-0.3691	-0.5618	-0.4639
<b>90-10 Ratio</b>	0.4134	0.4218	0.6973	0.5831	1.5336	0.8872

## 2.5 Mechanisms

As outlined in [Section 2.3](#), we analyse commute mode choices at the DA level as the main mechanism driving the results obtained in [Section 2.4](#). While commute mode is an imperfect measure of the number and type of trips made by British Columbians, we can rely on the same administrative level to the one used in the main analysis by retrieving information from the 2001, 2006, 2011, and 2016 Canadian censuses, thereby preserving granularity. In [Table 2.5.1](#), [Table 2.5.2](#), [Table C.1](#), and [Table C.2](#), we report TWFE-DID regression results employing the share of commuters using high-emissions, low-emissions, public transport, and zero-emissions commute modes, respectively. As the low-emissions transport mode is the sum of public transport and zero-emissions modes, we only report the results for low-emissions in the main text and present the sub-splits in the Appendix.

In all tables, column (1) is the baseline specification, a simple TWFE-DID regression with DA and year fixed effects and no controls, employing the full panel of DAs across census years. In column (4), we add weather controls for precipitation, maximum and minimum temperature, and wind speed, plus we control for the natural logarithm of population and median income. When employing the full pool of control DAs, the first result of note ([Table 2.5.1](#)) is that British Columbian DAs experience an average 4.2% reduction in the use of cars, taxis, and motorcycles, which rises to 4.7% when adding controls. This reduction is almost specular to the increase in the share of commuters using public transport, biking and walking to reach their workplace (see [Table 2.5.2](#)); moreover, as evidenced in [Table C.1](#) and [Table C.2](#), most of this increase (3.5-3.9%) is due to a higher reliance on public transport, while a



residual share of 0.5-0.7% is due to a switch to active commuting. All results are confirmed, and stronger in magnitude, when considering more restrictive specifications: columns (2) and (5) restrict the specifications in (1) and (4) to the DAs which receive positive weights in the main SDID regressions, in order to establish whether the mechanisms are effectively retrieved when employing the same set of observations on which the main ATT is estimated. Results are higher in magnitude by about 1%, jumping to a 5.3% reduction in high-emission commute modes in the case without controls. Here, the inclusion of control variables slightly dampens the impact to 5.2%; nonetheless, the specularity with the increase in low-emission commute modes is preserved. Finally, in columns (3) and (6) we further augment the TWFE-DID regressions by retrieving an including the weights from the main SDID regressions; we weigh all treatment observations equally and all control observations according to the value of  $\omega_i$  they receive after the data-driven SDID procedure. The magnitude of the increase in low-emission commute share increases further, to 5.5% in the case without covariates and is again dampened to 5.2% in the case with covariates. The hypothesis of a behavioural adjustment by BC citizens in response to the carbon tax is thus confirmed; residents of BC's DAs switch away from high-emissions commute modes towards low-emissions ones, with public transport as the main container for these substitutions.

**Table 2.5.1:** TWFE-DID results for high emissions commute mode

	High Emission Commute Mode					
	(1)	(2)	(3)	(4)	(5)	(6)
<b>DID</b>	-0.0417*** (0.0105)	-0.0527*** (0.0095)	-0.0549*** (0.0103)	-0.0466*** (0.0102)	-0.0519*** (0.0106)	-0.0516*** (0.0109)
DA FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Controls				✓	✓	✓
SDID control pool		✓	✓		✓	✓
SDID weights			✓			✓
R <sup>2</sup>	0.87184	0.83989	0.84360	0.87595	0.84508	0.84847
Adjusted R <sup>2</sup>	0.82896	0.78629	0.79124	0.83400	0.79267	0.79721
Observations	101,358	38,769	38,769	100,244	38,348	38,348

Notes: Standard errors clustered at the CMA level. \*\*\*:  $p < 0.01$ , \*\*:  $p < 0.05$ , \*:  $p < 0.1$

**Table 2.5.2:** TWFE-DID results for low emissions commute mode

	Low Emission Commute Mode					
	(1)	(2)	(3)	(4)	(5)	(6)
<b>DID</b>	0.0408*** (0.0111)	0.0516*** (0.0103)	0.0535*** (0.0110)	0.0457*** (0.0109)	0.0510*** (0.0113)	0.0506*** (0.0114)
DA FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Controls				✓	✓	✓
SDID control pool		✓	✓		✓	✓
SDID weights			✓			✓
R <sup>2</sup>	0.87321	0.84174	0.84532	0.87715	0.84674	0.84996
Adjusted R <sup>2</sup>	0.83078	0.78876	0.79354	0.83560	0.79490	0.79920
Observations	101,358	38,769	38,769	100,244	38,348	38,348

Notes: Standard error clustered at the CMA level. \*\*\*:  $p < 0.01$ , \*\*:  $p < 0.05$ , \*:  $p < 0.1$

## 2.6 Health Gains

In order to understand the magnitude of the economic co-benefits from air pollution reductions arising due to the 2008 carbon tax, in this section we convert the  $PM_{2.5}$  estimates from [Section 2.4](#) into a monetary quantification of the associated health gains. Notwithstanding the relatively low concentrations of particle pollution in the the British Columbian context, where pre-treatment air quality was of substantial better quality than in other North American locations (e.g. in the USA), it is important to note that the concept of “safe” thresholds for particle pollution concentrations is more normative than positive. Indeed, some studies (e.g. [Krewski \*et al.\*, 2009](#)) have highlighted that the marginal benefits from abatement may be nonlinear in baseline concentrations, with lower gains from abatement at higher levels of baseline air pollution. Hence, any improvement in air quality is likely to carry significant benefits in terms of reductions in mortality rates; moreover, the estimates reported in this section are a lower bound of the gains from local pollution reductions, as  $PM_{2.5}$  has been shown to have a multidimensional impact, ranging from health to productivity, to cognition and to long-term impacts on the formation of human capital.

Drawing from [Fowlie \*et al.\* \(2019\)](#) and [Carozzi and Roth \(2022\)](#), our approach consists of two steps. We first estimate the impact of a reduction in  $PM_{2.5}$  concentrations in terms of mortality reductions, using concentration-response (“hazard”) functions derived from the environmental health literature. Second, we retrieve the central estimate of the willingness to pay (WTP) to avoid a premature death from [Health Canada \(2021\)](#) and [Chestnut and De Civita \(2009\)](#)<sup>15</sup>, and multiply the mortality reductions estimated in the first step by the central estimate of the Value of a Statistical Life (VSL), equal to \$6.5 million in 2007 Canadian dollars, for each DA in the census metropolitan areas of Vancouver, Victoria, and Abbotsford.

The traditional form of the Cox proportional hazard model used in the environmental health literature is the log-linear regression reported in [Fowlie \*et al.\* \(2019\)](#):

$$\ln(\gamma) = \zeta + \alpha PM_{2.5} \quad (2.11)$$

Where  $\ln(\gamma)$  is the natural logarithm of mortality risk,  $\zeta = \ln(Z)$ , and  $PM_{2.5}$  are the local pollution concentrations. The term  $Z$  is a vector of covariates other than  $PM_{2.5}$  which impact mortality, and can be rewritten as  $Z = Z_0 + \exp(\beta_1 x_1 + \dots + \beta_n x_n)$ , with  $Z_0$  being the mortality risk when all covariates are zero. Indicating  $\gamma_0$  as the baseline mortality risk, and rearranging terms<sup>16</sup>, the change in mortality rate  $\Delta\gamma$

<sup>15</sup>It must be noted that the reported estimate for the Value of a Statistical Life does not reflect directly the economic value of an individually identified person’s life, but rather the aggregation of estimates of the WTP for a small reduction in mortality risk. Using the VSL central estimate of \$6,500,000, for example, the average Canadian would be willing to pay \$65 to reduce the risk of premature death by 1 out of 100,000.

<sup>16</sup>The derivation is as follows ([Carozzi and Roth, 2022](#)):

$$\Delta\gamma = Z(e^{\alpha PM_{2.5}^0} - e^{\alpha PM_{2.5}^1}) \rightarrow \Delta\gamma = Z e^{\alpha PM_{2.5}^0} \left[ 1 - e^{-\alpha(PM_{2.5}^0 - PM_{2.5}^1)} \right]$$

can be related to the change in pollution levels  $\Delta PM_{2.5}$  with the following equation:

$$\Delta\gamma = \gamma_0 \left( 1 - \frac{1}{e^{\alpha\Delta PM_{2.5}}} \right) \quad (2.12)$$

In order to find the total number of deaths for each DA associated with the above change in mortality rate  $\Delta\gamma$ , this quantity needs to be multiplied by the population of each DA<sup>17</sup>:

$$\Delta Deaths_i = Population_i \left[ \gamma_0 \left( 1 - \frac{1}{e^{\alpha\Delta PM_{2.5}}} \right) \right] \quad (2.13)$$

And finally, the monetary health gains in terms of mortality reductions at the DA level,  $\Delta Y_i$ , are obtained by multiplying the above estimates by the VSL figure of \$6.5 million CAD obtained from [Health Canada \(2021\)](#):

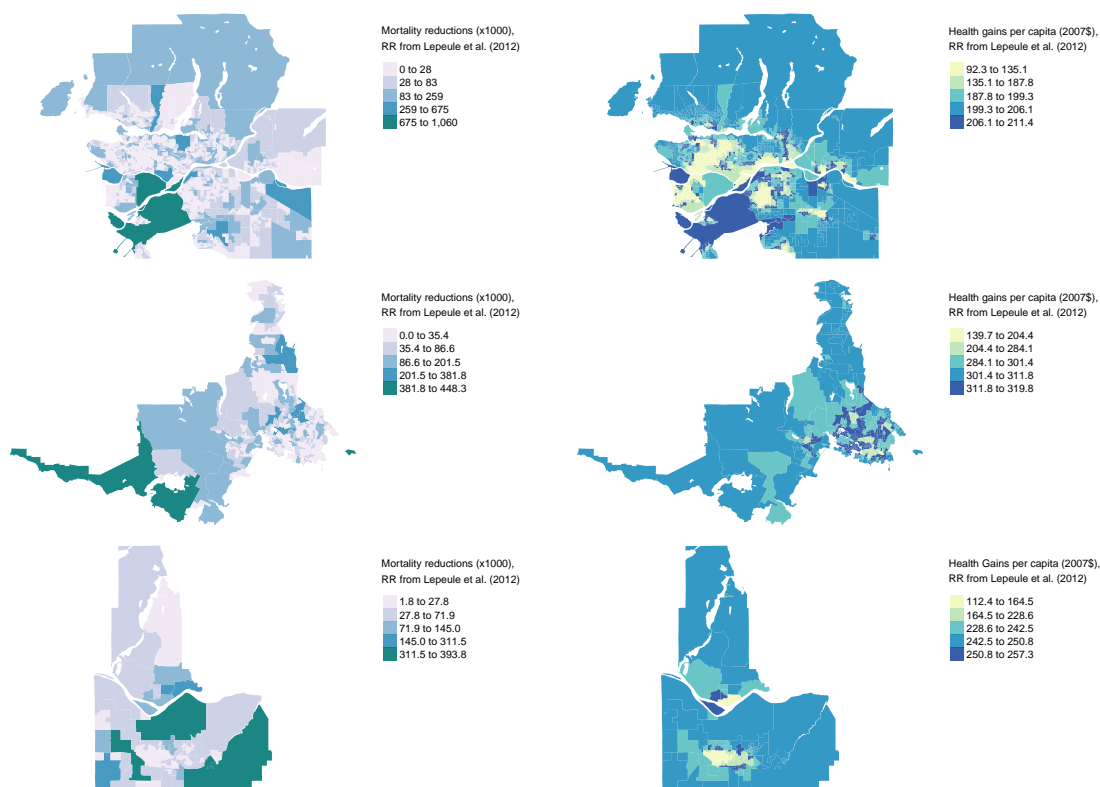
$$\Delta Y_i = VSL * \Delta Deaths_i \quad (2.14)$$

Hence, in order to estimate the model outlined in [Equation 2.12](#), and thus obtain mortality rate changes at the DA level, we first need to estimate the baseline mortality rate  $\gamma_0$ . Consistently with the literature, we obtain data for deaths due to lung cancers, all circulatory diseases, and all respiratory diseases from the ICD.10 selected causes of death at the CMA level from Statistics Canada. We divide total deaths due to the listed causes by total CMA population, and assign the resulting (baseline) mortality rates to all DAs in a given CMA. The parameter  $\alpha$  is usually not directly indicated in epidemiology studies, which instead report the relative risk (RR) increase due to a given increase in  $PM_{2.5}$ . For instance, [Lepeule et al. \(2012\)](#) report an all-cause RR of 1.14 associated with a  $\Delta PM_{2.5}$  of  $10 \mu g/m^3$ , while [Krewski et al. \(2009\)](#)'s estimate of RR is 1.06. However, it is straightforward to retrieve  $\alpha$  by exploiting the relationship between RR and  $\Delta PM_{2.5}$ , as reported in [Carozzi and Roth \(2022\)](#):  $\alpha = \ln(RR)/\Delta PM_{2.5}$ .

We employ these two estimates, in combination with the estimated  $PM_{2.5}$  reductions for each quintile of the pre-intervention  $PM_{2.5}$  distribution, in order to calculate the gains from mortality reductions at the DA level for the three CMAs included in the treated sample: Vancouver, Victoria and Abbotsford. In [Figure 2.6.1](#), we visually report the results of this exercise for each CMA, using  $RR = 1.14$  as estimated by [Lepeule et al. \(2012\)](#) (visual results using the RR estimate from [Krewski et al. \(2009\)](#) are reported in [Figure D.1](#)).

<sup>17</sup>We use the baseline population level, that is, the population of each DA in the year 2008.

The left panel maps the estimated mortality reductions per 1000 people (estimated according to Equation 2.13), while the right panel shows the associated per capita health gains, estimated via Equation 2.14. The median monetary gains due to the estimated reductions in  $PM_{2.5}$  are large: \$198 when using the Lepeule *et al.* (2012) RR and \$88 with the RR from Krewski *et al.* (2009)<sup>18</sup>.



**Figure 2.6.1:** Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from Lepeule *et al.* (2012), for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs.

To put this in context, the monetary value of per capita air quality co-benefits from the BC carbon tax amounts to 66% of the Low-income Climate Action Tax Credit, i.e. the carbon tax rebate for low-income families, for a family of four in year 2011; the total air quality co-benefit accruing to the same family would thus exceed the carbon tax rebate in monetary terms. The spatial distribution of these gains shows substantial heterogeneity: in particular, it is once again striking how air pollution co-benefits seem to be concentrated in peri-urban areas and positively correlated with income (see also Figure D.2). The results confirms that carbon taxation appears to be spatially regressive over urban areas, with greater co-benefits arising in higher income, low pollution tracts, and thus underpinning an increasing “pollution-income gap”, as also evidenced in Section 2.4.

<sup>18</sup>The same gains are \$402 and \$178, respectively, if calculated using the ATT estimated with the van Donkelaar *et al.* (2019)  $PM_{2.5}$  dataset instead of Meng *et al.* (2019).

## 2.7 Conclusions

A large body of literature has long posited that co-benefits arising from carbon taxation are large in magnitude, and may partially or fully offset the costs of complying with mitigation when converted to per capita monetary values. In order to gauge the entity of the full benefits of carbon taxation, it is essential to incorporate the monetary value of the reduction of local externalities in cost-benefit analyses, since policies aimed at the reduction of CO<sub>2</sub> emissions may imply substantial complementarities with reductions in local pollutants.

This paper undertakes the first rigorous empirical analysis of the impact of a carbon tax on particulate matter concentrations, focussing on PM<sub>2.5</sub>, showing that the introduction of carbon pricing can significantly improve local air quality. After the implementation of the 2008 carbon tax, PM<sub>2.5</sub> concentrations dropped by 5.2-10.9% in British Columbian Dissemination Areas, compared to a counterfactual obtained through the synthetic difference-in-differences estimator, which is able to produce a parallel pre-treatment trajectory for the average treated unit and its synthetic control. The estimated reductions are significantly heterogeneous across the geography of British Columbian census metropolitan areas, with greater reductions found in less polluted, less dense, peripheral areas and in richer neighbourhoods. These results highlight a spatial dimension of the regressive nature of carbon pricing: a carbon tax is indeed prone to exacerbate the pre-existing pollution gap, and the pollution-income gap. A significant driver of the air quality improvement is found in transport mode switching, with a 5.3% reduction in high-emission transport modes, mostly in favour of public transport: this insight is a signal that complementary policies such as an incentive for alternative transport mode may entail further reductions in air pollution if rolled out contemporaneously to a carbon tax. Moreover, the integration of multiple instruments aimed at multiple externalities may be able to give rise to increased benefits stemming from policy additionality.

Finally, the analysis converts the improvements in air quality into reductions in mortality rates and monetary health gains from co-benefits of carbon taxation. With a median estimate of the health gains of \$198 per capita, the health “savings” are large and comparable to the rebates offered to low-income families in British Columbia to mitigate the impact of the tax on their disposable income. Health benefits are heterogeneously distributed across metropolitan areas and accrue primarily to neighbourhoods in higher income brackets, once again highlighting the need for redistribution in the design of climate policy.

# References

- Abadie, Alberto (2021). “Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects”. *Journal of Economic Literature* 59.2, pp. 391–425.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2010). “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program”. *Journal of the American Statistical Association* 105.490, pp. 493–505.
- Abadie, Alberto and Gardeazabal, Javier (2003). “The Economic Costs of Conflict : A Case Study of the Basque Country”. *American Economic Review* 93.1, pp. 113–132.
- Abatzoglou, John, Dobrowski, Solomon, Parks, Sean, and Hegewisch, Katherine (2018). “TerraClimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015”. *Nature Scientific Data* 5, p. 170191.
- Aguilar-Gomez, Sandra, Dwyer, Holt, Graff Zivin, Joshua S, and Neidell, Matthew J. (2022). *This is Air: The "Non-Health" Effects of Air Pollution*. Working Paper 29848. National Bureau of Economic Research.
- Ambec, Stefan and Coria, Jessica (2013). “Prices vs quantities with multiple pollutants”. *Journal of Environmental Economics and Management* 66.1, pp. 123–140.
- Andersson, Julius J. (2019). “Carbon Taxes and CO2 Emissions: Sweden as a Case Study”. *American Economic Journal: Economic Policy* 11.4, pp. 1–30.
- Arkhangelsky, Dmitry, Athey, Susan, Hirshberg, David A., Imbens, Guido W., and Wager, Stefan (2021). “Synthetic Difference-in-Differences”. *American Economic Review* 111.12, pp. 4088–4118.
- Auffhammer, Maximilian, Bento, Antonio, and Lowe, Scott (2009). “Measuring the effects of the Clean Air Act Amendments on ambient PM10 concentrations: The critical importance of a spatially disaggregated analysis”. *Journal of Environmental Economics and Management* 58, pp. 15–26.
- Bondy, Malvina, Roth, Sefi, and Sager, Lutz (2020). “Crime Is in the Air: The Contemporaneous Relationship between Air Pollution and Crime”. *Journal of the Association of Environmental and Resource Economists* 7.3, pp. 555–585.
- Borck, Rainald and Schrauth, Philipp (Jan. 2021). “Population density and urban air quality”. *Regional Science and Urban Economics* 86, p. 103596.
- Carozzi, Felipe and Roth, Sefi (2022). “Dirty density: Air quality and the density of American cities”. *Journal of Environmental Economics and Management*, p. 102767.

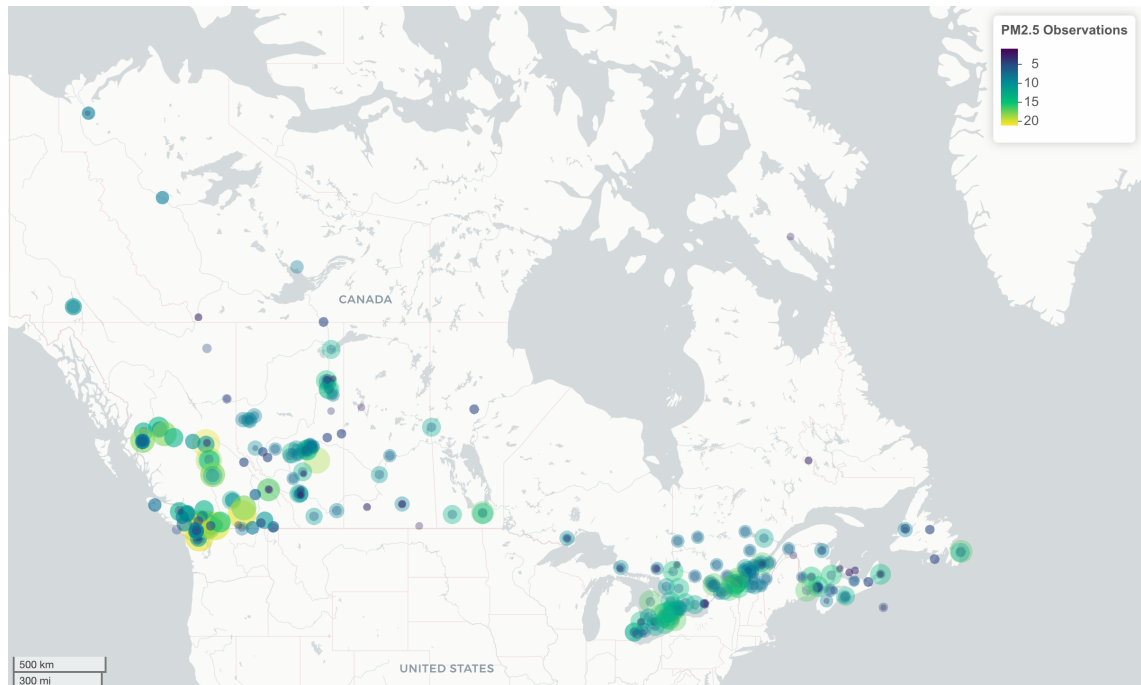
- Chestnut, L.G. and De Civita, P. (2009). “Economic valuation of mortality risk reduction : review and recommendations for policy and regulatory analysis”. *Government of Canada. Policy Research Initiative*.
- Currie, Janet, Voorheis, John, and Walker, Reed (2020). *What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality*. en. Tech. rep. w26659. Cambridge, MA: National Bureau of Economic Research, w26659. (Visited on 08/24/2022).
- Doudchenko, Nikolay and Imbens, Guido W. (2016). *Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis*. Working Paper 22791. National Bureau of Economic Research.
- Douenne, Thomas (2020). “The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy”. *The Energy Journal* 41.3, pp. 231–254.
- Ebenstein, Avraham, Lavy, Victor, and Roth, Sefi (2016). “The Long-Run Economic Consequences of High-Stakes Examinations: Evidence from Transitory Variation in Pollution”. *American Economic Journal: Applied Economics* 8.4, pp. 36–65.
- Fowlie, Meredith, Rubin, Edward, and Walker, Reed (2019). “Bringing Satellite-Based Air Quality Estimates Down to Earth”. *AEA Papers and Proceedings* 109, pp. 283–88.
- Freeman, Richard, Liang, Wenquan, Song, Ran, and Timmins, Christopher (2019). “Willingness to pay for clean air in China”. *Journal of Environmental Economics and Management* 94, pp. 188–216.
- Gehrsitz, Markus (May 2017). “The effect of low emission zones on air pollution and infant health”. *Journal of Environmental Economics and Management* 83, pp. 121–144.
- Graff Zivin, Joshua and Neidell, Matthew (2012). “The Impact of Pollution on Worker Productivity”. *American Economic Review* 102.7, pp. 3652–73.
- Health Canada (2021). “Health Impacts of Air Pollution in Canada: Estimates of morbidity and premature mortality outcomes – 2021 Report”.
- Heckman, James J. and Smith, Jeffrey A. (1999). “The Pre-Programme Earnings Dip and the Determinants of Participation in a Social Programme. Implications for Simple Programme Evaluation Strategies”. *The Economic Journal* 109.457, pp. 313–348.
- Henderson, J. Vernon, Storeygard, Adam, and Weil, David N. (2012). “Measuring Economic Growth from Outer Space”. *American Economic Review* 102.2.
- Jbaily, Abdulrahman, Zhou, Xiaodan, Liu, Jie, Lee, Ting-Hwan, Kamareddine, Leila, Verguet, Stéphane, and Dominici, Francesca (2022). “Air pollution exposure disparities across US population and income groups”. *Nature* 601.7892, pp. 228–233.
- Knittel, Christopher R. and Sandler, Ryan (2011). *Cleaning the Bathwater with the Baby: The Health Co-Benefits of Carbon Pricing in Transportation*. Working Papers 1115. Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research.
- Krewski, D. *et al.* (2009). “Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality”. *Research Reports, Health Effects Institute*. Discussion 115-36.140, pp. 5–114.



- Lepeule, J., Laden, F., Dockery, D., and Schwartz, J. (2012). “Chronic exposure to fine particles and mortality: an extended follow-up of the Harvard Six Cities study from 1974 to 2009”. *Environmental Health Perspectives* 120.7, pp. 965–970.
- Li, Mingwei, Zhang, Da, Li, Chiao-Ting, Mulvaney, Kathleen M., Selin, Noelle E., and Karplus, Valerie J. (2018). “Air quality co-benefits of carbon pricing in China”. *Nature Climate Change* 8.5, pp. 398–403.
- Meng, Jun, Li, Chi, Martin, Randall V., Donkelaar, Aaron van, Hystad, Perry, and Brauer, Michael (2019). “Estimated Long-Term (1981–2016) Concentrations of Ambient Fine Particulate Matter across North America from Chemical Transport Modeling, Satellite Remote Sensing, and Ground-Based Measurements”. en. *Environmental Science & Technology* 53.9, pp. 5071–5079.
- Nehiba, Cody (2022). “Correcting Heterogeneous Externalities: Evidence from Local Fuel Taxes”. en. *Journal of the Association of Environmental and Resource Economists* 9.3, pp. 495–529.
- Nordhaus, William (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Parry, Ian, Veung, Chandra, and Heine, Dirk (2015). “How Much Carbon Pricing is in Countries’ own Interests? The Critical Role of Co-Benefits”. *Climate Change Economics* 06.04, p. 1550019.
- Pestel, Nico and Wozny, Florian (2021). “Health effects of Low Emission Zones: Evidence from German hospitals”. *Journal of Environmental Economics and Management* 109, p. 102512.
- Rivers, Nicholas and Schaufele, Brandon (2015). “Salience of carbon taxes in the gasoline market”. *Journal of Environmental Economics and Management* 74, pp. 23–36.
- Rose, Amy N., McKee, Jacob J., Sims, Kelly M., Bright, Edward A., Reith, Andrew E., and Urban, Marie L. (2020). *LandScan 2019*.
- Saberian, Soodeh (2017). “The Negative Effect of the BC Carbon Tax on Vancouver Air Quality: A Good Climate for Bad Air?” *Working Paper*.
- Sager, Lutz and Singer, Gregor (2022). “Clean Identification? The Effects of the Clean Air Act on Air Pollution, Exposure Disparities and House Prices”. *Grantham Research Institute on Climate Change and the Environment Working Paper No. 376*, p. 57.
- Sarmiento, Luis, Wagner, Nicole, and Zaklan, Aleksandar (2022). “The Air Quality and Well-Being Effects of Low Emission Zones”. en. *SSRN Electronic Journal*.
- Shindell, Drew T., Lee, Yunha, and Faluvegi, Greg (2016). “Climate and health impacts of US emissions reductions consistent with 2°C”. *Nature Climate Change* 6.5, pp. 503–507.
- Vandyck, Toon, Keramidias, Kimon, Kitous, Alban, Spadaro, Joseph, Van Dingenen, Rita, Holland, Mike, and Saveyn, Bert (2018). “Air quality co-benefits for human health and agriculture counterbalance costs to meet Paris Agreement pledges”. *Nature Communications* 9.
- Weitzman, Martin L. (2016). *How a Minimum Carbon Price Commitment Might Help to Internalize the Global Warming Externality*. Working Paper 22197. National Bureau of Economic Research.

- Weitzman, Martin L. (2015). “Internationally-Tradable Permits Can Be Riskier for a Country than an Internally-Imposed Carbon Price”. *SSRN Electronic Journal*.
- Wen, Jeff and Burke, Marshall (2022). “Wildfire Smoke Exposure Worsens Learning Outcomes”. *Nature Sustainability* 5, pp. 920–921.
- West, Jason, Smith, Steven, Silva, Raquel, Naik, Vaishali, Zhang, Yuqiang, Adelman, Zachariah, Fry, Meridith, Anenberg, Susan, Horowitz, Larry, and Lamarque, Jean-François (2013). “Co-benefits of Global Greenhouse Gas Mitigation for Future Air Quality and Human Health”. *Nature Climate Change* 3.
- World Bank (2022). *State and Trends of Carbon Pricing 2022*. World Bank, Washington, DC.
- Xuecao, Li, Zhou, Yuyu, Zhao, Min, and Xia, Zhao (2020). “A harmonized global nighttime light dataset 1992–2018”. *Nature Scientific Data* 7.
- Zhang, Wen-Wen, Zhao, Bin, Ding, Dian, Sharp, Basil, Gu, Yu, Xu, Shi-Chun, Xing, Jia, Wang, Shu-Xiao, Liou, Kuo-Nan, and Rao, Lan-Lan (2021). “Co-benefits of subnationally differentiated carbon pricing policies in China: Alleviation of heavy PM2.5 pollution and improvement in environmental equity”. *Energy Policy* 149, p. 112060.
- van Donkelaar, Aaron, Martin, Randall V., Li, Chi, and Burnett, Richard T. (2019). “Regional Estimates of Chemical Composition of Fine Particulate Matter Using a Combined Geoscience-Statistical Method with Information from Satellites, Models, and Monitors”. *Environmental Science & Technology* 53.5, pp. 2595–2611.
- von Bergmann, Jens (2021). *tongfen: R package to Make Data Based on Different Geographies Comparable*. R package version 0.3.3. URL: <https://mountainmath.github.io/tongfen/>.
- von Bergmann, Jens, Shkolnik, Dmitry, and Jacobs, Aaron (2022). *cancensus: R package to access, retrieve, and work with Canadian Census data and geography*. R package version 0.5.3. URL: <https://mountainmath.github.io/cancensus/>.

## 2.A Additional Descriptive Statistics



**Figure 2.A.1:** Availability of PM<sub>2.5</sub> readings in the National Atmospheric Surveillance Program database between 2000 and 2018. Lighter colours indicate higher availability of readings, and hence, monitoring stations which were added earlier.

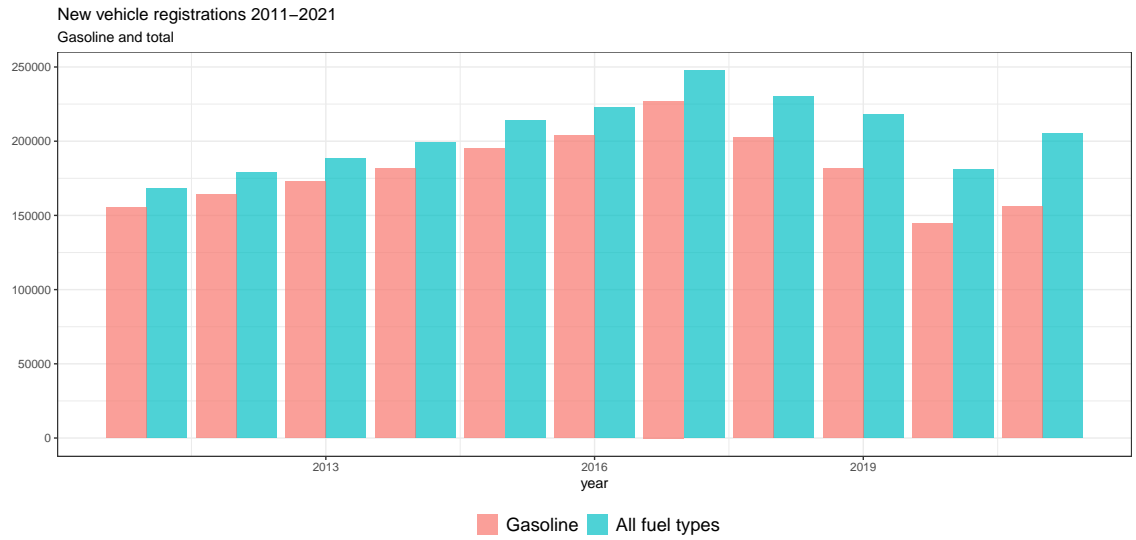


Figure 2.A.2: Gasoline vs Other fuel car sales in BC 2011-2021.


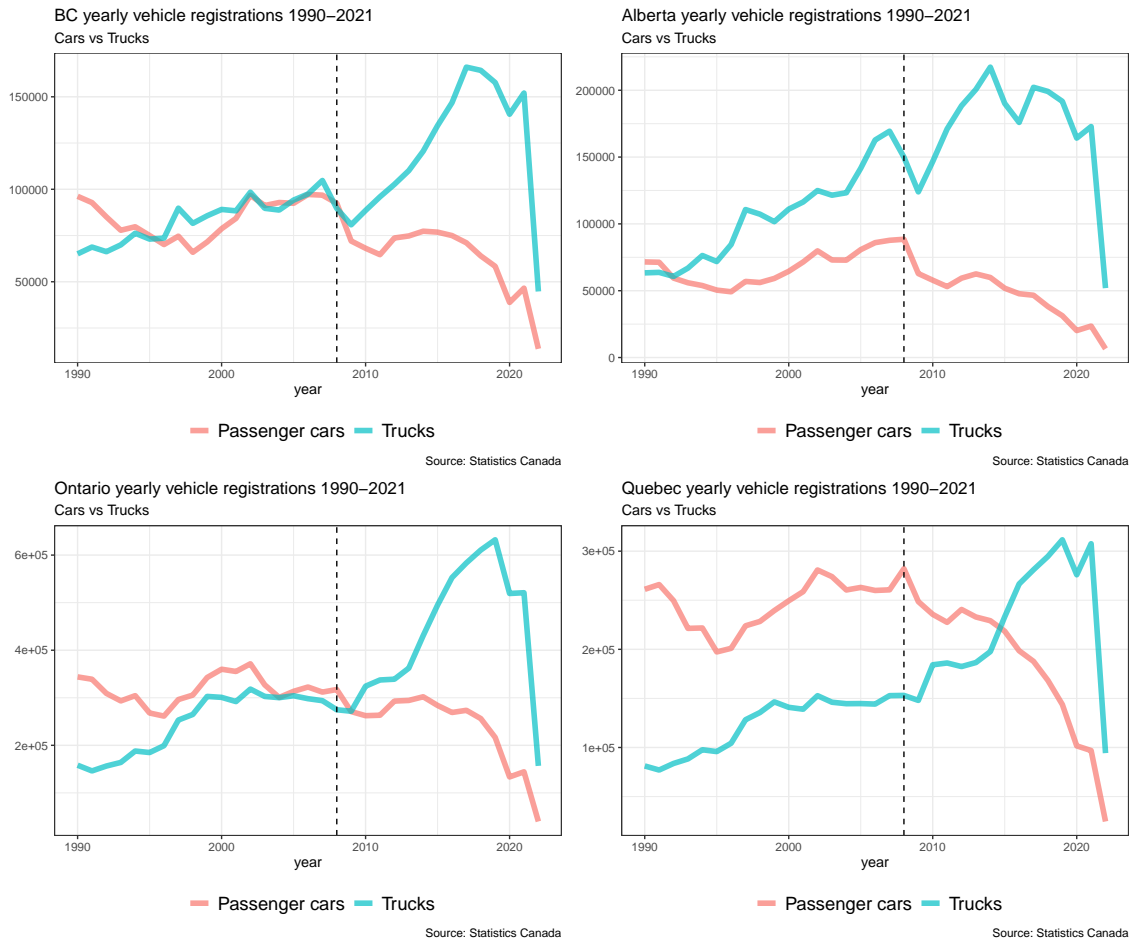


Figure 2.A.3: Minor fuel groups car sales in BC 2011-2021.



**Figure 2.A.4:** Passenger cars vs Truck and SUV sales, large Canadian Provinces, 1990-2021.

**Table 2.A.1:** Summary Statistics, 2000-2007


Variable	Control Provinces			British Columbia		
	N	Mean	SD	N	Mean	SD
PM <sub>2.5</sub> (van Donkelaar <i>et al.</i> , 2019)	175870	9.52	1.54	27920	8.06	1.19
PM <sub>2.5</sub> (Meng <i>et al.</i> , 2019)	175870	8.61	2.07	27920	6.95	1.39
Pop. Density (Rose <i>et al.</i> , 2020)	175912	3358.26	3375.33	27920	3169.94	2136.98
Road Density	175800	12.07	4.71	27904	11.77	4.50
NTL (Xuecao <i>et al.</i> , 2020)	175912	59.62	6.18	27920	53.00	17.38
Dwelling Value	43254	237314.81	196990.72	6885	417272.84	258798.32
Median Income	43978	26341.65	9088.59	6980	25055.65	8090.75
High Emission Commute %	43701	74.93	18.33	6926	77.64	16.37
Low Emission Commute %	43701	24.52	18.26	6926	21.62	16.26
Public Transport Commute %	43701	17.02	14.48	6926	13.06	10.68
Zero Emission Commute %	43701	7.50	9.43	6926	8.56	10.79
Precipitation (Abatzoglou <i>et al.</i> , 2018)	175768	74.05	21.74	27920	131.20	37.06
Max Temperature (Abatzoglou <i>et al.</i> , 2018)	175768	11.93	1.58	27920	14.55	0.66
Min Temperature (Abatzoglou <i>et al.</i> , 2018)	175768	1.74	2.50	27920	6.46	0.62
Wind Speed (Abatzoglou <i>et al.</i> , 2018)	175768	3.63	0.49	27920	2.98	0.16

**Table 2.A.2:** Summary Statistics, 2008-2018


Variable	Control Provinces			British Columbia		
	N	Mean	SD	N	Mean	SD
PM <sub>2.5</sub> (van Donkelaar <i>et al.</i> , 2019)	241865	8.15	1.51	38390	6.09	0.95
PM <sub>2.5</sub> (Meng <i>et al.</i> , 2019)	197888	7.35	1.73	31410	6.07	1.10
Population Density (Rose <i>et al.</i> , 2020)	241879	3614.39	3365.36	38390	3478.58	2305.69
Road Density	241725	12.07	4.71	38368	11.77	4.50
NTL (Xuecao <i>et al.</i> , 2020)	241879	60.31	6.87	38390	58.34	7.74
Dwelling Value	42437	415834.83	269717.66	6818	886776.76	584340.83
Median Income	43978	33324.06	11718.48	6980	31772.63	9765.79
High Emission Commute %	43806	74.39	20.20	6955	72.94	18.75
Low Emission Commute %	43806	25.06	20.12	6955	26.25	18.61
Public Transport Commute %	43806	18.85	15.89	6955	18.38	13.38
Zero Emission Commute %	43806	6.21	10.15	6955	7.87	11.32
Precipitation (Abatzoglou <i>et al.</i> , 2018)	241681	77.57	21.99	38390	134.58	37.33
Max Temperature (Abatzoglou <i>et al.</i> , 2018)	241681	12.32	1.76	38390	14.58	0.86
Min Temperature (Abatzoglou <i>et al.</i> , 2018)	241681	2.11	2.59	38390	6.56	0.80
Wind Speed (Abatzoglou <i>et al.</i> , 2018)	241681	3.64	0.48	38390	3.00	0.19

## 2.B Additional Regression Results

### 2.B.1 Composition of Main SDID Control




**Figure B.1:** Composition of the synthetic unit of [Figure 2.4.1](#).



**Figure B.2:** Composition of the synthetic unit of [Figure 2.4.2](#).

## 2.B.2 DAs in the Vancouver CMA




**Figure B.3:** Graphical results from DID, SCM and SDID for  $PM_{2.5}$  concentrations, with [Meng et al. \(2019\)](#) data, dataset restricted to DAs in the Vancouver CMA.

**Table B.1:** Summary of  $\hat{\tau}$  point estimates and standard errors from all estimation methods, dependent variable from [Meng et al. \(2019\)](#), dataset restricted to DAs in the Vancouver CMA.

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>
$\hat{\tau}$	0.4061	-0.1062	-0.3014
S.E.	0.0071	0.0702	0.0225
$N_{obs}$	422467	422467	422467



## 2.B.3 DAs matching NAPS Monitoring Stations




**Figure B.4:** Graphical results from DID, SCM and SDID for  $PM_{2.5}$  concentrations, with [Meng \*et al.\* \(2019\)](#) data, dataset restricted to DAs matching NAPS monitoring stations' locations.

**Table B.2:** Summary of  $\hat{\tau}$  point estimates and standard errors from all estimation methods, dependent variable from [Meng \*et al.\* \(2019\)](#), dataset restricted to DAs matching NAPS monitoring stations' locations.

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>
$\hat{\tau}$	0.132	-0.865	-0.288
S.E.	0.117	0.128	0.097
$N_{obs}$	2227	2227	2227

## 2.B.4 Post-treatment period limited to 2013



**Figure B.5:** Graphical results from DID, SCM and SDID for  $PM_{2.5}$  concentrations, with Meng *et al.* (2019) data, dataset restricted to 2013.

**Table B.3:** Summary of  $\hat{\tau}$  point estimates and standard errors from all estimation methods, dependent variable from Meng *et al.* (2019), dataset restricted to 2013.

	<i>DID</i>	<i>SCM</i>	<i>SDID</i>
$\hat{\tau}$	0.0547	-0.2723	-0.6703
S.E.	0.0081	0.0803	0.0341
$N_{obs}$	432939	432939	432939

## 2.C Additional Mechanisms Tables

**Table C.1:** TWFE-DID results for public transport

	Public Transport Commute Mode					
	(1)	(2)	(3)	(4)	(5)	(6)
<b>DID</b>	0.0352*** (0.0107)	0.0410*** (0.0107)	0.0417*** (0.0112)	0.0391*** (0.0115)	0.0422*** (0.0115)	0.0414*** (0.0111)
DA FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Controls				✓	✓	✓
SDID control pool		✓	✓		✓	✓
SDID weights			✓			✓
R <sup>2</sup>	0.83768	0.78668	0.78011	0.84196	0.79197	0.78571
Adjusted R <sup>2</sup>	0.78336	0.71526	0.70650	0.78851	0.72160	0.71322
Observations	101,358	38,769	38,769	100,244	38,348	38,348

*Notes:* Standard errors clustered at the CMA level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.1$

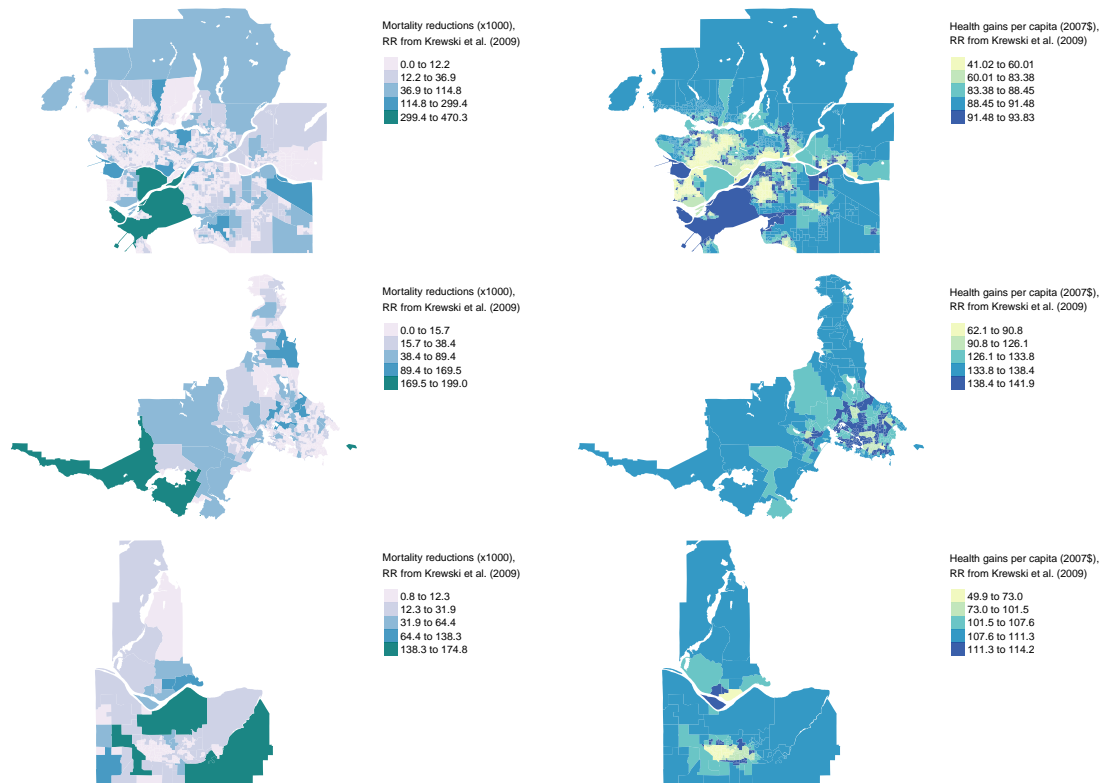
**Table C.2:** TWFE-DID results for zero emissions commute mode

	Zero Emission Commute Mode					
	(1)	(2)	(3)	(4)	(5)	(6)
<b>DID</b>	0.0057** (0.0025)	0.0106*** (0.0017)	0.0117*** (0.0021)	0.0066*** (0.0022)	0.0088*** (0.0016)	0.0092*** (0.0016)
DA FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Controls				✓	✓	✓
SDID control pool		✓	✓		✓	✓
SDID weights			✓			✓
R <sup>2</sup>	0.80811	0.80808	0.81877	0.81200	0.81355	0.82463
Adjusted R <sup>2</sup>	0.74390	0.74383	0.75810	0.74841	0.75047	0.76531
Observations	101,358	38,769	38,769	100,244	38,348	38,348

*Notes:* Standard errors clustered at the CMA level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.1$


## 2.D Additional Health Results

### 2.D.1 Estimates using RR from [Krewski et al. \(2009\)](#)



**Figure D.1:** Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from [Krewski et al. \(2009\)](#), for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAAs.

## 2.D.2 Health-income relationships



**Figure D.2:** Bivariate distribution of health gains using the RR from [Lepeule \*et al.\* \(2012\)](#) and median income for the Vancouver (top panel), Victoria (middle panel) and Abbotsford (bottom panel) CMAAs.

## Part II

# Causal Analyses of Tropical Deforestation

# Chapter 3

## The Unintended Impact of Colombia's Covid-19 Lockdown on Forest Fires

### Abstract

The covid-19 pandemic led to rapid and large-scale government intervention in economies and societies. A common policy response to covid-19 outbreaks has been the lockdown or quarantine. Designed to slow the spread of the disease, lockdowns have unintended consequences for the environment. This article examines the impact of Colombia's lockdown on forest fires, motivated by satellite data showing a particularly large upsurge of fires at around the time of lockdown implementation. We find that Colombia's lockdown is associated with an increase in forest fires compared to three different counterfactuals, constructed to simulate the expected number of fires in the absence of the lockdown. To varying degrees across Colombia's regions, the presence of armed groups is correlated with this fire upsurge. Mechanisms through which the lockdown might influence fire rates are discussed, including the mobilisation of armed groups and the reduction in the monitoring capacity of state and conservation organisations during the covid-19 outbreak. Given the fast-developing situation in Colombia, we conclude with some ideas for further research.

---

We thank the two reviewers at Environmental and Resource Economics for their helpful comments, Ian Bateman for his editorial guidance, and Leonardo Correa of the Fundación Paz y Reconciliación (PARES) for help in accessing data. Data used in the econometric analysis are available from the authors on request. Qualitative material used in this paper originated from the 'Bioresilience of Andean Forests in Colombia' research project, financed by NERC-UK.

*Note:* This chapter is available in print at [Environmental and Resource Economics](#).

### 3.1 Introduction

The global spread of Covid-19 in 2020 has had, and continues to have, a devastating impact on our societies and economies. In response, governments have intervened on a huge scale to try to slow and manage the spread of the disease, help those who get infected, and support economies. With the aim of slowing the spread of disease, mandatory ‘shelter-in-place’ restrictions on peoples’ movements, also known as lockdowns or quarantines, typically prevent people from leaving their homes or local areas for extended periods of time. By April 2020, lockdowns had become one of the commonest policy responses to Covid-19, affecting up to two-thirds of the global population (Bates *et al.*, 2020).

Evidence is emerging that suggests lockdowns have unintended environmental consequences, both negative and positive. For example, research undertaken in China suggests that lockdowns are associated with improvements in local air quality, likely due to sharp falls in road traffic and manufacturing activity (Liu *et al.*, 2020; Le *et al.*, 2020), while a lockdown-induced collapse in ecotourism revenues may have negatively affected local livelihoods, leading to an increase in wildlife poaching (The Guardian, 2020). Yet, the evidence base and hence, our understanding of how lockdowns might influence natural resource use, management and conservation, is still relatively weak. Our paper is an early, exploratory contribution, motivated by the release of satellite data showing a particularly large upsurge of forest fires in the Colombian Amazon that coincided with the emergence of Covid-19 in Latin America, in early-2020 (IDEAM, 2020; FCDS, 2020b).

We ask whether Colombia’s lockdown, which was implemented in stages between 14 March and 24 March 2020 and is projected to end on 15 July (at the time of writing in early-July), is associated with this observed upsurge of forest fires. As detailed in Section 3.2, forest governance and conservation in Colombia are intimately associated with the country’s long history of internal conflict and the militarisation of conflict areas. Thus, any analysis of forest change in Colombia needs to consider the role of its numerous armed groups. This we do by first exploiting spatial variation in the known locations of armed groups across Colombia in our econometric analysis, the methods for which are described in Section 3.3.

The observed number of forest fires are compared with three different counterfactuals (historical average, synthetic control, augmented synthetic control), constructed to simulate the expected number of fires in the absence of Colombia’s lockdown. Our results, presented in Section 3.4, suggest that the lockdown is associated with an increase in the number of fires. To varying degrees across Colombia’s regions, the presence of armed groups is found to be correlated with this increase. How and why Colombia’s lockdown might influence forest fires are questions that cannot be addressed by our econometric analysis. Therefore, in Section 3.5, we consider a number of possible mechanisms, including the mobilisation of Colombia’s armed groups and the monitoring capacity of state and conservation organisations, based on information from local stakeholders in lowland and Andean Colombia and key



informants connected with the Colombian Amazon, before and during the lockdown. The final section of the paper, [Section 3.6](#), concludes with some ideas for further work.

## 3.2 Background to Deforestation, Conflict and Covid-19 in Colombia

Colombia is one of the most biodiverse countries in the world, with five major biotic regions: Amazon, Andes, Caribbean, Orinoco and Pacific (see [Figure 3.A.1](#)). Between 1990 to 2016, more than six million hectares of natural forests were deforested ([IDEAM, 2018](#)). There was a surge in deforestation after a peace agreement with the guerrilla movement FARC-EP (The Revolutionary Armed Forces of Colombia — People's Army) was signed in 2016 (see e.g. [Clerici \*et al.\*, 2020](#)). The joint efforts of environmental organizations, state authorities, international cooperation, the media and communities helped reduce the annual deforestation rate nationally, by 10.1% in 2018, a trend that continued into 2019 ([IDEAM, 2018](#); [IDEAM, 2020](#)).

In common with many other parts of tropical Latin America, trees in Colombia are felled before any remaining forest is cleared by fire in preparation for new areas of crop cultivation and cattle pasture, although logging is not always followed by forest clearance via fire. This procedure takes advantage of seasonal climates, with felling often occurring in the wet season and fires subsequently started after weeks or months of less rain in the dry season. Most forest fires in the Colombian Amazon take place in the dry season, between November and April. When these dry season fires are started they can spread in forest areas where logging has not taken place previously. Forest fires are typically started by farmers or landless people seeking land for crop cultivation in order to feed their families and generate income, although other actors, such as armed groups, have also been implicated in forest fires.

Early estimates for 2020 suggest that rates of tree felling and deforestation in the Colombian Amazon are likely to reverse the gains of 2018 and 2019 ([MAAP, 2020](#); [SINCHI, 2020](#)), while Colombia's forest fire trends in the first half of 2020 imply that the country is on course to record one of its largest numbers of forest fires in recent years. As detailed in [Section 3.3](#), a huge increase in the number of forest fires was observed in March 2020 (12,953) compared to March 2019 (4,691) ([Semana Sostenible, 2020](#)). One explanation is that more rain than usual fell in December 2019, with Colombia's dry season starting later, in mid-January 2020 ([FCDS, 2020b](#)).

Forest governance and conservation in Colombia are associated with the militarisation of conflict areas. The current internal armed conflict dates back to a period called "the Violence", which lasted from 1948 until 1958 ([Guzmán Campos \*et al.\*, 1962](#)). After this period, bitterness at the lack of attention from the government in dealing with the conflict's underlying causes (land distribution and political exclusion), led to the rearming of guerrillas associated with the Liberal party and the conversion of some of these guerrillas into FARC-EP. Paramilitary groups were legally formed under the government's auspices in opposition to the guerrillas. The emergence of the

paramilitary groups was closely connected to the geographies of drug trafficking in Colombia (see e.g. Cubides, 1997; Gallego, 1990; Vargas, 1992). Areas under guerrilla control were branded “red zones” while the areas controlled by the paramilitaries were seen as zones of special control and military presence. People living in these areas came under the rule of these armed groups and often endured terrible hardship.

After the peace agreement with FARC-EP in 2016, the ELN (National Liberation Army) guerrillas, along with dissidents from FARC-EP, heavily-armed organized criminal organizations and neo-paramilitary groups associated with the political far right (see below<sup>1</sup>), expanded their territorial control in several forest areas, taking over areas that were previously controlled by FARC-EP. In some parts of Colombia, individual armed groups have secured full territorial control while in other parts different groups have been disputing control.

Until 2018, there was a rise in coca cultivation and deforestation rates in the Amazon, attributed to a lack of active government presence after the peace process (DeJusticia, 2018). The inability of the government in the post-conflict era to fill the power vacuum in areas previously occupied by FARC-EP, gave rise to the emergence of neo-paramilitary groups: the Urabeños, the Rastrojos and the Gulf Clan. Implicated in drug trafficking in association (or dispute) with Mexican drug cartels (PARES, 2018; Fundación Ideas para la Paz - FIP, 2019), these groups tended to operate on the Pacific coast and in some Inter-Andean areas.

Many areas contested by the armed groups were also priority areas for the government's military and conservation strategies. After 2016, the Colombian government promoted sustainable development and ecotourism, in an attempt to transform areas that were heavily affected by the armed conflict and previously under the full or partial control of FARC-EP. Under these government- and NGO-led schemes, some of which were implemented in National Natural Parks and areas of high biodiversity, former guerrillas were encouraged and trained for new roles, e.g. as forest guardians, tourist guides and organic farmers<sup>2</sup>.

The reduction in deforestation rates in 2018 and 2019 is partially attributed to Operation Artemisa, an initiative to curb deforestation in protected areas, led by the armed forces with the support of the Chief Prosecutor's Office, the Ministry of the Environment, and the Institute for Hydrology, Meteorology and Environmental Studies. Using satellite data, Geographical Information Systems, drones and field intelligence, some operations led to the arrest of actors, often local farmers and the landless, caught deforesting and starting fires illegally. Yet, the militarisation of the environmental agenda led to new rounds of conflict and protest. Peasant

---

<sup>1</sup>The “old” paramilitary groups were demobilized in 2004–2006 but some of the members of these groups joined forces with paramilitaries that chose not to demobilize to create smaller and more adaptable military organizations, known as neo-paramilitaries, BACRIM or Armed Organized Groups (GAOs).

<sup>2</sup>There are different types of protected area in Colombia: National Natural Parks, Regional Natural Parks, National Forest Reserves, Integrated Management Districts, Soil Protection Districts, Natural Reserves and Indigenous Territories.

organizations, environmental NGOs and human rights organizations have alleged that excessive force was used against local farmers and the landless, who were often treated as guerrilla collaborators. None of those financing and organizing the deforestation and forest fires, namely members of the armed groups, have been detained. Colombia's first case of Covid-19 was reported on 6 March 2020. The government's response to the outbreak that ensued was to implement a number of lockdown measures with the following timeline (Presidential Decree 749, 2020):

- 14 March: closure of border with Venezuela
- 15 March: suspension of all schools and universities
- 16 March: closure of all land and sea borders; curfews in several municipalities
- 17 March: declaration of state of emergency; mandatory isolation for all over-70s
- 20 March: announcement of nationwide quarantine, starting at midnight on 24 March

At the time of writing, Colombia's quarantine for all of the country's citizens was extended until 15 July, one of the world's most prolonged. Colombia's lockdown was effective in reducing mobility at its onset, soon after the closing of its land border with Venezuela (see [Figure 3.A.2](#) and [Figure 3.A.3](#)). As of 7 July 2020, Colombia reported 124,494 cases of Covid-19, of whom 4,359 have died.

## 3.3 Data and Methods

### 3.3.1 Data Description

Our econometric analysis in [Section 3.4](#) primarily utilises data for forest fires detected by the Visible Infrared Imaging Radiometer Suite (VIIRS) on board the Suomi NPP satellite. The spatial resolution of VIIRS is 375 metres resolution per pixel, a higher resolution than that of the Moderate Resolution Imaging Spectroradiometer (MODIS) sensor (1000m resolution per pixel). Thus, VIIRS can detect fires that MODIS might overlook, although fire data from the latter are used as an additional predictor in our application of the synthetic control method (SCM; see below). Daily data for forest fires detected by both VIIRS and MODIS are sourced from the National Aeronautics and Space Administration (NASA)<sup>3</sup>.

While MODIS data are available from 2001 onwards, VIIRS began operating in 2012; therefore, our data cover all of Colombia over the period between 1 January 2012 and 28 May 2020. We use the daily count of VIIRS and MODIS fire hotspots, a flow variable, and the daily cumulative sum of fire hotspots, a stock variable. Additional analysis makes use of the sum of daily Fire Radiative Power (FRP), also reported in the VIIRS and MODIS products, which accounts for heterogeneity in fire hotspots' size and intensity. Fire hotspots are aggregated at the country level (for the analyses that cover all of Colombia), and the municipality level (allowing us to examine regional heterogeneity). Given heterogeneity in the biophysical and ecological characteristics of the country, including climatic conditions and forest types, the municipalities are grouped by biotic region.

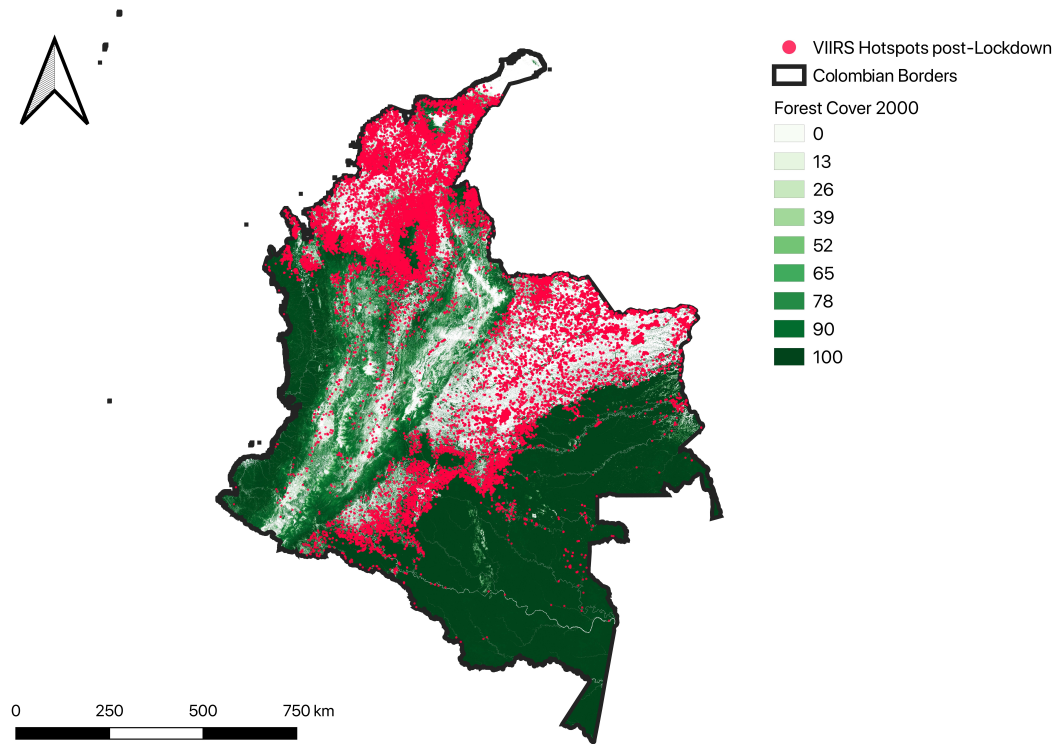
[Figure 3.3.1](#) shows the location of fire hotspots during Colombia's lockdown, up until 28 May 2020. Although fires can be observed across the country, in all biotic regions, they are particularly concentrated in northern, central and south-western areas. The northern and central zones are more seasonally dry, with large areas of dry forest and savanna that are naturally more fire-prone, but there are also many fires in the Andean valleys, and in and near the Amazon frontier. The lockdown restrictions mandated that the entire population of Colombia should stay home yet the patterns of fire hotspots in [Figure 3.3.1](#) suggest that people, at least in some parts of the country, were ignoring these restrictions.

Our 2012 to 2020 time-period, while motivated by data limitations, implies a reasonable level of confidence in ruling out major climatic shocks as the sole drivers of extreme fire seasons. We expect to observe some yearly fire variability in our sample, although as shown in [Figure 3.3.3](#) (daily count of fires) and [Figure 3.3.2](#) (cumulative daily number of fires), the dispersion of the time series is relatively contained. There is, however, an unusual spike in fire hotspots in both the count and cumulative trends starting from 14 March 2020, when the border with Venezuela was closed, the first step in Colombia's lockdown response to the Covid-19 pandemic.

---

<sup>3</sup>Available from [NASA MODIS](#) and [NASA Earthdata](#).


We are unable to augment the dataset with the inclusion of relevant climatic covariates (temperature, rainfall, wind speed, etc.) as controls due to the near-real time nature of our analysis.




**Figure 3.3.1:** Location of fire hotspots in Colombia, 14 March–28 May 2020. *Sources:* Hansen *et al.* (2013), NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/Radiometer Suite (VIIRS)

Thus, a major caveat of our analysis is the possibility that any effects we find are mediated by extreme climatic shocks during the lockdown period or by an idiosyncratic alteration in the timing of the start of the dry season. Visual inspection of the trends in cumulative hotspots (Figure 3.3.2) and fire counts (Figure 3.3.3) does not lend full support to these alternative explanations. Indeed, the patterns for the 2020 fire season indicate spikes similar to earlier years prior to mid-March and an unusual upsurge in hotspots during the Covid-19 lockdown. Also, 2020 was already the third-highest fire season in the record as of 14 March 2020, having deviated from the long-run historical mean since early-February, thereby reducing the likelihood that the observed increase in March was solely due to a late dry season.

Our discussion of how and why Colombia's Covid-19 lockdown might influence deforestation, presented in Section 3.5, is based on secondary sources, interview data and information obtained from research networks and key informants associated with the interdisciplinary research project “BioResilience: Biodiversity resilience and



**Figure 3.3.2:** Cumulative number of forest fires by year in Colombia, 2012–2020. *Note:* 95% Confidence interval around the historical mean shaded in grey. *Source:* NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/Radiometer Suite (VIIRS)



**Figure 3.3.3:** Daily number of forest fires by year in Colombia, 2012–2020. *Source:* NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group, 2020. Visible and Infrared Imager/ Radiometer Suite (VIIRS)

ecosystem services in post-conflict socio-ecological systems in Colombia”<sup>4</sup>. Since 2018, extensive socio-cultural fieldwork has been conducted in two areas representative of two different socio-ecological systems in the central-eastern Andean mountain range: the lowlands of the Middle Magdalena region (lowland rain forest merging with lower montane (cloud) forest at higher altitudes, above 1500m) and the highlands of the National Natural Park Chingaza.

### 3.3.2 Methodology

We first estimate “excess fires” over the lockdown period in Colombia covered in this paper, that is, from 14 March 2020 until 28 May 2020. Our measure of excess fires follows the methodology used to estimate excess mortality, a tool used in epidemiology to describe the number of deaths exceeding what would have been expected under “normal” conditions. Specifically, we use the methodology used to calculate excess deaths due to Covid-19 by The Economist and the Financial Times (for an overview, see [Roser \*et al.\* \(2020\)](#)), and adapt it to our purposes. Thus, the number of excess forest fires, EF, during the lockdown period is calculated by subtracting the mean number of fires, MF, during this same period between 2012 and 2019, i.e. 14 March–28 May, from the total number of fires observed, OF, between 14 March and 28 May, 2020:

$$EF_{14March-28May,2020} = OF_{14March-28May,2020} - MF_{14March-28May,2012-2019} \quad (3.1)$$

Therefore, our first empirical approach simply compares the observed number of fires against a historical average number of fires. As shown in [Figure 3.3.2](#) and [Figure 3.3.3](#), however, 2020 deviated from the mean historical fire trend well before the beginning of the lockdown, biasing this comparison upwards: cumulative fires on 14 March 2020 were already in excess of the 95% confidence interval around the long-run mean. Since this discrepancy could have been driven by idiosyncratic climatic factors specific to 2020, such a comparison is only useful for descriptive purposes, identifying this year’s fire season as anomalous with respect to prior ones. Thus, the historical mean of fire trends is not a statistically-grounded counterfactual for 2020 fire observations because it violates the foundational assumption of parallel trends prior to treatment.

To account for this problem, we apply two further approaches, akin to difference-in-differences, which compare the observed number of fires with a counterfactual constructed to simulate the expected number of fires in the absence of Colombia’s Covid-19 lockdown. Both methods have advantages over the use of the historical average in terms of how the underlying distribution of the historical fire data is treated and in accounting for pre-lockdown time trends in forest fires. Moreover, both methods are geared towards the construction of a counterfactual that closely tracks fire trends in 2020, thereby ensuring that fire trends for the treated and control units are optimally matched for the whole pre-treatment period, conditional on the

---

<sup>4</sup>[BioResilience Webpage](#).

feasibility of said matching.

Our second approach is the SCM. First developed by [Abadie and Gardeazabal \(2003\)](#), the SCM estimates an artificial counterfactual for the single treated unit via a data-driven method that employs minimal assumptions on its underlying data distribution. A synthetic control, our counterfactual, is a weighted average of the available control units. The weights for these units are estimated by minimizing the difference between the counterfactual and what was actually observed during the pre-treatment period.

Following [Modi \*et al.\* \(2020\)](#), we construct the synthetic control unit for Colombia's 2020 fire trends from a weighted average of fire trends in prior years. On the one hand, by using historical time periods in a single country rather than multiple countries, this procedure has the advantage of ruling out cross-country differences that cannot be summarised by the predictors of choice (e.g. different structures of the forestry sectors, differences in monitoring and enforcement, different timings of the wet and dry seasons). On the other hand, it is possible that events specific to 2020, perhaps not previously observed in the fire record, are entirely responsible for the deviation of the actual trend from its synthetic counterfactual.

Another concern arises from the inspection of the time series used as outcome variables. Indeed, as reported by [Masini and Medeiros \(2019\)](#), the SCM suffers from issues of over-rejection of the null hypothesis of no effect when the data are non-stationary, as is the case for the daily count of fire observations ([Figure 3.3.3](#)). For this reason, we use the cumulative count of fire observations, rather than daily fire counts, as our outcome variable.

To partially account for the issues connected to the data-generating process encountered with the SCM, we adopt the augmented SCM (ASCM, [Ben-Michael \*et al.\* \(2021\)](#)). Our third approach addresses a concern about the SCM, namely that it may not provide a meaningful estimate of excess forest fires if the trajectory of fires in the synthetic control unit does not closely match the trajectory of the lockdown treatment unit prior to the intervention ([Abadie \*et al.\*, 2015](#)). In particular, we follow [Ben-Michael \*et al.\* \(2021\)](#) and [Cole \*et al.\* \(2020\)](#) in implementing a ridge-regularised outcome regression model to estimate and correct for the bias arising from discrepancies in pre-intervention fit between the treated and synthetic units. Also, the ASCM is able to describe dispersion around its point estimate by employing the “average squared placebo gap”, which makes use of the standard leave-one-out SCM estimates in calculating the SCM noise variance ([Ben-Michael \*et al.\*, 2021](#)).

The synthetic and augmented synthetic controls are constructed by adapting the methodology from [Modi \*et al.\* \(2020\)](#) and employ 2012–2019 data as the “donor pool”. The number of pre-treatment days is  $T = 73$ . We employ MODIS fire observations, VIIRS Fire Radiative Power (FRP) and four lags of the dependent variable ( $n = 6$ ) as predictors. Both approaches transparently report the observations receiving non-zero weights in the construction of the artificial counterfactuals. For the SCM,



the 2020 trends are reconstructed from a weighted combination of trends recorded in 2018, 2013 and 2016<sup>5</sup>. The ASCM algorithm allows negative weights to be placed on donor observations thus making them slightly less interpretable. Nonetheless, 2018 and 2013 are again the most important fire seasons used in the construction of the counterfactual, followed by 2019 and 2016<sup>6</sup>.

We perform two robustness checks on our country-level results. First, an in-time placebo test (Abadie *et al.*, 2010; Abadie *et al.*, 2011; Abadie *et al.*, 2015) for both the SCM and ASCM, to check whether the optimisation algorithm provides significant results in the absence of an intervention. If a test of no intervention effects fails to reject the null hypothesis, then the method is poorly identified. We impose a placebo lockdown on 19 February that ends on 13 March, letting the matching procedure run up until 19 February. Second, we replicate the analysis using Fire Radiative Power (FRP) as an alternative dependent variable. By measuring the radiative intensity of fire hotspots, this outcome variable minimises the possibility that our results stem from more frequent, but less intense, fire observations.

To test for the number of forest fires conditional on the presence of armed groups in our regional analysis, data on the known locations of armed groups in Colombia are digitally coded into our dataset from maps originally created in 2019 by the Peace and Reconciliation Foundation (PARES). We focus on two of the main groups, with broad geographic reach in Colombia: FARC-EP dissidents and the Gulf Clan neo-paramilitaries (Figure 3.A.4).

Qualitative insights presented in Section 3.5 were generated from fieldwork involving ethnographies, participant observation, interviews and workshops. Dialogue with the inhabitants of the highlands as well as with those of the lowlands continued during Colombia's lockdown, via mobile phone and other electronic means. This dialogue, though not initially motivated by our research question, gives clear local perspectives on land-use change during the lockdown period, enabled by a high level of trust that has been built between local people, including community leaders, and members of the BioResilience research team. The BioResilience team is also embedded in research and civil society networks across the country, which have generated insights in other regions beyond the Andes, in particular, the Amazon. For security reasons, key informants and stakeholders are not cited in the text unless their views have already been made public, e.g. via NGO reports.

---

<sup>5</sup>This is reassuring because these three fire seasons have the largest numbers of fires in our dataset. Hence, 2018, 2013 and 2016 are best-positioned to reproduce the 2020 season, respectively, receiving weights of 0.508, 0.437 and 0.055. All the other years receive next-to-zero weights.

<sup>6</sup>2019 is the fourth-highest fire year in the record, thereby validating the performance of the ASCM.

### 3.4 Results

We provide exploratory, quantitative evidence for the unintended impact of Colombia's lockdown on forest fires, beginning with our results for all of Colombia. The actual time series of cumulative fire observations in 2020 is examined vis-à-vis the 2012–2019 historical mean, the synthetic control counterfactual and an augmented synthetic control estimated via a ridge regression in the pre-treatment period.


As noted in [Section 3.3](#), the historical mean does not represent an adequate counterfactual for 2020 fire trends. Indeed, even if the 2020 series is located in the 95% confidence interval around the mean up until February (see [Figure 3.3.2](#)), the trends diverge from January onwards, and exhibit fundamentally different slopes at the start of the treatment in mid-March. Thus, any comparison that uses the historical mean as a counterfactual substantially overestimates the upsurge in 2020 fires, as evidenced in [Figure 3.4.1](#), which shows the results of our country-level analysis.

The historical average is shown by the dotted line in [Figure 3.4.1](#) while the counterfactual fire trends generated by the SCM and ASCM are shown by the black and blue dashed lines, respectively. The lines generated by the SCM and ASCM clearly improve upon the simple historical mean. Indeed, these two fire trends are much more closely matched to actual fire observations (line shaded red) in 2020, up until 14 March, than the historical average. After 14 March, they diverge dramatically thus indicating evidence of a clear upsurge in cumulative daily fires. Note that the 95% confidence interval around the augmented synthetic control, shaded grey in [Figure 3.4.1](#), does not overlap with the 2020 actual series, thus identifying a statistically significant divergence of the 2020 fire season from weighted combinations of previous years' fire rates.


As of 28 May, the discrepancy between the actual 2020 fire season and its synthetic counterfactual totals 35,212 fires, a number that falls to 13,019–52,781 (the point estimate is 32,900) with respect to the augmented synthetic control. Notably, 80.4% of this difference (28,396 fires) is recorded within 1 month from the closure of the border with Venezuela (58.6–84.6% or 7,631–44,661 fires when employing the ASCM, with a point estimate of 78.7% or 25,894 fires), which indicates either that the lockdown created particular incentives to start forest fires and/or that specific climatic conditions have postponed the fire season to coincide precisely with the period of lockdown.

Our two robustness checks provide support for our country-level results. First, the in-time placebo test results, shown in [Figure 3.4.2](#), suggest that both the SCM and ASCM fail to identify an upsurge in fires coinciding with the placebo lockdown period, between 19 February and 13 March, thereby validating our SCM and ASCM procedures. Second, results from replicating our analysis using Fire Radiative Power (FRP) as an alternative dependent variable ([Figure 3.B.1](#)) are consistent with those in [Figure 3.4.1](#).

[Figure 3.4.2](#) suggests that we are almost certainly capturing the pre-lockdown trend



**Figure 3.4.1:** Actual series, historical mean, SCM and ASCM, for all of Colombia. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.



**Figure 3.4.2:** In-time Placebo Test for SCM and ASCM, with lockdown beginning on 19 February 2020. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.

correctly. Yet, we remain somewhat cautious about our results in [Figure 3.4.1](#) due to the non-stationarity caveats raised by [Masini and Medeiros \(2019\)](#) and the unavailability of climatic data. Thus, we cannot completely discount the possibility that our results, while robust and significant, could be plagued by oversized tests of no intervention effects or, unlikely as it may seem, mediated by extreme climatic events coinciding with the lockdown.

Although the cumulative number of fires has increased across the whole of Colombia, this fire surge is heterogeneous across regions and by presence or absence of armed groups. Results are generated by region with a focus on the two of the most biodiverse regions for which we also have qualitative insights: Amazon and Andes. Because their national-level command structures are not known with certainty, we consider the presence or absence of one or both of the FARC-EP dissidents and the Gulf Clan neo-paramilitaries at the regional scale. Note that due to the possibility of unobserved confounders, the following results should be interpreted as showing the extent to which the presence or absence of armed groups is correlated with forest fires.


The Amazon region, which is not naturally fire-prone, experienced an upsurge in forest fires during the process of lockdown ([Figure 3.4.3](#)). We observe militarised municipalities controlled by FARC-EP dissidents and municipalities where no known presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries is recorded in our dataset<sup>7</sup>. Municipalities controlled by FARC-EP dissidents exhibit significantly higher fire trends during the lockdown with respect to the synthetic (3,163 more fires as of 28 May) and augmented synthetic (3,078–3,493 more fires) controls ([Figure 3.4.3a](#)).

The contribution of the first month of lockdown to the fire upsurge in the Amazon is even starker than that for the whole country, accounting for 93.8% and 94.7% (the point estimate) with respect to the SCM and ASCM, respectively. Interestingly, and in contrast to the whole country, the fire upsurge begins to manifest after the government's announcement of a national quarantine on March 20, and levels off once the quarantine took effect, after 24 March. A significant upsurge in fires, even if more contained, is also observed in municipalities where neither FARC-EP dissidents nor Gulf Clan neo-paramilitaries are known to be present. From [Figure 3.4.3b](#), the lockdown resulted in 672 and 427–923 more fire hotspots with respect to the SCM and ASCM, respectively, again primarily in the first month of lockdown and especially after the announcement of the national quarantine.


In the Andes ([Figure 3.4.4](#)), both FARC-EP dissidents and the Gulf Clan are present in some but not all municipalities, and sometimes together in the same municipalities. Municipalities solely controlled by FARC-EP dissidents do not show significant increases in fires ([Figure 3.4.4a](#)). Here, the improved performance of the SCM and ASCM with respect to the simple historical mean is apparent and protects us against a false positive result. The ASCM, in particular, shields us against a

---

<sup>7</sup>Two municipalities are controlled by the Gulf Clan, as can be seen in Appendix [Figure 3.A.4](#), but their cumulative fire count never exceeds 60 between January and June 2020. Therefore, we exclude them from the dataset.




(a) Areas with FARC-EP dissidents.




(b) Areas with no known presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries.

**Figure 3.4.3:** Results for the Amazon region, conditional on armed group presence. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.


biased interpretation of the differences in trends between the actual series and the synthetic counterfactuals via the calculation of a 95% confidence interval. As shown in Figure 3.4.4b, the 95% confidence interval does not overlap with the actual trend during the first month of the lockdown, suggestive of a significant fire upsurge in municipalities controlled by the Gulf Clan. By 28 May, however, the observed 2020 rates are comparable with weighted combinations of previous years' fire seasons.




(a) Areas with FARC-EP dissidents.



(b) Areas with Gulf Clan neo-paramilitaries.



(c) Areas with FARC-EP dissidents and Gulf Clan neo-paramilitaries.



(d) Areas with no presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries.

**Figure 3.4.4:** Results for the Andes region, conditional on armed group presence. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.

Andean municipalities that register the compresence of FARC-EP dissidents and the Gulf Clan neo-paramilitaries experienced an upsurge in fires, which is both dramatic and significant (Figure 3.4.4c), albeit on a much smaller scale than the increases observed in the Amazon. Here, the null hypothesis of no lockdown effects is rejected in favour of the alternative of an increase accounting for 412–658 more fires with respect to the ASCM as of 28 May.

There are Andean municipalities that have no known presence of either the Gulf Clan or the FARC-EP dissidents, at least according to our dataset (Figure 3.4.4d). The ASCM improves significantly upon the SCM and guarantees a good pre-treatment fit. We identify a zero effect from the lockdown in driving up fire rates, with the actual series sitting within the 95% confidence interval around the ridge-augmented synthetic counterfactual.

The other regions (Caribbean, Orinoco, Pacific), the results of which are shown in Section 3.B, display heterogeneity in terms of forest fire trends. In the Caribbean

region, the presence of either the Gulf Clan alone (Figure 3.B.2a), or in combination with FARC-EP dissidents (Figure 3.B.2b), is correlated with an upsurge of forest fires during the lockdown, although the counterfactuals are subject to wide uncertainty when we consider the 95% confidence interval around the ASCM fire trends. Interestingly, the patterns observed in the Caribbean are, to some extent, also observed in the Orinoco (Figure 3.B.3b). The presence of FARC-EP dissidents alone is, similar to the Andes (Figure 3.4.4a), not associated with a fire upsurge (Figure 3.B.3a). Yet, municipalities where neither FARC-EP dissidents nor the Gulf Clan are known to be present (Figure 3.B.3d), appears to be associated with a huge fire upsurge (4,712 to 6,501 fires with respect to the ASCM as of 28 May). The Gulf Clan is omnipresent in the Pacific (Figure 3.B.4) and similar to the Caribbean and Orinoco regions, this particular armed group is associated with an upsurge of fires during the lockdown.

### 3.5 Discussion

In this section, we discuss how and why the Covid-19 outbreak in Colombia and its lockdown might relate to the particularly large and significant upsurge of fires observed in early-2020. We begin with a number of general, possible mechanisms before exploring those that might help explain spatial heterogeneity in our econometric results with respect to the Amazon and Andean regions, and the role of armed groups.

The Covid-19 outbreak is likely to have had the general effect of reducing individuals' incentives to travel and interact with other people, thus changing behaviour. A mandatory lockdown goes a step further by imposing a potential cost on individuals if moving around and interacting with other people breaks any laws punishable by, e.g. fines. In Colombia, people increasingly became less mobile after 14 March even though the national, mandated lockdown ("national quarantine") did not come into effect until 25 March. After 24 March, mobility actually began picking up again.

Further costs are imposed when lockdowns reduce jobs and incomes. If this reduces the ability or willingness to purchase commodities, such as timber and beef, at given prices, then any resultant fall in demand could help alleviate pressure on forests. As of June 2020, the effects of the Covid-19 pandemic on commodity markets appear to have been mixed (Mongabay, 2020). If lockdowns prevent farmers and other actors from clearing forests, e.g. for new cattle pastures, this may also reduce pressures on forests, at least in the short term.

Our econometric results suggest that the opposite happened in Colombia, indeed that the lockdown may have even increased incentives to start fires illegally. This could occur if there is a fall in the cost of getting caught due to weaker forest law enforcement, especially in forest areas where governance is already fragile. Reports of weakening law enforcement in Latin America's forests since the Covid-19 outbreak began (e.g. British Broadcasting Company (BBC), 2020; FCDS, 2020a; El País, 2020) suggests the possibility that the behaviour of Colombia's enforcement agencies might have changed in response to the outbreak and/or lockdown.

We found that Colombia's lockdown is associated with an upsurge in forest fires between 14 March and 24 March, that is, before the start of the national quarantine yet during the period when people were less mobile. Just after the announcement of the national quarantine on 20 March, both the Chief Prosecutor's Office and the regional environmental agencies strictly limited or stopped their officials' field visits, including to deforestation hotspots, not only to protect themselves but also to shield local communities. The armed forces and other agencies involved in Operation Artemisa partially suspended their operations against deforestation, although it is likely that this suspension occurred before the Covid-19 outbreak given that the last known, reported operation took place in October 2019. Also, the suspension was apparently motivated by criticism of Operation Artemisa from civil society and the media as well as a new focus on the forced eradication of coca in the Inter-Andean forests (see below).



Along with the possibility that the upsurge of forest fires occurred due to a slowdown of state enforcement, there are a number of alternative explanations. First, forest fires might have been started by farmers and landless people in anticipation of being locked down under the national quarantine. Second, the period after 14 March was already a time when the attention of the media (and government) was focused exclusively on the pandemic. Although plausible, we have little further information on these two explanations and hence, focus our discussion below on a third possible explanation, namely the mobilization of armed groups. In parts of the Colombian Amazon region, where the state already had a limited presence, armed groups and criminal networks, described by local environmental organisations as “land grabbers” and “mafias come bosque” (forest-devouring mafias) (Krause, 2019), took advantage of the lockdown to strengthen their territorial control and expand their activities, particularly forest-based ones (FCDS, 2020b). The underlying motivation of the armed groups and criminal networks was to profit from deforestation and generate income, e.g. from cattle ranching. Yet, our econometric analysis suggests that the upsurge in fires in the Amazon, which occurred at around the same time as the slowdown in state enforcement activities, was similar regardless of whether or not municipalities were under the control of FARC-EP dissidents. But as we lack data showing the location of municipalities controlled by criminal networks, we cannot evaluate whether, in municipalities where FARC-EP dissidents were absent, the fire upsurge might be correlated with any of these criminal networks.

That FARC-EP dissidents have strengthened their territorial control and influence in the absence of effective government control in the Amazon region, have expanded to new areas of the Amazon and have taken advantage of the lockdown to “burn the jungle”, is increasingly well-documented (Semana Sostenible, 2020). The National Natural Parks System temporarily closed all its parks on 16 March 2020 but in February, before the lockdown, FARC-EP dissidents forced the government to remove its rangers from a number of parks in the Amazon (FCDS, 2020a). Also, reports suggest that, pre-lockdown, FARC-EP dissidents collaborated with other actors to occupy and clear forest in La Macarena National Park, and local people in forest reserve areas were coerced by FARC-EP to cut or burn down large forest areas for the expansion of cattle ranching and coca production. Such activities are known to have continued during the lockdown.

The tightening of control by armed non-state actors in the Colombian Amazon has made it more difficult for civilian government agencies and NGOs to operate in the region. For the first time, FARC-EP dissidents have threatened the Amazon Vision programme – the Colombian government’s programme for sustainable development in the Amazon – as well as a number of other developmental and environmental actors, including NGOs and agencies affiliated with international development schemes.

Our econometric results for the Andes are heterogeneous. In some areas, the forests were unaffected, either by logging or fire, particularly in municipalities where neither the Gulf Clan nor the FARC-EP dissidents had control, for example, in the highland

Andean forests of Cundinamarca. In this part of the high Andes, the local authorities closed the inter-municipal roads during the lockdown thus interrupting the transport of cargo products, except for food, fuels and medicines. Affected commodities included timber, mined extractives such as sand, limestone, and coal, and other resources pertaining to the construction sector, which often involve dredging or grinding of the Andean mountains.

Municipalities in the Andes controlled by the Gulf Clan, particularly those controlled by both the Clan neo-paramilitaries and FARC-EP dissidents, experienced a sharp and significant increase in fires during the lockdown. Whether these two groups were in conflict or working with one another is unclear. Beyond these two groups, there are also other armed groups implicated in the starting of fires. For example, the municipality of Puerto Boyacá (Figure 3.B.5) has a long history of being militarised<sup>8</sup> and was partially abandoned by the government after 2016. Since 2016, organized criminal networks, along with the Gulf Clan neo-paramilitaries, have been in dispute over control of Puerto Boyacá. An important transit zone for drug trafficking in the lowland Andean forests, the national army entered Puerto Boyacá in early-April 2020 to carry out operations to eradicate coca crops and dismantle cocaine laboratories<sup>9</sup>. Forced eradication involved burning part of the surrounding forest, a process that has occurred in other coca growing areas in the Inter-Andean forests.

---

<sup>8</sup>Puerto Boyacá hosted a government-sponsored experiment, a “laboratory of military operations” during the 1980s, in an effort to deprive the FARC-EP of territorial control by using paramilitary groups. This model was exported to other parts of the county, and is known as the Puerto Boyacá paramilitary model. Throughout the 1990s, this model was expanded to other Colombian regions. New paramilitary blocs were established, which incriminated civilians, accusing them of being guerrilla collaborators, perpetuating massacres on this basis, and exercising control over the local economy.

<sup>9</sup>This occurred without consultation with, or participation of, local people or their organizations. Part of the 2016 peace agreement between the Colombian government and FARC-EP includes a right granted to local people to participate in voluntary programmes to eradicate illicit crops. However, during the lockdown the government preferred to focus its efforts on the forced eradication of illicit crops without the involvement of local people.

### 3.6 Conclusion

In the context of what is a rapidly developing situation in Colombia, our study is an early contribution regarding whether and how a common policy response to the Covid-19 outbreak, namely the lockdown, influences forest fire rates. Our econometric analysis clearly exposed the abnormality of the March–May 2020 fire season. Indeed, the maximum spike in fire observations in Colombia is usually observed around January–February, a trend from which 2020 does not diverge. Yet, as our results show, the 2020 fire season exceeds the intensity of the fire seasons in the previous 8 years by a significant and alarming amount. Although fires occur naturally in some of Colombia's ecosystems, such as those in the Caribbean and Orinoco regions, this upsurge generated additional carbon dioxide emissions and could be potentially detrimental to biodiversity. Moreover, damaged and degraded forests are associated with an increased risk of future virus spillovers (e.g. Olivero *et al.*, 2017; Rulli *et al.*, 2017).

The public policy implications of the fire upsurge will become clearer as time passes. That said, it is already clear that an increase in carbon dioxide emissions may make it harder for Colombia to meet its ambitions to reduce emissions from deforestation and forest degradation (REDD+). Other policy implications will emerge with further research. First, future work should be able to confirm whether our results hold with the inclusion of climatic data, as and when such data become available. Alternative explanations for the upsurge could also be evaluated, including the fact that the upsurge occurred while the country was preoccupied with the pandemic and the possibility that forest fires were started in anticipation of being locked down. The question then is who might have started these fires. We conjecture that it is likely to have been the same actors who were implicated in activities that involved the burning of forests prior to the Covid-19 outbreak. This includes poorer farmers and the landless—whether they had agency or not—but also larger landowners, criminal networks and armed groups.

From our regional analysis, territories controlled by the Gulf Clan neo-paramilitaries, either alone or with FARC-EP dissidents, experienced a significant increase in fire rates during the lockdown. This analysis could be improved with detailed data on the Clan, along with other armed groups, criminal networks, assorted mafias, and their activities. The involvement of some of these groups in the production of high-value commodities, like beef and cocaine, is likely to be central to incentivising forest clearing behaviour. Such incentives could be boosted by reduced government capacity to monitor and enforce the rules against illegal fire-setting. The Colombian government prior to the lockdown not only monitored fire hotspots and deforestation but also the armed groups themselves. Thus, if lockdown-induced reductions in the government's monitoring and enforcement capacity has played a role in increasing the incentives of armed groups to clear forest, it is imperative that such capacity is reinstated when the Covid-19 outbreak recedes.

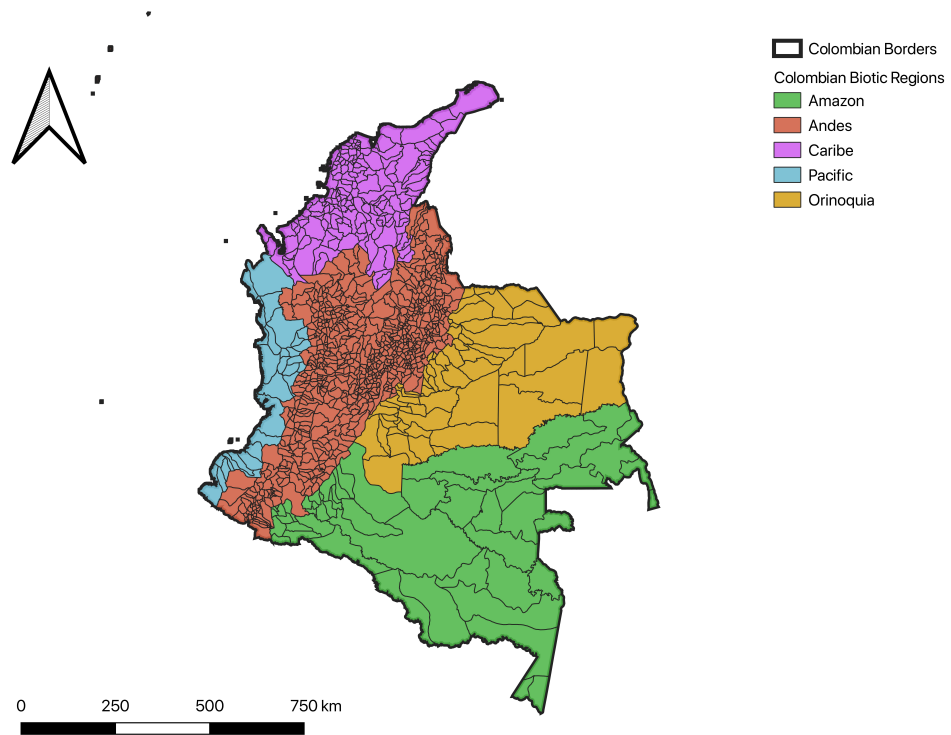
# References

- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2015). “Comparative Politics and the Synthetic Control Method”. *American Journal of Political Science* 59.2, pp. 495–510.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2011). “Synth: An R package for synthetic control methods in comparative case studies”. *Journal of Statistical Software* 42.13, pp. 1–17.
- Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2010). “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program”. *Journal of the American Statistical Association* 105.490, pp. 493–505.
- Abadie, Alberto and Gardeazabal, Javier (2003). “The Economic Costs of Conflict : A Case Study of the Basque Country”. *American Economic Review* 93.1, pp. 113–132.
- Bates, Amanda, Primack, Richard, Moraga, Paula, and Duarte, Carlos (June 2020). “COVID-19 pandemic and associated lockdown as a “Global Human Confinement Experiment” to investigate biodiversity conservation”. *Biological Conservation* 248, p. 108665.
- Ben-Michael, Eli, Feller, Avi, and Rothstein, Jesse (2021). “The Augmented Synthetic Control Method”. *Journal of the American Statistical Association* 116.536, pp. 1789–1803.
- British Broadcasting Company (BBC) (2020). *Brazil’s Amazon: Surge in deforestation as military prepares to deploy*. URL: <https://www.bbc.com/news/world-latin-america-52595030> (visited on 05/08/2020).
- Clerici, Nicola *et al.* (2020). “Deforestation in Colombian protected areas increased during post-conflict periods”. *Scientific Reports* 10, p. 4971.
- Cole, Matthew A., Elliott, Robert J. R., and Liu, Bowen (2020). *The Impact of the Wuhan Covid-19 Lockdown on Air Pollution and Health: A Machine Learning and Augmented Synthetic Control Approach*. Discussion Papers 20-09. Department of Economics, University of Birmingham.
- Cubides, F. (1997). “Los paramilitares y su estrategia, Ensayo”. *Documento de Trabajo N°8 por el Programa Paz Pública del CEDE de la Universidad de los Andes* 8.
- DeJusticia (2018). *Voces desde el Cocal. Colección Dejusticia, Bogotá: Editorial Dejusticia*. URL: <https://www.dejusticia.org/wp-content/uploads/2020/05/CATALOGO-2020.pdf> (visited on 03/20/2020).
- El País (2020). *La voraz deforestación durante la pandemia dispara las alarmas de los ambientalistas en Colombia*. URL: <https://elpais.com/sociedad/2020-05->

- [13/la-voraz-deforestacion-durante-la-pandemia-dispara-las-alarmas-de-los-ambientalistas-en-colombia.html](#) (visited on 05/15/2020).
- FCDS (2020a). *La voraz deforestación durante la pandemia dispara las alarmas de los ambientalistas en Colombia*. URL: <https://fcds.org.co/la-voraz-deforestacion-durante-la-pandemia-dispara-las-alarmas-de-los-ambientalistas-en-colombia/> (visited on 05/10/2020).
- FCDS (2020b). *Reporte Deforestación en la Amazonia Colombiana 2020*. URL: <https://fcds.org.co/publicaciones/reportes-fcde-deforestacion-amazonia-colombiana-2020/> (visited on 04/06/2020).
- Fundación Ideas para la Paz - FIP (2019). *Dinámicas de la confrontación armada y afectación humanitaria*. URL: <http://ideaspaz.org/especiales/infografias/homicidios-ene-sep.html>. (visited on 06/08/2020).
- Gallego, M.C. (1990). “Autodefensas, Paramilitares y Narcotraficantes en Colombia: origen, Desarrollo y Consolidación, el caso de Puerto Boyacá”. *Documentos Periodísticos, Bogota, Colombia*.
- Guzmán Campos, Germán, Fals Borda, Orlando, and Umaña Luna, Eduardo (1962). “La violencia en Colombia”. *Estudio de un proceso social* 1.
- Hansen, M. C. *et al.* (2013). “High-Resolution Global Maps of 21st-Century Forest Cover Change”. *Science* 342.6160, pp. 850–853.
- IDEAM (2018). *Obtenido de Resultados Monitoreo de Deforestación 2017*. URL: [http://documentacion.ideam.gov.co/openbiblio/bvirtual/023835/Resultados\\_Monitoreo\\_Deforestacion\\_2017.pdf](http://documentacion.ideam.gov.co/openbiblio/bvirtual/023835/Resultados_Monitoreo_Deforestacion_2017.pdf) (visited on 03/21/2020).
- IDEAM (2020). *Sistema de monitoreo de bosques y carbono SMBYC*. URL: <http://smbyc.ideam.gov.co/MonitoreoBC-WEB/pub/alertasDeforestacion.jsp?0.7620455894427884> (visited on 06/07/2020).
- Krause, Torsten (2019). “Forest governance in post-agreement Colombia”. *Natural resource conflicts and sustainable development*. Routledge, pp. 114–127.
- Le, Tianhao, Wang, Yuan, Liu, Lang, Yang, Jiani, Yung, Yuk L., Li, Guohui, and Seinfeld, John H. (2020). “Unexpected air pollution with marked emission reductions during the COVID-19 outbreak in China”. *Science* 369.6504, pp. 702–706.
- Liu, Fei *et al.* (2020). “Abrupt decline in tropospheric nitrogen dioxide over China after the outbreak of COVID-19”. *Science Advances* 6.28, eabc2992.
- MAAP (2020). *Deforestación de la Amazonía Colombiana 2020*. URL: <https://maaproject.org/> (visited on 06/15/2020).
- Masini, Ricardo and Medeiros, Marcelo C. (2019). “Counterfactual Analysis With Artificial Controls: Inference, High Dimensions, and Nonstationarity”. *Journal of the American Statistical Association* 116.536, pp. 1773–1788.
- Modi, C., Böhm, V., Ferraro, S., Stein, G., and Seljak, U. (2020). *How deadly is COVID-19? A rigorous analysis of excess mortality and age-dependent fatality rates in Italy*. URL: <https://www.medrxiv.org/content/10.1101/2020.04.15.20067074v3.full.pdf> (visited on 05/26/2020).
- Mongabay (2020). *As COVID-19 spreads, commodity markets rumble*. URL: <https://news.mongabay.com/2020/04/as-covid-19-spreads-commodity-markets-rumble>. (visited on 06/03/2020).
- Olivero, Jesús *et al.* (2017). “Recent loss of closed forests is associated with Ebola virus disease outbreaks OPEN”. *Scientific Reports* 7, pp. 1–9.

- PARES (2018). *Zonas Post-FARC*. URL: <https://pares.com.co/2018/06/25/zonas-postfarc/>. (visited on 06/08/2020).
- Roser, M., Ritchie, H., Ortiz-Ospina, E., and Hasell, J. (2020). *Excess mortality from the Coronavirus pandemic (COVID-19)*. URL: <https://ourworldindata.org/excess-mortality-covid> (visited on 06/09/2020).
- Rulli, Maria Cristina, Santini, Monia, Hayman, David TS, and D’Odorico, Paolo (2017). “The nexus between forest fragmentation in Africa and Ebola virus disease outbreaks”. *Scientific reports* 7.1, pp. 1–8.
- SINCHI (2020). *Reporte de Puntos de Calor*. URL: <http://siatac.co/web/guest/reportes> (visited on 03/27/2020).
- Semana Sostenible (2020). *Están aprovechando la cuarentena para quemar la selva*. URL: <https://sostenibilidad.semmana.com/impacto/articulo/estan-aprovechando-la-cuarentena-para-quemar-la-selva-corporamazonia/49489> (visited on 05/24/2020).
- The Guardian (2020). *Conservation in crisis: ecotourism collapse threatens communities and wildlife*. URL: <https://www.theguardian.com/environment/2020/may/05/conservation-in-crisis-covid-19-coronavirus-ecotourism-collapse-threatens-communities-and-wildlife-aoe> (visited on 05/07/2020).
- Vargas, V.A. (1992). “Magdalena Medio Santandereano, Colonización y conflicto armado.” *CINEP, Colombia*.


### 3.A Additional Figures



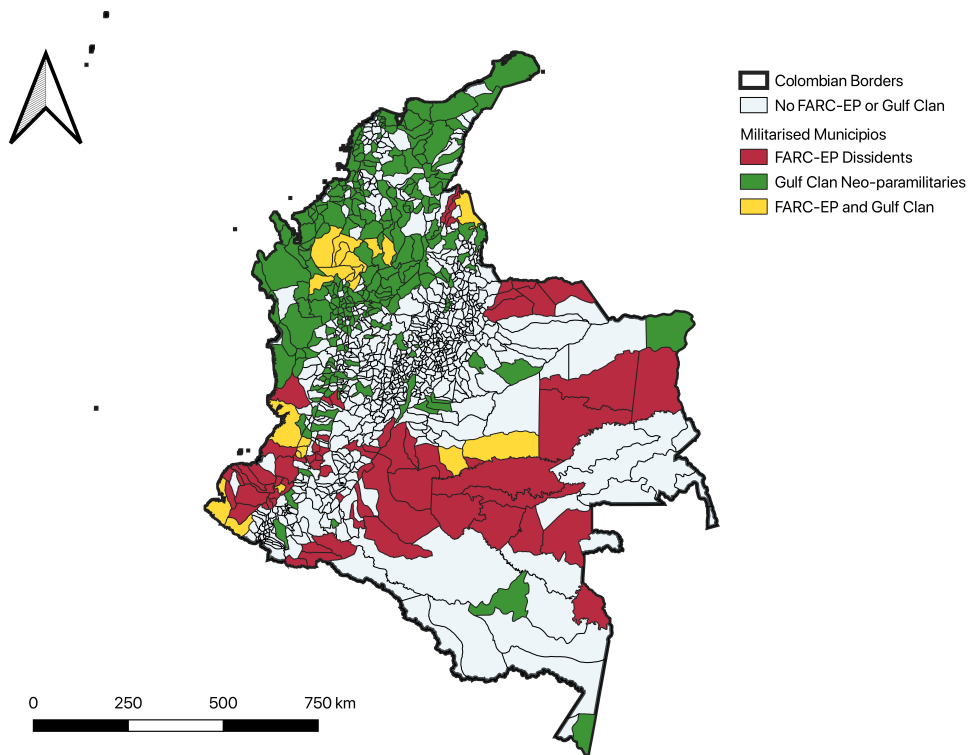
**Figure 3.A.1:** Colombia's biotic regions. Areas within regions denote the boundaries of Colombian municipalities.



**Figure 3.A.2:** Mobility by destination in Colombia, 2020.




**Figure 3.A.3:** Mobility, walking or driving, Colombia, 2020.




**Figure 3.A.4:** Known locations of FARC-EP dissidents and the Gulf Clan in Colombia by municipality, 2019.




### 3.B Additional Results




**Figure 3.B.1:** Alternative dependent variable, Fire Radiative Power (FRP). *Note:* 95% Confidence Interval around the ASCM series shaded in grey.



(a) Areas with Gulf Clan neo-paramilitaries.

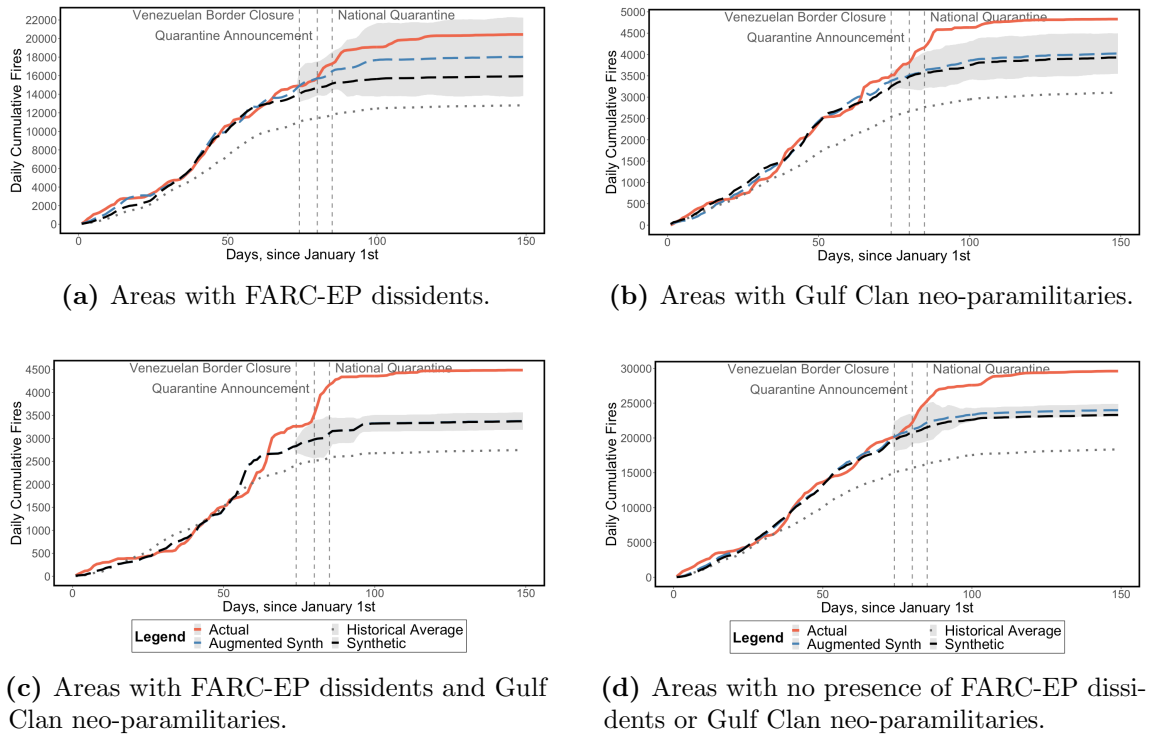


(b) Areas with FARC-EP dissidents and Gulf Clan neo-paramilitaries.




(c) Areas with no presence of FARC-EP dissidents or Gulf Clan neo-paramilitaries.

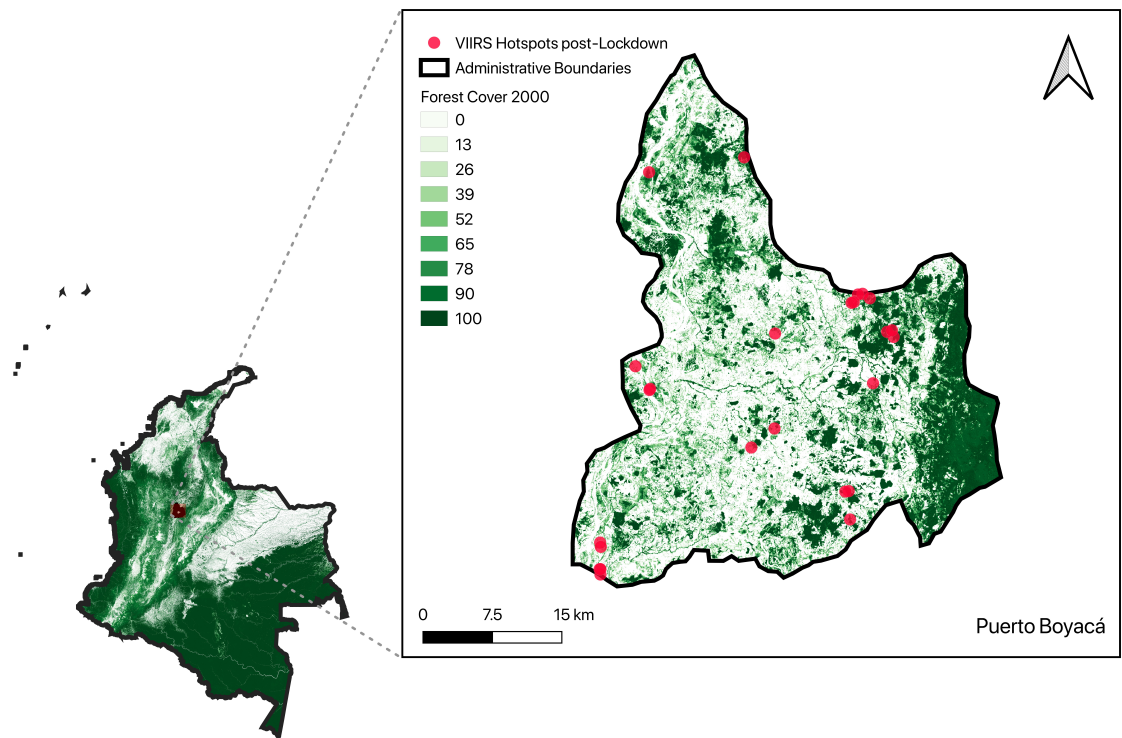
**Figure 3.B.2:** Results for the Caribbean region, conditional on armed group presence. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.



**Figure 3.B.3:** Results for the Orinoco region, conditional on armed group presence. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.



**Figure 3.B.4:** Results for the Pacific region, conditional on armed group presence. Only areas with Gulf Clan neo-paramilitaries. *Note:* 95% Confidence Interval around the ASCM series shaded in grey.



**Figure 3.B.5:** Location of fire hotspots in Puerto Boyacá, 14 March–28 May 2020.

# Chapter 4

## Carbon emissions reductions from Indonesia's Moratorium on forest concessions are cost-effective yet contribute little to Paris pledges

### Abstract

International initiatives for reducing carbon emissions from deforestation and forest degradation (REDD+) could make critical, cost-effective contributions to tropical countries' Nationally Determined Contributions. Norway, a key donor of such initiatives, had a REDD+ partnership with Indonesia, offering results-based payments in exchange for emissions reductions calculated against a historical baseline. Central to this partnership was an area-based Moratorium on new oil palm, timber and logging concessions in primary and peatland forests. We evaluate the effectiveness of the Moratorium between 2011 and 2018 by applying a matched triple difference strategy to a unique panel dataset. Treated dryland forest inside Moratorium areas retained at most, an average of 0.65% higher forest cover compared to untreated dryland forest outside the Moratorium. By contrast, carbon-rich peatland forest was unaffected by the Moratorium. Cumulative avoided dryland deforestation from 2011 until 2018 translates into 67.8-86.9 million tons of emissions reductions, implying an effective carbon price below Norway's US\$5 per ton price. Based on Norway's price, our estimated cumulative emissions reductions are equivalent to a payment of US\$339-434.5 million. Annually, our estimates suggest a 3-4 percent contribution to Indonesia's NDC commitment of a 29% emissions reduction by 2030. Despite the Indonesia-Norway partnership ending in 2021, reducing emissions from deforestation remains critical for meeting this commitment. Future area-based REDD+ initiatives could build on the Moratorium's outcomes by reforming its incentives and institutional arrangements, particularly in peatland forest areas.

---

For helpful discussions and comments we thank: Eugenie Dugoua, Paul Ferraro, Sefi Roth and Tom Smith; participants at conferences hosted by LEEP (Exeter) and BIOECON (Cambridge); research seminars hosted by the Department of Geography and Environment and the Saw Swee Hock Southeast Asia Centre (SEAC) at the LSE; the LEEPout, SWEEET, SWEEP and IBAHCM webinars hosted, respectively, by the Land, Environment, Economics and Policy Institute, the Institute for Research in Management and Economics, the RFF-CMCC European Institute on Economics and the Environment, and Glasgow University. We also thank Ellen Bruzelius Backer (Government of Norway's International Climate and Forest Initiative), along with Celine Lofthus Gaasrud and Marianne Johansen (Royal Norwegian Embassy in Jakarta), for further insights, and SEAC for funding. Finally, we thank the reviewers. The usual disclaimer applies.

*Note:* This chapter is available in print at [PNAS](#).

## Significance Statement

More than a decade after the global adoption of REDD+ as a climate change mitigation strategy, countries have started accessing results-based payments. However, the extent to which payments are actually based on results is unknown, necessitating programme evaluations to establish the contribution of REDD+ to the Paris NDCs. We undertake a micro-econometric evaluation of one of the most globally-significant REDD+ initiatives, Indonesia's Moratorium on forest concessions, in which a payment has been awarded. At the agreed US\$5/tCO<sub>2</sub>-eq, the value of our estimated cumulative carbon emissions far exceeds the proposed payment from the donor, Norway. Although cost-effective, the emissions reductions only contribute 3-4% of Indonesia's NDC. This contribution could be increased in new initiatives with better-designed incentives and institutional arrangements.

### 4.1 Introduction

Deforestation and forest degradation account for approximately 10% of global greenhouse gas emissions (IPCC, 2014). Recognising the importance of slowing deforestation in efforts to mitigate global warming, an international framework for Reducing Emissions from Deforestation and Forest Degradation (REDD+) was established in 2007 at the 13th Conference of the Parties (COP) (UNFCCC, 2008). At this COP, Norway's government announced its International Climate and Forest Initiative, pledging up to \$300 million every year towards REDD+. Norway's funds have been channelled through a number of negotiated, bilateral deals with countries hosting tropical forests. Countries include Brazil, Guyana, Tanzania and the setting for our study, Indonesia, where between 2000 and 2010 lowland evergreen forests and peat swamp forests were deforested by 1.2% and 2.2% per year, respectively (Miettinen *et al.*, 2011).

Indonesia is one of the world's largest GHG emitters. Between 2000 and 2016, approximately 50% of Indonesia's annual emissions were generated from deforestation, forest degradation, peatland decomposition and peat fires, accounting for around a quarter of global emissions from these sources (Government of Indonesia, 2018; FAO, 2021). The country's partnership with Norway, established in 2010, included a pledge of US\$1 billion to fund "results-based" REDD+ payments (Government of Norway, 2010). Central to this partnership was a Moratorium on the granting of new concession licenses by district governments for the conversion of primary dryland and peatland forest into new palm oil, timber and logging concessions (Government of Indonesia, 2011; Purnomo, 2012). Such concessions, operating across the archipelago (Figure 4.A.1), have been estimated to be responsible for almost half of Indonesia's forest loss (Abood *et al.*, 2015; Austin *et al.*, 2019). Implemented in 2011, the Moratorium initially covered 69 million ha of forest across the country (LTS International, 2018), most of Indonesia's forest estate (*SI Appendix*, Fig. S2). Additional restrictions on the conversion of peatland forest, affecting all concession types, were implemented across Indonesia in 2017 (Alisjahbana and Busch, 2017), and in late-2018, a new three-year Moratorium on new palm oil concessions was also imposed nationally (Mongabay, 2019a).

In 2017, Indonesia reportedly reduced emissions from deforestation and forest degradation by 11.2 MtCO<sub>2</sub>-eq (Mongabay, 2019b). Norway subsequently announced that it would pay Indonesia US\$56.2 million (Mongabay, 2020b) based on a carbon price of US\$5/tCO<sub>2</sub>-eq. In this article, we ask whether Norway is getting carbon value for its money, that is, whether this payment is actually based on results. The extent to which the Moratorium has had any meaningful impact on deforestation has been the subject of intense debate in Indonesia, particularly after it became permanent in 2019 (Mongabay, 2019b). Although the REDD+ partnership was terminated by Indonesia's government in 2021 (Mongabay, 2021b), an effective Moratorium since 2011 could contribute to meeting Indonesia's Nationally Determined Contribution (NDC) commitment of reducing GHG emissions by 29% unconditionally (and up to 41% conditionally) by 2030 (Government of Indonesia, 2018). Indeed, large-scale REDD+ initiatives play a potentially critical role in global climate change mitigation efforts (Roopsind *et al.*, 2019; Duchelle *et al.*, 2018) and more than a decade after the 13th COP, Indonesia is among a number of countries that has begun moving towards REDD+ implementation and accessing results-based payments (Maniatis *et al.*, 2019).

Norway's US\$1 billion pledge to Indonesia emphasises the global role of Norway in the design and funding of international REDD+ strategies (Angelsen, 2017). This pledge acted as an incentive to Indonesia's national government to enforce the Moratorium. A measurement, reporting and verification (MRV) system was developed and although district governments had an enforcement role (Mongabay, 2019a), there was little evidence of coordination, or of a plan to share benefits, between the national and district governments (Austin *et al.*, 2014; LTS International, 2018). Effective coordination might have prevented, or at least influenced, changes in the Moratorium's boundaries due to the re-designation of forestland by district governments (Enrici and Hubacek, 2016; Mongabay, 2019a). That re-designated forestland was often subsequently licensed out to concessionaires has raised concerns about corruption among government officials (Mongabay, 2019a). Long endemic in Indonesia's forest sector (Enrici and Hubacek, 2016; Palmer, 2005; Boer, 2020), corruption exacerbates the country's weak capacity to monitor and enforce forest regulations, characterised by, e.g. limited budgets and personnel (Meehan and Tacconi, 2017; Tacconi *et al.*, 2019). In sum, we anticipate little or no impact of the Moratorium on deforestation. A best-case scenario from previous research, an *ex ante* simulation of the Moratorium as the counterfactual to actual land uses between 2000 and 2010 (prior to the implementation of the Moratorium) and assuming 100% compliance (or effectiveness) (Busch *et al.*, 2015), indicates a maximum 3.5% reduction in deforestation and a 7.2% reduction in emissions.

While this best-case scenario was based on econometric analysis, the payment for the emissions reduction in 2017 was estimated by comparing, for the whole country, the amount of deforestation observed in 2017 against a historical baseline based on the average annual level of deforestation observed between 2006-16 (MoEF, 2019). The use of a historical baseline as a counterfactual provides weak evidence that the Moratorium has had a causal effect on REDD+ outcomes because deforestation in any given year will vary due to stochastic natural processes (e.g. weather and fires)

and economic factors (e.g. demand for commodities). Thus, it is highly unlikely that observed deviations from the average deforestation rate can be meaningfully related to the performance of the Moratorium and instead such deviations could be over- or under-estimated by chance.

Our analysis begins with the observation that any measurable policy effect could only have been generated by forest areas covered by the Moratorium. Estimating a policy effect requires a comparison of forest areas covered by the Moratorium with a counterfactual that mimics what would have happened in those areas had the Moratorium not been implemented. The intention of the counterfactual is to ensure that all other (non-Moratorium) factors relevant for determining the economic viability of palm oil and timber production in forest areas (Angelsen and Kaimowitz, 1999; Geist and Lambin, 2002; Busch and Ferretti-Gallon, 2017) are the same. Similar to networks of protected areas, Moratorium areas were not randomly assigned (Andam *et al.*, 2008; Joppa and Pfaff, 2010) and the small size of the reductions simulated by Busch *et al.* (2015) imply relatively low returns from forest conversion in these areas. Estimating the impact of the Moratorium on deforestation is therefore hampered by pre-existing differences in levels and the likely trajectory of deforestation between Moratorium and non-Moratorium areas. Furthermore, district governments could continue to issue licenses for new concessions in forestland outside Moratorium areas. One response to the Moratorium might be for licenses planned for Moratorium areas to be granted in non-Moratorium areas instead. Such spatial spillovers (“leakage”), a common concern in forest conservation (e.g. Andam *et al.*, 2008; Joppa and Pfaff, 2010; Meyfroidt and Lambin, 2009; Meyfroidt *et al.*, 2010; Ostwald and Henders, 2014), also confound estimates of Moratorium impact, potentially making it look more successful compared to non-Moratorium areas, when in fact activities have just been displaced. By testing for leakage, much can be learned about the processes that govern successful – or poor – performances. None of these confounding effects is specifically accounted for in the estimates of the payment offered by Norway’s government (Section 4.A).

To address confounding factors and isolate the Moratorium’s impacts on deforestation and associated emissions, we undertake a programme evaluation using quasi-experimental methods: a matched triple difference strategy (Chabé-Ferret and Subervie, 2013) applied to Global Forest Change data (Hansen *et al.*, 2013) at the 1.2km-by-1.2km scale between 2004 and 2018 (Section 4.4). A matched difference-in-differences estimator controls for the different levels of forest cover in Moratorium and non-Moratorium areas, and removes the deforestation trend in non-Moratorium areas matched by 1.2km<sup>2</sup> grid cells. Matching on, e.g. proximity to markets and topography, means that Moratorium and non-Moratorium grid cells have similar probabilities of concession-driven forest loss. The triple difference step removes any remaining deviations in deforestation trends prior to the Moratorium commencing in 2011, which otherwise would be attributed to the Moratorium (Chabé-Ferret, 2015; Chabé-Ferret, 2017). The impact that remains once these confounding factors have been addressed can then be attributed to the Moratorium. In principle, only estimates that have taken seriously the non-random assignment of the Moratorium



should inform the Norwegian government’s results-based payments for emissions reductions. We test the robustness of our estimates in several ways (Section 4.4). In particular, the potential for leakage as a response to the Moratorium is tested with a regression discontinuity analysis applied to the boundaries of Moratorium areas each year between 2005-2018. Thus, we estimate the differences in deforestation rates between each side of the Moratorium’s boundaries. By comparing these differences from before to after the start of the Moratorium in 2011, our analysis provides suggestive evidence for or against the presence of spillovers.


After evaluating the robustness of model estimates, we convert the Moratorium’s impacts on forest cover change into carbon dioxide equivalents for comparison with estimates generated by the Indonesia-Norway partnership. On the basis of our programme evaluation approach, we conclude that Norway’s government would be getting carbon value for money but the emissions reductions generated by the Moratorium contribute relatively little to Indonesia’s NDC commitments.

## 4.2 Results

### 4.2.1 Forest cover trends in Indonesia

Our outcome variable is “forest cover”, either dryland or peatland, in hectares. Concessions established within the Moratorium’s 2011 boundaries prior to the start of the Moratorium were legally allowed to continue operating, business-as-usual, after 2011. Thus, in Figure 4.2.1 we distinguish between forest cover trends observed in forest areas located outside (panels A and B) and inside (C and D) concessions. Overall, the proportion of forest cover has declined, both inside and outside the Moratorium’s boundaries, by approximately 10-15 percentage points between 2000 and 2018. By comparing trends inside the Moratorium to those outside, we observe that the rate of decline differs.

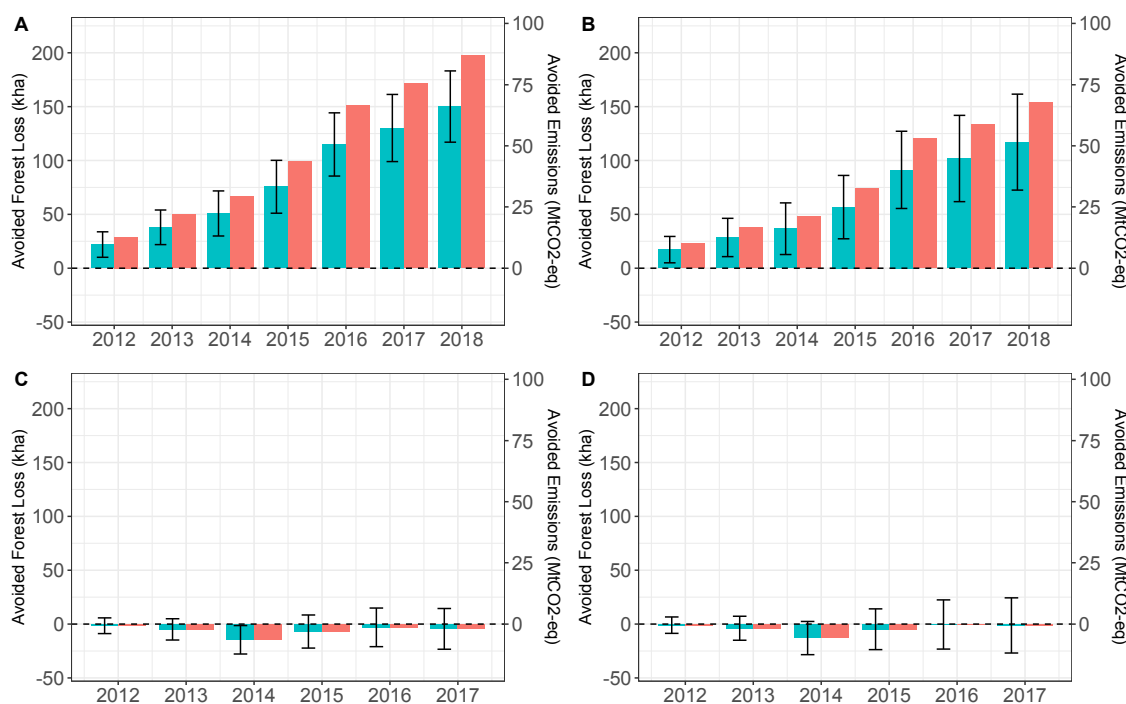
The extent of forest cover outside concessions is, on average, higher inside the Moratorium compared to outside, and a steeper decline in forest cover is observed outside the Moratorium compared to inside (Figure 4.2.1, A and B). These trends suggest that the Moratorium’s impacts from 2011 onward can only realistically stem from differences in the negative trends in forest cover between Moratorium and non-Moratorium areas. A similar pattern of decline is also observed in the trends for concessions. The extent of dryland forest cover is almost the same when we compare concessions inside the Moratorium to those outside (Figure 4.2.1, C). By contrast, the extent of peatland forest cover is higher in concessions outside the Moratorium compared to concessions inside (Fig. Figure 4.2.1, D). Unaffected by the Moratorium in principle, we use concessions as a placebo (falsification) test in our empirical analysis (Section 4.4).



**Figure 4.2.1:** Forest cover trends inside and outside the Moratorium, 2000-2018: non-concession dryland grid cells (A); non-concession peatland grid cells (B); concession dryland grid cells (C); concession peatland grid cells (D). Shaded areas denote treatment period. Grid cells in (A) and (B) also exclude forest in protected areas.

### 4.2.2 National-level effects of the Moratorium

Cumulative avoided forest loss and carbon emissions are estimated separately for dryland forest (Figure 4.2.2, A and B) and peatland forest (Figure 4.2.2, C and D), based on estimates of the average treatment effect on the treated (ATT). In the spirit of Andam *et al.* (2008), two estimators are presented to show the range of plausible results: the non-parametric difference-in-difference (“DD”: upper-bound estimate) and triple difference approaches (“DDD”: lower-bound estimate). We prefer the DDD estimator because it has a correction for non-parallel trends implied by the partial failure of the DD estimator of our placebo tests (placebo treatments prior to 2011), suggesting non-parallel trends even after matching. The differences in estimates shown in Figure 4.2.2 demonstrate the importance of examining the parallel trends assumption. Further placebo tests using concessions provide further support for the DDD approach (Section 4.4).



**Figure 4.2.2:** Cumulative avoided forest loss (’000 ha) and avoided carbon dioxide emissions (MtCO<sub>2</sub>-eq): dryland forest DD, 2012-2018 (A); dryland forest DDD, 2012-2018 (B); peatland forest DD, 2012-2017 (C); peatland DDD, 2012-2017 (D). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI.

Our ATT estimates are equivalent to the amount of forest loss, in hectares, avoided in each grid cell of 144 ha. The ATT for dryland forest (Table 4.B.1, Table 4.B.2) range between (with  $p \leq 0.000$  where it is not reported): 0.108 ( $p=0.006$ ) and 0.137 in 2011-12; 0.178 ( $p=0.002$ ) and 0.237 in 2011-13; 0.229 ( $p=0.003$ ) and 0.318 in 2011-14;

0.354 and 0.472 in 2011-15; 0.571 and 0.718 in 2011-16; 0.637 and 0.813 in 2011-17; and, 0.732 and 0.938 in 2011-18. Despite the magnitude of impact increasing steadily since 2011, the Moratorium was effective in protecting no more than one hectare of dryland forest in each grid cell by 2018, or about 0.651% of a cell. Our ATT estimates for peatland forest (Table 4.B.3, Table 4.B.4) are close to zero and not statistically significant at conventional levels.


The Moratorium has had a relatively small cumulative impact in preventing deforestation in dryland forest covered by the Moratorium relative to comparable forest areas outside the Moratorium (Figure 4.2.2, A and B; Table 4.B.5, Table 4.B.6): 17,248-21,967 ha in 2011-12; 28,533-37,972 ha in 2011-13; 36,672-50,830 ha in 2011-14; 56,725-75,603 ha in 2011-15; 91,303-114,901 ha in 2011-16; 101,851-130,168 ha in 2011-17; and, 117,053-150,089 ha in 2011-18. Put in context, avoided dryland forest loss represents, at most, 0.03% of all land covered by the Moratorium in 2012, rising over seven-fold to 0.22% by 2018. Our estimates of avoided dryland forest loss translate into cumulative carbon emission reductions of: 10.0-12.7 MtCO<sub>2</sub>-eq in 2011-12; 16.5-22.0 MtCO<sub>2</sub>-eq in 2011-13; 21.2-29.4 MtCO<sub>2</sub>-eq in 2011-14; 32.9-43.8 MtCO<sub>2</sub>-eq in 2011-15; 52.9-66.6 MtCO<sub>2</sub>-eq in 2011-16; 59.0-75.4 MtCO<sub>2</sub>-eq in 2011-17; and, 67.8-86.9 MtCO<sub>2</sub>-eq in 2011-18 (*SI Appendix*, Table S5-S6). By contrast, the Moratorium had null effects on peatland forest (Figure 4.2.2, C and D; Table 4.B.7, Table 4.B.8), implying a high likelihood of few if any carbon emissions savings, including those from peat drainage and peat fires.

Our results in Figure 4.2.2 are checked for their robustness (Section 4.4). First, we trim the sample on the basis of forest cover extent at the grid-cell scale (DD and DDD; Table 4.B.9-Table 4.B.12). Second, we apply Coarsened Exact Matching models (DDD only; Table 4.B.13, Table 4.B.14). Third, we apply a wider caliper (0.001) (DDD only; Table 4.B.15, Table 4.B.16) and 1:2, 1:3 and 1:5 nearest-neighbour matching (DDD only; Table 4.B.17, Table 4.B.18). Fourth, we estimate the ATT using only observations above an elevation of 1,000 metres, followed by all elevations (DDD only; Table 4.B.19, Table 4.B.20). These four checks are shown for 2011-2017 only, although the results for other years are also consistent with those in Figure 4.2.2. Our estimates of forest loss and carbon emissions avoided above an elevation of 1,000 metres are either very low (peatland) or not statistically significant (dryland) (Figure 4.B.1). Finally, using forest cover data based on a tighter definition of forest – “intact primary” forests with no detectable signs of human-caused alteration or fragmentation (Margono *et al.*, 2014; Turubanova *et al.*, 2018) – we find patterns of avoided forest loss and carbon emissions that are consistent with those in Figure 4.2.2 (Table 4.B.21-Table 4.C.2). Unsurprisingly, estimates of the extent of cumulative avoided dryland forest loss and carbon emissions are lower, around 25-35% of our 2018 estimates in Figure 4.2.2 (see also Figure 4.B.2).

### 4.2.3 Testing for leakage

Our estimated effects in Figure 4.2.2 are relatively small, implying a low probability of upward bias due to leakage. We report the results of a regression discontinuity

analysis along the Moratorium’s boundaries, including the treatment effects at the discontinuity each year, between 2005 and 2018 (Figure 4.2.3). The point estimate of this Local Average Treatment Effect (LATE) is always positive after 2013, implying higher deforestation outside the Moratorium. Yet, the point estimates are also positive in 2008, 2009 and 2011, and hover around zero for the entire sample period, excluding an upwards deviation in trend after the Moratorium was implemented. The 95% confidence intervals for the LATE always include zero, thus failing to identify significant leakage from the Moratorium to the surrounding areas, either before or after 2011. Our analysis over various bandwidths suggests no evidence of leakage even within some considerable distance from the Moratorium (Section 4.E).



**Figure 4.2.3:** Regression discontinuity LATE (with Calonico *et al.* (2014) bandwidth), 2005-2018. Scatterplots of all the observations within and outside the Moratorium’s boundaries are shown in Section 4.E. Error bars denote the 95% CI.

#### 4.2.4 Meeting Indonesia’s NDC commitments and the effective carbon price

We put our results into perspective by comparing our estimates of emissions reductions between 2011-17 with Indonesia’s aggregate emissions and its NDC emissions reduction commitments (Table 4.2.1). Our estimated average annual emissions reductions, of around 10.4-13.0 MtCO<sub>2</sub>-eq (including our peatland forest estimates and factoring in below-ground carbon), at most, comprise only 0.38-0.47% of annual aggregate emissions (all sectors), and 0.83-1.05% of emissions from the forest sector. The emissions avoided due to the Moratorium comprise 10.3-12.9% and 3.1-3.8% of the 29% (unconditional) NDC target in 2020 and 2030, respectively. These shares fall to 7.8-9.8% (2020) and 2.0-2.5% (2030) when we consider the 41% (conditional) NDC target.

Norway agreed to pay Indonesia US\$56.2 million for the Indonesia-Norway partnership's estimate of 11.2 MtCO<sub>2</sub>-eq of avoided emissions in 2017, including emissions from avoided peat fires and peat decomposition. Dividing this payment by our estimated average annual emissions reduction during the 2011-17 period gives effective carbon price ranges of, respectively, US\$4.3-5.4/tCO<sub>2</sub>-eq and US\$1.8-2.3/tCO<sub>2</sub>-eq, with and without the share of the payment for avoided peat fires and decomposition. All of these estimates are within range of Norway's proposed carbon price, US\$5/tCO<sub>2</sub>-eq. Applied to our estimates of cumulative emissions reductions over the entire 2011-17 period, Norway's payment has effectively bought emissions reductions at less than US\$1/tCO<sub>2</sub>-eq, thus representing value for money from Norway's perspective. Indeed, from a global perspective these are very cost-effective emissions reductions.

**Table 4.2.1:** The Moratorium’s contribution to Indonesia’s NDC commitments and the effective carbon price, 2011-2017

<b>Estimator</b>	<b>DD</b>	<b>DDD</b>
Avoided emissions (above-ground carbon only, MtCO <sub>2</sub> -eq)		
<b>Dryland</b>	75.4	59.0
<b>Peatland</b>	-2.0	-0.6
<b>Total</b>	73.4	58.4
<b>Annual average (2011-2017)</b>	10.5	8.3
Avoided emissions (above- & below-ground carbon, MtCO <sub>2</sub> -eq)		
<b>Total</b>	91.1	72.5
<b>Annual average (2011-2017)</b>	13.0	10.4
% of Indonesia’s emissions (above-ground carbon only)		
<b>% all emissions (2.2GtCO<sub>2</sub>/year)</b>	0.47	0.38
<b>% forest emissions (1.0GtCO<sub>2</sub>/year)</b>	1.05	0.83
% of Indonesia’s emissions (above- & below-ground carbon)		
<b>% all emissions (2.2GtCO<sub>2</sub>/year)</b>	0.59	0.47
<b>% forest emissions (1.0GtCO<sub>2</sub>/year)</b>	1.30	1.04
Comparison with Indonesia’s NDC 2030 commitments (%)		
<b>% unconditional (29%, 2020)</b>	12.9	10.3
<b>% conditional (41%, 2020)</b>	9.8	7.8
<b>% unconditional (29%, 2030)</b>	3.8	3.1
<b>% conditional (41%, 2030)</b>	2.5	2.0
Effective carbon price (US\$/tCO <sub>2</sub> )		
<b>With peatland payments (US\$56m/total)</b>	\$0.6	\$0.8
<b>With peatland payments (US\$56m/annual average)</b>	\$4.3	\$5.4
<b>No peatland payments (US\$24m/total)</b>	\$0.3	\$0.3
<b>No peatland payments (US\$24m/annual average)</b>	\$1.8	\$2.3

*Notes:* “DD” denotes that the results are derived from the ATT estimated using the non-parametric difference-in-difference approach. “DDD” denotes that the results are derived from the ATT estimated using the non-parametric triple difference approach. Underlying ATT estimates are for 2011-2017; those for peatland are not significantly different from zero. Details of all calculations in the table are in [Section 4.4](#).

## 4.3 Discussion

### 4.3.1 Impacts of Indonesia's Moratorium on forest loss and emissions

The centerpiece of one of Norway's pioneering REDD+ partnerships, Indonesia's Moratorium and the associated US\$ 1 billion pledge, represented an ambitious scaling-up of tropical forest conservation efforts. We found a relatively small effect of the Moratorium in slowing deforestation. Our cumulative estimates to 2018, using a quasi-experimental programme evaluation approach, are at the lower-bound of estimates in previous research (Busch *et al.*, 2015). We also found evidence of a positive impact on dryland forest that materialised earlier, between 2012 and 2016 (Fig. 4.2.2), a period that has not been assessed for emissions reductions by the Indonesia-Norway partnership. The magnitude of our estimated impacts, and our regression discontinuity results, suggests either that Moratorium areas were mostly economically marginal and, compared to matched control areas, unlikely to experience a large effect from the Moratorium, or that deforestation has continued largely unchecked by the Moratorium. The general secular decline in forest cover in all areas (Fig. 4.2.1) suggests the latter explanation is more likely.

What matters for REDD+ is how these positive impacts on forest cover translate into carbon emissions reductions. Our estimates in Table 4.2.1 are in line with a projection that the Moratorium had the potential to cumulatively reduce emissions by nearly 200 MtCO<sub>2</sub>-eq by 2030 (Wijaya *et al.*, 2017), at an estimated annual average of 9.4 MtCO<sub>2</sub>-eq. Note, however, that our impact estimates stem from differences between two declining paths of forest cover over time. While straightforward to measure against a business-as-usual target, this just delays emissions from deforestation. To stop emissions permanently, deforestation needs to be halted not slowed. With this in mind, our estimates accounted for, at most, around 13% and 4% of Indonesia's NDC (unconditional) commitment to reduce GHG emissions by 29% in 2020 and 41% in 2030, respectively. Our estimates suggest that Indonesia is unlikely to meet this commitment given that most of it (17.2 percentage points, or three-fifths of the 29% target) is supposed to be met via the country's forest sector (Government of Indonesia, 2018). Peatland forest loss and peat fires are key contributors to Indonesia's share of global, forest-based emissions yet our results suggest that the Moratorium has had no meaningful impact on peatland forests. Norway's proposed payment of US\$56.2 million included emissions reductions from avoided peat decomposition and fires in 2017. Our peatland results imply that, viewed purely in terms of performance, this share of the payment could be justifiably withheld.

### 4.3.2 Comparing estimates of impact

That our estimates of impact differ from those calculated by the Indonesia-Norway partnership is primarily due to the choice of baseline, or counterfactual, against which impact was measured. For emissions reductions between 2018-2020, the Indonesia-Norway partnership planned to adopt a historical baseline similar to the one the



partnership used for estimating reductions in 2017 (MoEF, 2019). How this baseline is estimated originates from the calculation of Indonesia’s arguably more generous Forest Reference Emission Level (FREL). Submitted to the UNFCCC in 2016, the FREL is based on the average annual deforestation rate between 1990 and 2012 (Government of Indonesia, 2016; Government of Indonesia, 2018). Indonesia’s FREL provided the basis for a proposed US\$103.8 million payment from the Green Climate Fund (GCF) for an estimated 20.3 MtCO<sub>2</sub>-eq reduction in carbon emissions between 2014-2016 (Mongabay, 2020a).

Although both the Norwegian and GCF payments are supposed to be “results-based”, the baselines used to estimate these emissions reductions emerged as a consequence of political negotiations and are subject to precisely the biases that we attempted to eliminate in our programme evaluation approach. Thus, they are arguably independent of performance (West *et al.*, 2020), that is, the counterfactuals constructed by the Indonesia-Norway partnership and the GCF could be driven entirely by stochastic shocks and economic factors unrelated to REDD+ efforts. Our causal framework explicitly attempted to balance these biases. On the basis of historical performance, our results suggest that Indonesia could legitimately claim an even larger payment from Norway, up to US\$339-434.5 million on the basis of a US\$5/tCO<sub>2</sub>-eq carbon price, for cumulative emissions reductions between 2011 and 2018. Fig. 4.2.2 shows a predominantly steady cumulative effect of the Moratorium over time, particularly from 2016 onward. The Moratorium seems to have had a causal impact on avoided deforestation.

### 4.3.3 Beyond the Moratorium

The partnership underlying the Moratorium was unilaterally terminated by Indonesia in 2021, apparently due, in part, to delays in release of the payment by Norway (Mongabay, 2021b). Even though this implies that the Moratorium is unlikely to continue, at least not in its current form, large-scale, area-based initiatives, in the form of jurisdictional REDD+ schemes, are likely in the future. For example, the Lowering Emissions by Accelerating Forest finance (LEAF) private-public coalition was established in 2021 to mobilise at least US\$1 billion for area-based tropical forest conservation (Mongabay, 2020a). Such initiatives could usefully learn from how the Moratorium performed with respect to emissions reductions.

First, the small size of the Moratorium’s impact suggests limited compliance. Improving compliance might increase impact yet patronage linkages, between large-scale industrial plantation companies and local politicians, ensured weak monitoring and enforcement (Varkkey, 2013). Future REDD+ initiatives could help bolster local monitoring and enforcement capabilities. The ongoing One Map process, to resolve inconsistencies resulting from the use of different data and maps by creating a national standard of land cover and usage, could further help strengthen transparency and improve forest governance (LTS International, 2018).

Second, future initiatives could also help incentivise reductions in emissions from deforestation by local forest users not originally targeted by the Moratorium, such

as smallholders engaged in logging and palm oil production (Clough *et al.*, 2016; Ferraro and Simorangkir, 2020), who reportedly contributed one-fifth of nationwide forest loss between 2001 and 2016 (Austin *et al.*, 2019). The forestland claims of local forest users have been strengthened by the Village Law (6/2014) (Antlöv *et al.*, 2016), which, combined with the millions of hectares of forest pledged for “social forestry” initiatives, suggest new conservation opportunities (Myers *et al.*, 2017). These opportunities could be aligned with the goals of future initiatives, although there remains a risk that continued policy layering could exacerbate ambiguity with respect to forest regulations and enforcement (Erbaugh and Nurrochmat, 2019). Also, given ongoing uncertainties over forest users’ land rights, new initiatives should pay careful attention to representation and recognition notions of justice as a means of legitimising REDD+ at the local scale (Myers *et al.*, 2018).

Third, the Moratorium lacked formal allocation and distributional mechanisms. Our results are based on the aggregate ATT but the district-level ATT indicate where the positive effects were likely generated (Figure 4.B.3). Consistent with calls from forest-rich districts for an “ecological fiscal transfer” scheme based on ecological performance (Jakarta Post, 2019), payments from future initiatives could be distributed to districts that demonstrate emissions reductions. The Regional Governance Law (23/2014), however, shifted control over forests from district to provincial governments and established a greater administrative role for forest management units (Sahide *et al.*, 2016). Thus, both provincial and district governments are likely to play a role in any benefits transfer system, perhaps via intergovernmental fiscal transfers (IFT) (Nurfatriani *et al.*, 2015). Presently focused on timber production revenues, Indonesia’s forest sector IFT could, in theory, be used to transfer REDD+ funds (Irawan *et al.*, 2013).

#### 4.3.4 Delivering carbon value for money and meeting the NDC commitments

The value of estimated cumulative emissions reductions, even on the basis of a relatively low carbon price of US\$5/tCO<sub>2</sub>-eq, comfortably exceeds the amount Norway agreed to pay for emissions reductions in 2017. From the perspective of Norway’s government, and the global community, Norway’s payment could be characterised in terms of abatement costs, that is, the sum that Indonesia’s government is willing to accept (WTA) to reduce emissions from deforestation. However, the global benefit, and in principle the willingness to pay (WTP) for emissions reductions, is the social cost of carbon (SCC). The SCC is estimated to range between approximately US\$40-200/tCO<sub>2</sub>-eq, e.g. (Hänsel *et al.*, 2020). The question of how the discrepancy between WTP and WTA is shared between donor and recipient countries was resolved by the Indonesia-Norway partnership, moving far closer to the WTA than to the WTP. Thus, Norway, and indeed the global community, would be getting “good value” for emissions reductions in Indonesia. The Moratorium appeared to be cost-effective but with a very skewed share of the global surplus transferred to a carbon-rich yet poor country. Although Indonesia accepted the US\$5/tCO<sub>2</sub>-eq price offered by Norway, this is arguably an unfair distribution of the surplus given estimates of the SCC.

With the ending of the Indonesia-Norway partnership, Indonesia could negotiate a higher price with, for example, an initiative like LEAF, which is offering a minimum of US\$10/tCO<sub>2</sub>-eq (Mongabay, 2021a).

Our analysis emphasises that emissions reductions, although cost-effective, still require large transfers, even at low carbon prices. Much steeper emissions reductions are clearly needed to reach the NDC targets but the cost of these reductions is unlikely to be met by a single country or initiative. Indeed, Indonesia estimates that to meet its NDC commitments, around US\$5.5 billion is required between 2018 and 2030 for the country's forest sector alone (Government of Indonesia, 2018). When REDD+ first emerged in the 2000s, there were initial calls for US\$10-15 billion of funding per year to cut global deforestation by half (Stern, 2006). These funding needs were based on opportunity cost calculations, which will be higher for high-value agricultural commodities such as palm oil. Yet, by the 2010s pledges to the value of only US\$10 billion had been made for REDD+ (Norman and Nakhouda, 2015).

It was hoped that a global climate agreement, incorporating a cap-and-trade system, would generate sufficient and sustainable sources of finance for the protection of tropical forest carbon stocks. Given that such a system has yet to materialise, it has fallen upon individual countries to voluntarily finance REDD+ initiatives around the world. Norway's contribution to date exceeds that of all other countries but is insufficient to protect tropical forest carbon stocks at a scale necessary to meaningfully contribute to global climate change mitigation efforts. As the world's attention moves beyond the COP26 in Glasgow, where an ambitious global commitment was made to halt deforestation by 2030, the critical climate role of forests needs to be matched by a global willingness to pay for it.

## 4.4 Materials and Methods

### 4.4.1 Data

Our outcome variable is forest cover in hectares, with data spanning the period 2000-2018 drawn from the Global Forest Change dataset (Hansen *et al.*, 2013). The 2004-2018 period is selected for matching data in our analysis while data for the 2000-2004 period are used in separate placebo tests that determine whether there are violations in the parallel trends assumption (see below). The forest cover data, obtained for the whole of Indonesia, are spatially explicit and not defined according to different forest classes. Our units of analysis are grid cells of 1.2km-by-1.2km (144 ha), which accommodate 1,600 pixels at a resolution of 30m-by-30m (i.e. Landsat8 pixel size). A scale of 1.2km-by-1.2km is chosen because it allows for a similar scale across the different sources of data used in our analysis and minimises the risk of grid cells overlapping treatment and control areas. For each grid cell and for each year of our study period, we count the number of pixels where forest loss is recorded and then convert pixels to hectares by multiplying by 0.09. We account for the precise fraction of a pixel that falls within a grid cell.

Tree cover in the year 2000, the base year of the Global Forest Change dataset (Hansen *et al.*, 2013), is defined as canopy closure for all vegetation taller than 5m in height, and is encoded as a percentage per pixel, in the range 0–100. Forest loss during the period 2001–2018 is defined as a stand-replacement disturbance, or a change from a forest to non-forest state at the pixel scale. Encoded as either 0 (no loss) or else a value in the range 1–18, representing loss detected in the year 2001–2018, respectively. A pixel is categorised as forested if its canopy cover is greater or equal to 25%, below which a pixel changes its state from forest to non-forest (Hansen *et al.*, 2010). We obtain peat depth data from Gumbricht *et al.* (2017), which are used to subdivide grid cells into peatland and dryland forest types. These two types are analysed separately due to their different ecological characteristics that are relevant to changes in forest cover and carbon emissions.

Forest areas covered by the Moratorium supposedly included all of Indonesia’s primary and peatland forests, and are determined using Moratorium shapefiles obtained from the Ministry of Forestry of the Republic of Indonesia (2011). Our treatment group includes forest areas within the Moratorium boundaries that were established in 2011 (Figure 4.A.2). Since 2011, the Moratorium’s boundaries have shifted due to forestland being re-designated by Indonesia’s district governments and dropped out of the Moratorium before typically being licensed out to concessionaires (Gaveau *et al.*, 2013; Marlier *et al.*, 2015; Enrici and Hubacek, 2016)(Section 4.A). Although legal, this re-designation of forestland is effectively a behavioural response to the Moratorium and hence, should be included in estimates of impact. Thus, the 2011 Moratorium boundaries are assumed constant throughout our treatment periods. Shapefiles for the location of palm oil, timber and logging concessions established before the start of the Moratorium were originally obtained from the Ministry of Forestry (2014). Digitised by Greenpeace, this was the most comprehensive source of concessions data available.

We compile a matrix of control variables and time-invariant characteristics at the grid-cell level by combining different sources of geo-referenced data: information on altitude, slope, and distance from major roads from the WorldPop repository (WorldPop, 2018); grid-cell level travel time to major cities (Nelson, 2008); and, a population trend based on counts for the years 2004 and 2010 (Rose *et al.*, 2020). Grid-cell level above-ground carbon stock values are estimated by dividing data on above-ground biomass density, from the Global Forest Watch dataset, by 0.5 (based on Baccini *et al.* (2012)).

The cumulative impacts of the Moratorium on forest cover and carbon emissions are estimated each year in the periods 2011-2018 and 2011-2017 for dryland and peatland forest, respectively. The latter is estimated only up until 2017 due to the additional restrictions on peatland forest conversion implemented in 2017 (Alisjahbana and Busch, 2017).

### 4.4.2 Empirical approach

The Moratorium mandated that district governments stop issuing new concession licenses in forest areas covered by the Moratorium. Evaluating the causal impact of the Moratorium is complicated by selection bias: forest areas covered by the Moratorium differ in their observable and unobservable characteristics. Any imbalance in these characteristics implies that simply comparing the extent of forest cover in the Moratorium and non-Moratorium areas will capture pre-existing imbalances in, for example, their suitability for palm oil or timber production, thus confounding the estimate of the treatment effect. Descriptive statistics illustrate the differences between the Moratorium and non-Moratorium forest areas (Table 4.C.1-Table 4.C.8). The values of several variables that determine the suitability of grid cells for new concessions differ between Moratorium and non-Moratorium areas. For example, in the unmatched data, Moratorium grid cells have, on average, a higher elevation and are further away from roads and cities than non-Moratorium cells, making them less suitable for concessions, other things equal. It is necessary to address the imbalance of these observable characteristics to estimate the causal effect of the Moratorium on forest cover.

We adopt a difference-in-differences (DD) research design to control for observable and unobservable confounding characteristics in the estimation of the treatment effect. Empirical testing leads us to prefer a matched triple difference estimator (DDD). Via matching Moratorium and non-Moratorium cells on the basis of their observable characteristics, we argue that the confounding effect of unobservable characteristics is also controlled for, thus generating an unbiased estimate of the treatment effect (Imbens, 2014; Wooldridge, 2010; Abadie and Imbens, 2008). Prior to estimation, we adjust the Moratorium and non-Moratorium samples by excluding cells which are unlikely to become concessions for agronomic or jurisdictional reasons. We first exclude cells which are part of the Indonesian protected area network, both within and outside the Moratorium, as conversion in these cells is already strictly prohibited. Second, we focus on unconverted dryland and peatland forest outside of concessions. We then remove all cells outside of concessions with an elevation of 1,000m or more above sea level. The likelihood of these cells being a realistic proposition for a concession in either Moratorium or non-Moratorium areas is close to zero because above 1,000m land is unsuitable for palm oil cultivation (Austin *et al.*, 2015) and for *Acacia mangium*, the main tree species employed for the production of wood pulp and paper (Krisnawati *et al.*, 2011). These adjustments represent our first attempt to balance the sample in terms of the likelihood of concessions being granted. The resulting dataset has 567,634 cells, of which there are 160,012 treated observations (28.2%). Next, we deploy a matching procedure to match individual grid cells in Moratorium areas with counterfactual grid cells in non-Moratorium areas, and vice versa. Our main results use propensity score, one-to-one caliper matching for this purpose. After matching, we retain 152,118 treated cells (7,894 dropped cells), and 198,794 control cells, for a total of 358,806 cells. The dropped cells represent a reduction of 2.2% (4.9% of the treated group) due to the exclusion of imprecise (outside of the caliper) matches. This matched dataset is used for both parametric and non-parametric estimators, to ensure balanced characteristics be-

tween the treatment and control groups and to facilitate easier comparisons between estimators (Imbens, 2014). When analysing the Moratorium's impact on different forest types (dryland, peatland) the propensity score is estimated separately and different matched samples arise.

Having balanced the sample in this way, the identification of a causal estimate of the Moratorium's impact on forest cover stems from a DD research design. It is well-known that under their identifying assumptions, DD estimators identify the Average Treatment on the Treated (ATT) (e.g. Lee, 2005, Ch. 4), which can be understood as the impact of the Moratorium on forest areas covered by the Moratorium and hence, is a policy-relevant treatment effect. With the dropping of cells, we are no longer identifying the ATT, although dropping fewer than 5% of treated observations arguably results in a close approximation to the ATT. Where  $d_i$  indicates whether an individual grid cell  $i$  is in the treatment group ( $=1$ ) or not ( $=0$ ), and  $Y_{1iT_1}$  and  $Y_{0iT_1}$  are, respectively, the potential outcomes (forest cover) in the treated ( $=1$ ) and untreated ( $=0$ ) states for grid cell  $i$  at post-treatment time  $T_1$ , the ATT is defined as:

$$ATT = E[Y_{1iT_1} - Y_{0iT_1} | d_i = 1] \quad (4.1)$$

Following Lee (2005, p. 101), if the potential outcomes have the separable form  $Y_{kit} = \mu_{kit} + \lambda_i + u_{kit}$  for treatment states  $k = 0, 1$ , and individual, grid-level fixed effect  $\lambda_i$ , the ATT is identified by a DD estimator of the form:

$$DD = E[\Delta_{T_1, T_0} Y_{it} | d_i = 1] - E[\Delta_{T_1, T_0} Y_{it} | d_i = 0] \quad (4.2)$$

where  $Y_{it}$  is the observed data and  $\Delta_{T_1, T_0}$  is the change operator between the pre-treatment period  $T_0$  and  $T_1$  (see also Section 4.D). The DD estimator controls for individual fixed effects  $\lambda_i$  by taking differences at the individual level. A necessary condition to identify the ATT in this way is the parallel trends assumption:

$$E[\Delta_{T_1, T_0} u_{1it} | d_i = 1] - E[\Delta_{T_1, T_0} u_{0it} | d_i = 0] = 0 \quad (4.3)$$

meaning that the unobservable characteristics determining forest cover must be identical in expectation, otherwise they will confound the estimate of impact.

The DD estimator in Equation 4.2 can be estimated parametrically or non-parametrically (e.g. matching). The parametric DD estimator we use takes the following form, and is estimated using a fixed-effects estimator:

$$Y_{it} = \alpha + \sum_{s=\tau}^T \beta_{1s} D_{sit} + \sum_{k=2}^n \beta_k X_{kit} + \lambda_i + \theta_t + \varepsilon_{it} \quad (4.4)$$

where  $Y_{it}$  is forest cover in (non-concession) grid cell  $i$  in year  $t$ ,  $D_{it}$  is the time-varying Moratorium treatment indicator,  $X_{kit}$  are  $n$  potential time-varying control variables. The  $\beta_{1s}$  coefficients represent the DD estimates of ATT for each post-treatment year

$s$  between 2012 and 2018. This basic model controls for time-invariant characteristics via the individual, grid-level fixed effects,  $\lambda_i$ , and time fixed effects,  $\theta_t$ , which capture shocks common to all grid cells, such as weather shocks.

We select among a number of different parametric and non-parametric (matched) DD and triple difference (DDD) estimators through a four-step process of model selection. In the first step, we estimate parametric DD models of the form described in Equation 4.13, including district-by-year trends, pre-treatment forest cover-by-year interactions, and both clustered (at the district level) and Conley standard errors (this accounts for spatial autocorrelation using the *fixest* package in R), but no further control variables ( $X_{kit}$ ), before comparing these estimates to a propensity score, one-to-one caliper, matched DD estimator (Table 4.D.1). Clustering and Conley standard errors do not affect the results. To rule out the possibility that the balance in the sample of observable characteristics between Moratorium and non-Moratorium areas causes differences between the parametric and non-parametric estimates, we use the same matched sample for the parametric DD estimation as is used for matched DD estimators (see Imbens, 2014). We then undertake sensitivity analysis on the non-parametric estimators by relaxing the precision of the matching procedure in two ways: (i) widening the caliper; and, (ii) sampling matches without replacement. The sensitivity analysis suggests that the matching estimates are sensitive only to extreme reductions in precision of the matching (no replacement or no caliper) (Table 4.D.1). Extreme sensitivity of the parametric estimator to the inclusion of district-by-year trends suggests that heterogeneity across grid cells is a potentially important confounding factor. Matching estimators deal more flexibly with heterogeneity, and in matched DD estimators this can include heterogeneous trends (Heckman *et al.*, 1997). There is also empirical evidence to show that the parallel trends assumption is more likely to hold with matched DD rather than parametric DD (e.g. Ryan *et al.*, 2018). To account for heterogeneous trends, our matching procedure matches Moratorium and non-Moratorium grid-cells very precisely, and explicitly, on pre-treatment trends. Given the sensitivity of the parametric estimator and the fact that matching estimators are better equipped to deal with heterogeneity we opt for a non-parametric DD approach, among which we include propensity score matched DD.

Our central estimates use propensity score, one-to-one, caliper matching. The matching variables we use capture important differences between the Moratorium and non-Moratorium grid cells, their dynamics and suitability for future concessions. We use pre-treatment values of: distance to concessions (palm oil, timber and logging), distance to roads and cities, population (2005 and 2010), forest cover for each year from 2005 to 2010, elevation, slope, peat depth and above-ground carbon stock in the year 2000. Matching on pre-treatment outcomes (forest cover) and population in more than one pre-treatment year (2005-2010 for forest cover, 2005 and 2010 for population) attempts to control for heterogeneous pre-treatment trends and levels between Moratorium and non-Moratorium areas.

The matched DD estimator takes the following form and estimates  $ATT_{DD,T}$  for time horizon  $T$  using forest cover data  $Y_{it}^M$  from the Moratorium grid cell  $i$  matched

with data  $Y_{it}^{j,NM}$  in cell  $j$  from the non-Moratorium grid cells:

$$\widehat{ATT}_{DD,T} = \frac{1}{N_M} \sum_i I_0 \left[ [Y_{i,T}^M - Y_{i,T_1}^M] - [Y_{i,T}^{j,NM} - Y_{i,T_1}^{j,NM}] \right] \quad (4.5)$$

where  $I_0$  is an indicator variable that is equal to 1 if a grid cell  $i$  in the Moratorium has a counterfactual grid cell  $j$  in the non-Moratorium area whose propensity scores  $p_i$  and  $p_j$  fall within the caliper:

$$|p_i - p_j| < \epsilon$$

where  $\epsilon$  is a predetermined distance in propensity score space. The caliper defines the set of one-to-one matches from the non-Moratorium area,  $C(i)$ , such that  $j \in C(i)$ . Therefore,  $I_0 = I(\min |p_i - p_j| : |p_i - p_j| < \epsilon)$ . Our large sample allows us to choose a very precise caliper (0.0001) without substantially reducing the sample size.

In the second step, we subject this matched DD estimator to two separate placebo (falsification) tests in time, both of which model a placebo Moratorium implemented in 2004. Placebo Test 1 retains the matches from the real pre-treatment phase (2005-2010) while Placebo Test 2 re-estimates the propensity scores for the pre-placebo-treatment phase (2000-2004) (see [Table 4.D.2-Table 4.D.5](#) and [Figure 4.D.1](#)). In some cases, the null hypothesis is rejected for the matched DD estimator, suggesting a failure of the parallel trends assumption. For this reason, in the third step, we use a matched triple differences (DDD) estimator inspired by ([Chabé-Ferret and Subervie, 2013](#); [Chabé-Ferret, 2017](#)), which applies a correction for non-parallel trends. The DDD estimator takes the following form:

$$\widehat{ATT}_{DDD,T} = \frac{1}{N_M} \sum_i I_0 \left\{ \left[ [Y_{i,T}^M - Y_{i,T_1}^M] - [Y_{i,T}^{j,NM} - Y_{i,T_1}^{j,NM}] \right] - \left[ [Y_{i,T_1'}^M - Y_{i,T_0}^M] - [Y_{i,T_1'}^{j,NM} - Y_{i,T_0}^{j,NM}] \right] \left( \frac{T-T_1}{T_1'-T_0} \right) \right\} \quad (4.6)$$

The third line reflects the correction for the non-parallel pre-treatment trends between  $T_0$  and  $T_1'$  with a correction  $\left( \frac{T-T_1}{T_1'-T_0} \right)$  to adjust the trend correction for potentially different pre- and post-treatment time horizons. The matched DD results come from the estimator in [Equation 4.5](#), and the matched DDD results come from the estimator in [Equation 4.6](#). Finally, we subject the DDD estimator to a placebo in time test of the parallel trends assumption, Placebo Test 3 ([Table 4.D.6](#)). The years used to estimate [Equation 4.5](#) and [Equation 4.6](#) are  $T_0 = 2004$ ,  $T_1' = 2010$ ,  $T_1 = 2011$  and  $T = \text{endpoint year}$ . For the placebo tests the placebo treatment



year is 2005 and the pre- and post-treatment periods considered are, respectively, 2000-2004 and 2005-2010, with sensitivity using 2005-2011. In each case, we use the *Matching* routine in R (Sekhon, 2011).

Finally, in step four, we undertake a spatial placebo test to evaluate the robustness of the matched DDD estimator. The spatial placebo test uses the DDD estimator in Equation 4.6 and applies it to concession Moratorium and concession non-Moratorium grid cells for dryland and peatland forest. In theory, the Moratorium has no effect on forest cover in concession grid cells because conversion is still allowed on concessions in Moratorium areas that were granted pre-Moratorium. This observation forms the basis of the null hypothesis, with the alternative hypothesis that Moratorium and non-Moratorium concession forest cover is evolving in different ways, hence falsifying the parallel trends assumption. Results of the spatial placebo test suggest support for the null hypothesis (Table 4.D.7-Table 4.D.8).

The resulting, preferred estimator is the propensity score, one-to-one caliper matched DDD estimator. Covariate balance tables are reported for all specifications (Table 4.D.11-Table 4.D.20). Standard errors are calculated according to the consistent Abadie-Imbens procedure for matching (Abadie and Imbens, 2006).

We undertake several robustness checks on our estimates in Figure 4.2.2. These estimates are based on untrimmed samples, so we first obtain estimates of the ATT for a sample with a stricter definition of forest cover at the grid-cell level: 30% (or more) and 60% (or more) forested pixels in a grid cell, in 2005. These samples are obtained by estimating the mean forest cover of 2005 pixels in the grid cells. Second, we conduct a robustness analysis on alternative approaches to matching which do not rely on the propensity score, namely Coarsened Exact Matching (CEM) (King and Zeng, 2006; Iacus *et al.*, 2012). CEM addresses the possibility that propensity score matching may introduce biases due to the way in which it reduces the dimensionality of the matching problem to matching on a single dimension: the propensity score (e.g. King and Nielsen, 2019). We undertake CEM using the same matching variables as before. As with propensity score matching, sample sizes are also sensitive to choices of matching variables, the coarseness of matching and other implementation decisions in CEM. For this reason, following Iacus *et al.* (2009) and Iacus *et al.* (2012), we undertake four separate CEM routines in which the control variables are used either in the exact matching algorithm, or as control variables in a covariate adjustment step after matching has occurred (Iacus *et al.*, 2012). The DDD CEM covariate-adjusted estimate of ATT over the period of the Moratorium from period  $T_1$  and  $T$ , is the estimate of  $\beta_1$  in the following:

$$\Delta_{T_1, T}^{DDD} Y_i = \alpha + \beta_1 d_i + \sum_{k=2}^n \beta_k X_{ki} + \epsilon_i \quad (4.7)$$

The  $X_{ki}$  are the pre-treatment control variables that provide covariate adjustment, which we exclude from the matching algorithm.  $\Delta_{T_1, T}^{DDD} Y_i$  is the individual-level matched triple-difference (corrected for non-parallel trends) in forest cover that is

constructed by the CEM matching algorithm. We use the *cem* routine in R to undertake the matching, and the *att* routine in R to obtain this covariate-adjusted estimate of the ATT.

Third, we conduct sensitivity analyses on the assumptions used in the central non-parametric DDD propensity score matching estimates. These analyses focus on: (1) caliper width choices, where calipers of 0.01 and 0.001 are used in one-to-one nearest neighbour matching; and, (2)  $k > 1$  nearest neighbours analysis, in which we test the sensitivity of our results to matching on  $k = 2, 3, 5$  nearest neighbours instead of one-to-one. Fourth, we test the sensitivity of the sample being limited to grid cells below an elevation of 1,000m. We estimate separate ATT for: (1) cells at all elevations; and, (2) cells only found above 1,000m (covariate balancing reported in [Table 4.D.21-Table 4.D.24](#)). Finally, we replicate the analysis in [Figure 4.2.2](#) using a tighter definition of forest, specifically, “intact primary” forest, with no detectable signs of human-caused alteration or fragmentation ([Margono \*et al.\*, 2014](#); [Turubanova \*et al.\*, 2018](#)) (covariate balancing reported in [Table 4.D.25](#) and [Table 4.D.26](#)).

A final identification assumption of the DD and DDD estimators is the Stable Unit Treatment Value Assumption (SUTVA): the treatment should not cause leakage to untreated forest areas. Such leakage/spillover effects are a common confounder in the evaluation of area-based policies (e.g. [Andam \*et al.\*, 2008](#); [Gaveau \*et al.\*, 2009](#); [Joppa and Pfaff, 2010](#); [Nelson and Chomitz, 2011](#)). If protection via the Moratorium induces the displacement of forest clearing to outside the Moratorium’s boundaries, deforestation rates inside and outside these boundaries are contemporaneously affected in opposite directions resulting in treatment effects of a higher magnitude than the “true” effects. To check for leakage, we use a Regression Discontinuity Design (RDD) ([Lee and Lemieux, 2010](#); [Calonico \*et al.\*, 2014](#)) with a sharp cut-off at the Moratorium’s boundaries. RDD estimates the Local Average Treatment Effect (LATE) of the Moratorium in the proximity of its boundaries with non-Moratorium land.

We specify separate linear polynomials on both sides of the boundary cut-off following [Gelman and Imbens \(2019\)](#) and [Burgess \*et al.\* \(2019\)](#), and estimate treatment effects via OLS regressions with robust standard errors clustered at the district administrative level. Our preferred results are obtained via separate linear polynomials and optimal bandwidth selection (through the [Calonico \*et al.\* \(2014\)](#) method). We also examine an alternative bandwidth selection algorithm ([Imbens and Kalyanaraman, 2012](#)), additional fixed bandwidths of 5, 10 and 20 km from the Moratorium’s boundaries, and specifications using separate quadratic polynomials of distance ([Figure 4.E.1-Figure 4.E.7](#)). Given the spatial configuration of the Moratorium ([Figure 4.A.2](#)), our bandwidths cover a wide extent of non-Moratorium cells in which leakage could feasibly occur.

### 4.4.3 Estimating carbon emissions

We obtain grid cell-level CO<sub>2</sub>-eq estimates by multiplying the average carbon stock (in tC/ha) per grid cell by our ATT estimates of avoided deforestation (ha). These results are then extrapolated to the total number of Moratorium grid cells in the sample, and converted to tons of carbon equivalent, tCO<sub>2</sub>-eq, by applying a conversion factor of 3.67 (Zarin *et al.*, 2016). Below-ground carbon stocks in Table 4.2.1 are approximated by multiplying above-ground carbon by 1.24 following (Busch *et al.*, 2015).

In Table 4.2.1, total emissions of 2.2 GtCO<sub>2</sub>-eq per year are based on the projected business-as-usual (BAU) emissions for all sectors excluding LULUCF plus estimated annual emissions from LULUCF (1.0 GtCO<sub>2</sub>-eq per year) according to Government of Indonesia (2018) and reported by Climate Action Tracker. The estimate of 1.0 GtCO<sub>2</sub>-eq per year is an annual average over 2001-2018. Indonesia's NDC commitments involve an emissions reduction path which implies specific reductions against the BAU in 2020: an unconditional (without overseas assistance) 29% reduction in emissions compared to projected BAU emissions in 2030, and a 41% conditional (with overseas assistance) emissions reduction.

We measure our estimated avoided deforestation against the 2030 commitment's implied emissions reduction for 2020. The BAU emissions (non-LULUCF) in 2020 is projected to be 1.2 GtCO<sub>2</sub>-eq. The conditional (unconditional) emissions reduction in 2020 represents a 8% (11%) reduction against 2020 BAU emissions. The 2030 29% conditional (unconditional) NDC implies a target of 1.12 (1.09) GtCO<sub>2</sub>-eq in 2020, a reduction of 100 (132.9) MtCO<sub>2</sub>-eq. The impact of the Moratorium is calculated as a percentage of the 100 (132.9) MtCO<sub>2</sub>-eq reduction. BAU (non-LULUCF) and unconditional (conditional) pledged emissions in 2030 are 2.2 GtCO<sub>2</sub>-eq and 1.8 (1.6) GtCO<sub>2</sub>-eq, implying a reduction of 328 (527) MtCO<sub>2</sub>-eq in 2030. Our estimated Moratorium impact is calculated as a percentage of these figures.

Norway's agreed payment of US\$ 56.2 million includes emissions reductions from peatland forest in 2017. Without this the payment falls to US\$ 37 million net of the 35% "set aside factor" (Section 4.A)<sup>1</sup>.

---

<sup>1</sup>All replication materials, including processed datasets, regression routines, and R scripts used to generate our results are available in the [Harvard Dataverse](#).

# References

- Abadie, Alberto and Imbens, Guido W (2006). “Large sample properties of matching estimators”. *Econometrica* 74.1, pp. 235–267.
- Abadie, Alberto and Imbens, Guido W (2008). “Notes and comments on the failure of the bootstrap for matching estimators”. *Econometrica* 76.6, pp. 1537–1557.
- Abood, Sinan A, Ser, Janice, Lee, Huay, Burivalova, Zuzana, Garcia-ulloa, John, and Koh, Lian Pin (2015). “Relative contributions of the logging, fiber, oil palm, and mining industries to forest loss in Indonesia”. *Conservation Letters* 8, pp. 58–67.
- Alisjahbana, Armida S. and Busch, Jonah M. (2017). “Forestry, Forest Fires, and Climate Change in Indonesia”. *Bulletin of Indonesian Economic Studies* 53, pp. 111–136.
- Andam, Kwaw S., Ferraro, Paul J., Pfaff, Alexander, Sanchez-Azofeifa, G. Arturo, and Robalino, Juan A. (2008). “Measuring the effectiveness of protected area networks in reducing deforestation”. *Proceedings of the National Academy of Sciences* 105, pp. 16089–16094.
- Angelsen, A. (2017). “REDD+ as Results-based Aid: General Lessons and Bilateral Agreements of Norway”. *Review of Development Economics* 2, pp. 237–264.
- Angelsen, Arild and Kaimowitz, David (1999). “Rethinking the causes of deforestation: lessons from economic models.” *The World Bank Research Observer* 14, pp. 73–98.
- Antlöv, H., Wetterberg, A., and Dharmawan, L (2016). “Village governance, community life and the 2014 Village Law in Indonesia”. *Bulletin of Indonesian Economic Studies* 52, pp. 161–183.
- Austin, Kemen G., Kasibhatla, Prasad S., Urban, Dean L., Stolle, Fred, and Vincent, Jeffrey (2015). “Reconciling oil palm expansion and climate change mitigation in Kalimantan, Indonesia”. *PLOS ONE* 10.5, pp. 1–17.
- Austin, Kemen G., Schwantes, Amanda, Gu, Yaofeng, and Kasibhatla, Prasad S. (2019). “What causes deforestation in Indonesia?” *Environmental Research Letters* 14.2, p. 024007.
- Austin, Kemen, Alisjahbana, Ariana, Darusman, Taryono, Rachmat, Boediono, Budianto, Bambang Eko, Purba, Christian, Indrarto, Giorgio Budi, Pohnan, Erica, Putraditama, Andika, and Stolle, Fred (2014). *Indonesia’s Forest Moratorium: Impacts and Next Steps*. Tech. rep. January 2014. World Resources Institute, pp. 1–16.
- Baccini, A. *et al.* (2012). “Estimated carbon dioxide emissions from tropical deforestation improved by carbon-density maps”. *Nature Climate Change* 2.3, pp. 182–185.
- Boer, Henry J (2020). “Power, REDD+ and reforming forest governance in Indonesia”. *Third World Quarterly* 41, pp. 783–800.

- Burgess, Robin, Costa, Francisco, and Olken, Ben (Aug. 2019). *The Brazilian Amazon's double reversal of fortune*. Working Paper.
- Busch, Jonah and Ferretti-Gallon, Kalifi (2017). "What Drives Deforestation and What Stops It? A Meta-Analysis". *Review of Environmental Economics and Policy* 11, pp. 3–23.
- Busch, Jonah *et al.* (2015). "Reductions in emissions from deforestation from Indonesia's moratorium on new oil palm, timber, and logging concessions". *Proceedings of the National Academy of Sciences* 112, pp. 1328–1333.
- Calonico, Sebastian, Cattaneo, Matias D., and Titiunik, Rocio (2014). "Robust non-parametric confidence intervals for regression-discontinuity designs". *Econometrica* 82.6, pp. 2295–2326.
- Chabé-Ferret, Sylvain (2015). "Analysis of the bias of Matching and Difference-in-Difference under alternative earnings and selection processes". *Journal of Econometrics* 185.1, pp. 110–123.
- Chabé-Ferret, Sylvain (2017). *Should We Combine Difference In Differences with Conditioning on Pre-Treatment Outcomes?* Tech. rep. Toulouse: Toulouse School of Economics, INRA and IAST.
- Chabé-Ferret, Sylvain and Subervie, Julie (2013). "How much green for the buck? Estimating additional and windfall effects of French agro-environmental schemes by DID-matching". *Journal of Environmental Economics and Management* 65.1, pp. 12–27.
- Clough, Y. *et al.* (2016). "Land-use choices follow profitability at the expense of ecological functions in Indonesian smallholder landscapes". *Nature Communications* 7, p. 13137.
- Duchelle, Amy E., Simonet, Gabriela, Sunderlin, William D., and Wunder, Sven (2018). "What is REDD+ achieving on the ground?" *Current Opinion in Environmental Sustainability* 32, pp. 134–140.
- Enrici, Ashley and Hubacek, Klaus (2016). "Business as usual in Indonesia: governance factors effecting the acceleration of the deforestation rate after the introduction of REDD+" *Energy, Ecology and Environment* 1, pp. 183–196.
- Erbaugh, J.T. and Nurrochmat, D.R. (2019). "Paradigm shift and business as usual through policy layering: Forest-related policy change in Indonesia (1999-2016)". *Land Use Policy* 86, 136–146.
- FAO (2021). *FAOSTAT*. URL: <http://www.fao.org/faostat/en/#home>.
- Ferraro, Paul J. and Simorangkir, Rhita (2020). "Conditional cash transfers to alleviate poverty also reduced deforestation in Indonesia". *Science Advances* 6, eaaz1298.
- Gaveau, David L A *et al.* (2013). "Reconciling forest conservation and logging in Indonesian Borneo". *PLoS ONE* 8, b.
- Gaveau, David L.A., Epting, Justin, Lyne, Owen, Linkie, Matthew, Kumara, Indra, Kanninen, Markku, and Leader-Williams, Nigel (2009). "Evaluating whether protected areas reduce tropical deforestation in Sumatra". *Journal of Biogeography* 36.11, pp. 2165–2175.
- Geist, Helmut J. and Lambin, Eric F. (2002). "Proximate Causes and Underlying Driving Forces of Tropical Deforestation: Tropical forests are disappearing as the

- result of many pressures, both local and regional, acting in various combinations in different geographical locations”. *BioScience* 52, pp. 143–150.
- Gelman, Andrew and Imbens, Guido (2019). “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs”. *Journal of Business & Economic Statistics* 37.3, pp. 447–456.
- Government of Indonesia (2018). *Indonesia Second Biennial Update Report Under the United Nations Framework Convention on Climate Change*. Directorate General of Climate Change, Ministry of Environment and Forestry, Government of Indonesia report.
- Government of Indonesia (2016). *National Forest Reference Emission Level for Deforestation and Forest Degradation: Submission by Indonesia*. Post Technical Assessment by UNFCCC. Directorate General of Climate Change. The Ministry of Environment and Forestry, Government of Indonesia.
- Government of Indonesia (2011). *Presidential Instruction No. 10/2011: Moratorium on granting of new licenses and improvement of natural primary forest and peatland governance*. Jakarta, Indonesia.
- Government of Norway (2010). *Letter of Intent on Cooperation on Reducing Greenhouse Gas Emissions from Deforestation and Forest Degradation*. Ministry of the Environment, Oslo, Norway.
- Gumbricht, T., Román-Cuesta, R.M., Verchot, L.V., Herold, M., Wittmann, F., Householder, E., Herold, N., and Murdiyarto, D. (2017). *Tropical and Subtropical Wetlands Distribution version 2*. Version V3. DOI: [10.17528/CIFOR/DATA.00058](https://doi.org/10.17528/CIFOR/DATA.00058). URL: <https://doi.org/10.17528/CIFOR/DATA.00058>.
- Hansen, M. C., Stehman, S. V., and Potapov, P. V. (2010). “Quantification of global gross forest cover loss”. *Proceedings of the National Academy of Sciences* 107.19, 8650—8655.
- Hansen, M. C. *et al.* (2013). “High-Resolution Global Maps of 21st-Century Forest Cover Change”. *Science* 342.6160, pp. 850–853.
- Heckman, James J., Ichimura, Hidehiko, and Todd, Petra (1997). “Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme”. *Review of Economic Studies* 64, pp. 605–654.
- Hänsel, Martin C., Drupp, Moritz A., Johansson, Daniel J. A., Nesje, Frikk, Azar, Christian, Freeman, Mark C., Groom, Ben, and Sterner, Thomas (2020). “Climate economics support for the UN climate targets”. *Nature Climate Change* 10, pp. 781–789.
- IPCC (2014). *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Tech. rep. Geneva, Switzerland: IPCC.
- Iacus, Stefano M, King, Gary, and Porro, Giuseppe (2009). “CEM: Software for Coarsened Exact Matching”. *Journal of Statistical Software* 30.9.
- Iacus, Stefano M, King, Gary, and Porro, Giuseppe (2012). “Causal inference without balance checking: Coarsened Exact Matching”. *Political analysis* 20.1, pp. 1–24.
- Imbens, Guido (2014). *Matching Methods in Practice: Three Examples*. Working Paper 19959. National Bureau of Economic Research.

- Imbens, Guido and Kalyanaraman, Karthik (2012). “Optimal Bandwidth Choice for the Regression Discontinuity Estimator”. *The Review of Economic Studies* 79.3, pp. 933–959.
- Irawan, S., Tacconi, L., and Ring, I. (2013). “Designing intergovernmental fiscal transfers for conservation: The case of REDD+ revenue distribution to local governments in Indonesia”. *Land Use Policy* 36, 47–59.
- Jakarta Post (2019). *5-step plan to protect our forests*. URL: <https://www.thejakartapost.com/academia/2019/06/03/5-step-plan-to-protect-our-forests.html>.
- Joppa, Lucas and Pfaff, Alexander (2010). “Reassessing the forest impacts of protection”. *Annals of the New York Academy of Sciences* 1185.1, pp. 135–149.
- King, Gary and Nielsen, Richard (2019). “Why propensity scores should not be used for matching”. *Political Analysis* 27.4, pp. 435–454.
- King, Gary and Zeng, Langche (2006). “The dangers of extreme counterfactuals”. *Political Analysis* 14.2, pp. 131–159.
- Krisnawati, Haruni, Kallio, Maarit, and Kanninen, Markku (2011). *Acacia Mangium Willd.: ecology, silviculture and productivity*. Tech. rep. Bogor, Indonesia: CIFOR.
- LTS International (2018). *Third independent review of the Indonesia-Norway cooperation on reducing greenhouse gas emissions from REDD+*.
- Lee, David S. and Lemieux, Thomas (2010). “Regression discontinuity designs in economics”. *Journal of Economic Literature* 48.2, pp. 281–355.
- Lee, Myoung-jae (2005). *Micro-Econometrics for Policy, Program and Treatment Effects*. Advanced Texts in Econometrics. Oxford: Oxford University Press. ISBN: 9780199267699.
- Maniatis, Danae, Scriven, Joel, Jonckheere, Inge, Laughlin, Jennifer, and Todd, Kimberly (2019). “Toward REDD+ implementation”. *Annual Review of Environment and Resources* 44, pp. 373–398.
- Margono, Belinda Arunarwati, Potapov, Peter V., Turubanova, Svetlana, Stolle, Fred, and Hansen, Matthew C. (2014). “Primary forest cover loss in Indonesia over 2000–2012”. *Nature Climate Change* 4, pp. 730–735.
- Marlier, Miriam E, Defries, Ruth, Pennington, Derric, and Nelson, Erik (2015). “Future fire emissions associated with projected land use change in Sumatra”. *Global Change Biology* 21, pp. 345–362.
- Meehan, Fiona and Tacconi, Luca (2017). “A framework to assess the impacts of corruption on forests and prioritize responses”. *Land Use Policy* 60, pp. 113–122.
- Meyfroidt, Patrick and Lambin, Eric F. (2009). “Forest transition in Vietnam and displacement of deforestation abroad”. *Proceedings of the National Academy of Sciences* 106, 16139–16144.
- Meyfroidt, Patrick, Rudel, Thomas K., and Lambin, Eric F. (2010). “Forest transitions, trade, and the global displacement of land use”. *Proceedings of the National Academy of Sciences* 107, 20917–20922.
- Miettinen, Jukka, Shi, Chenghua, and Liew, Soo Chin (2011). “Deforestation rates in insular Southeast Asia between 2000 and 2010”. *Global Change Biology* 17, pp. 2261–2270.

- Ministry of Forestry of the Republic of Indonesia (2011). *Indicative Map of New Revised Permit Delays 5*. URL: <http://appgis.dephut.go.id/appgis/download.aspx>.
- Ministry of Forestry (2014). *Ministerial Decree SK 5984/Menhut-VI/BRPUK/2014*. URL: [www.dephut.go.id/index.php/news/details/9671](http://www.dephut.go.id/index.php/news/details/9671) (accessed July 2017).
- MoEF (2019). *Emissions Reduction Report for the Indonesia-Norway Partnership*. government report. Directorate General of Climate Change, Ministry of Environment and Forestry, Republic of Indonesia.
- Mongabay (2020a). *Experts question integrity of Indonesia’s claim of avoided deforestation*. URL: <https://news.mongabay.com/2020/09/green-climate-fund-indonesia-redd-deforestation>.
- Mongabay (2021a). *Governments, companies pledge \$1 billion for tropical forests*. URL: <https://news.mongabay.com/2021/04/governments-companies-pledge-1-billion-for-tropical-forests/>.
- Mongabay (2019a). *Indonesia forest-clearing ban is made permanent, but labeled “propaganda”*. URL: <https://news.mongabay.com/2019/08/indonesia-forest-clearing-ban-is-made-permanent-but-labeled-propaganda/>.
- Mongabay (2021b). *Indonesia terminates agreement with Norway on \$1b REDD+ scheme*. URL: <https://news.mongabay.com/2021/09/indonesia-terminates-agreement-with-norway-on-1b-redd-scheme/>.
- Mongabay (2019b). *Indonesia to get first payment from Norway under USD1bn REDD+ scheme*. URL: <https://news.mongabay.com/2019/02/indonesia-to-get-first-payment-from-norway-under-1b-redd-scheme/>.
- Mongabay (2020b). *Indonesia to receive \$56m payment from Norway for reducing deforestation*. URL: <https://news.mongabay.com/2020/05/indonesia-norway-redd-payment-deforestation-carbon-emission-climate-change/>.
- Myers, R., Intarini, D., Sirait, M.T., and Maryudi, A. (2017). “Claiming the forest: Inclusions and exclusions under Indonesia’s ‘new’ forest policies on customary forests”. *Land Use Policy* 66, 205—213.
- Myers, R., Larson, A.M., Ravikumar, A., Kowler, L.F., Yang, A., and Trench, T. (2018). “Messiness of forest governance: How technical approaches suppress politics in REDD+ and conservation projects”. *Global Environmental Change* 50, pp. 314—324.
- Nelson, A. (2008). *Travel time to major cities: A global map of accessibility*. Tech. rep. DOI: [10.2788/95835](https://doi.org/10.2788/95835).
- Nelson, Andrew and Chomitz, Kenneth M. (2011). “Effectiveness of strict vs. multiple use protected areas in reducing tropical forest fires: A global analysis using matching methods”. *PLoS ONE* 6.8.
- Norman, Marigold and Nakhoda, Smita (2015). *The State of REDD+ Finance*. Tech. rep. 378. Center for Global Development, Washington D.C.
- Nurfatriani, F., Darusman, D., Nurrochmat, D.R., Yustika, A.E., and Muttaqin, M.Z. (2015). “Redesigning Indonesian forest fiscal policy to support forest conservation”. *Forest Policy and Economics* 61, 39—50.
- Ostwald, Madelene and Henders, Sabine (2014). “Making two parallel land-use sector debates meet: Carbon leakage and indirect land-use change”. *Land Use Policy* 36, 533—554.



- Palmer, Charles (2005). “The Nature of Corruption in Forest Management”. *World Economics* 6.2, pp. 1–10.
- Purnomo, Agus (2012). *Protecting Our Forests, Pros-Cons of Moratorium on Forests and Peatlands’ Policy*. Jakarta, Indonesia: Kepustakaan Populer Gramedia.
- Roopsind, Anand, Sohngen, Brent, and Brandt, Jodi (2019). “Evidence that a national REDD+ program reduces tree cover loss and carbon emissions in a high forest cover, low deforestation country”. *Proceedings of the National Academy of Sciences* 116, pp. 24492–24499.
- Rose, Amy N., McKee, Jacob J., Sims, Kelly M., Bright, Edward A., Reith, Andrew E., and Urban, Marie L. (2020). *LandScan 2019*.
- Ryan, Andrew M., Evangelos, Kontopantelis, Ariel, Linden, and Burgess, James F. (2018). “Now trending: Coping with non-parallel trends in difference-in-differences analysis”. *Statistical Methods in Medical Research* 28.12, pp. 3697–3711.
- Sahide, M.A.K., Maryudi, A., Supratman, S., and Giessen, L. (2016). “Is Indonesia utilising its international partners? The driving forces behind Forest Management Units”. *Forest Policy and Economics* 69, 11–20.
- Sekhon, Jasjeet S. (2011). “Multivariate and propensity score matching software with automated balance optimization: The matching package for R”. *Journal of statistical software* 42.7.
- Stern, Nicholas (2006). *The Economics of Climate Change: The Stern Review*. Her Majesty’s Treasury of the UK Government.
- Tacconi, Luca, Rodrigues, Rafael J, and Maryudi, Ahmad (2019). “Law enforcement and deforestation: Lessons for Indonesia from Brazil”. *Forest Policy and Economics* 108, p. 101943.
- Turubanova, S., Popapov, P. V., Tyukavina, A., and Hansen, M. C. (2018). “Ongoing primary forest loss in Brazil, Democratic Republic of the Congo, and Indonesia”. *Environmental Research Letters* 13, p. 074028.
- UNFCCC (2008). *UNITED Report of the Conference of the Parties on its thirteenth session, held in Bali from 3 to 15 December 2007 Part Two : Action taken by the Conference of the Parties at its thirteenth session Decisions adopted by the Conference of the Parties*. Tech. rep., pp. 1–60.
- Varkkey, Helena (2013). “Patronage politics, plantation fires and transboundary haze”. *Environmental Hazards* 12, pp. 37–41.
- West, Thales A. P., Borner, Jan, Sills, Erin O., and Kontoleon, Andreas (2020). “Overstated carbon emission reductions from voluntary REDD+ projects in the Brazilian Amazon”. *Proceedings of the National Academy of Sciences* 117, pp. 24188–24194.
- Wijaya, A., Chrysolite, H., Ge, M., Wibowo, C.K., Pradana A. Utami, A.F., and Austin, K. (2017). *How Can Indonesia Achieve its Climate Change Mitigation Goal? An Analysis of Potential Emissions Reductions from Energy and Land-Use Policies*. Tech. rep. World Resources Institute (Jakarta). URL: [www.wri.org/publication/how-can-indonesia-achieve-its-climate-goal](http://www.wri.org/publication/how-can-indonesia-achieve-its-climate-goal).
- Wooldridge, Jeffrey M. (2010). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.
- WorldPop (2018). *WorldPop - School of Geography and Environmental Science, University of Southampton; Department of Geography and Geosciences, University of Louisville; Departement de Geographie, Universite de Namur and Center for*

*International Earth Science Information Network (CIESIN), Columbia University (2018). Global High Resolution Population Denominators Project - Funded by The Bill and Melinda Gates Foundation (OPP1134076).* DOI: <https://dx.doi.org/10.5258/SOTON/WP00644>.

Zarin, D. J. *et al.* (2016). “Can carbon emissions from tropical deforestation drop by 50% in 5 years?” *Global Change Biology* 22.4, 1336?1347. DOI: <https://doi.org/10.1111/gcb.13153>.

## 4.A Background to the Moratorium

### 4.A.1 The Indonesia-Norway REDD+ partnership

On 26 May 2010, a Letter of Intent (LoI) was signed by the governments of Indonesia and Norway, establishing the Indonesia-Norway REDD+ partnership. As a part of this partnership, Norway pledged to provide US\$1 billion to Indonesia in exchange for verified reductions in GHG emissions from deforestation. Three phases were specified in the LoI. The first aimed to establish the necessary institutions and capacity, the second to transform managerial systems while verified emission reductions were to be delivered in the third. Based on the LoI, Indonesia was supposed to be on track to receive its first result-based payment in 2014 but the transition to phase three was delayed due to various political and institutional changes in Indonesia ([Mongabay, 2020](#)).

To operationalise the three phases, various REDD+ Task Forces were appointed between 2010 and 2013 by the then President of Indonesia, Susilo Bambang Yudhoyono, alongside a ministerial-level REDD+ Management Agency (BP REDD+) that was created in 2013. In 2015, President Joko Widodo, who was elected in 2014, dissolved BP REDD+ and a new Ministry of Environment and Forestry (KLHK) was formed from the Ministry of Environment and the Ministry of Forestry. This new ministry absorbed the National Council on Climate Change (DNPI) and BP REDD+. Thus, the responsibility for delivering GHG emission reductions fell to KLHK, the capacity of which has been supported through cooperation with Kemitraan (the Partnership for Governance Reform in Indonesia) and the World Resources Institute Indonesia (WRI).

Further delaying Indonesia's transition into phase three was the long time that was needed to create an integrated measurement, reporting and verification (MRV) system sufficient to account for progress in reducing emissions from deforestation ([LTS International, 2018](#)). As of 2018, there was no MRV system in place. A key reason for the long gestation of this system was the context and purpose of national MRV capacity, which has evolved since the LoI was signed in 2010. Specifically, a divergence arose between the commitment to economy-wide emission reductions in line with national policy, coordinated by the national planning agency (Bappenas), and emissions reductions commitments as part of Indonesia's Nationally Determined Contribution (NDC). The latter is led by KLHK but is coordinated with other agencies and sub-national entities that regulate or impact forest and peatland. In addition, there were and continue to be uncertainties regarding forest definitions, boundaries and baselines.

Similar to the MRV system, the Indonesia-Norway partnership has had to create and develop other, critical institutional arrangements off the back of relatively weak, pre-existing arrangements. The most-recent review of the partnership, published in 2018 ([LTS International, 2018](#)), makes clear that forest law enforcement remains patchy at best. The partnership has been financing various capacity-building initiatives in KLHK. For example, it has been developing capacity to detect networks of


individuals, holding companies and licensees who may have violated the Moratorium, to track prosecutions, recover assets and detect corruption. The partnership has also been involved in efforts to strengthen civil society capacity in the use of high technology and fieldwork to support independent monitoring and reporting on forests. To manage financial flows linked to REDD+, including the first payment announced in 2019 for emissions reductions in 2017 (see below), the Indonesian Environmental Estate Fund (BPDLH) has been established, overseen by the Ministry of Finance.

As of 2018, there has been relatively little progress on the implementation of sub-national REDD+ strategies and the institutional arrangements of the Indonesia-Norway REDD+ partnership remained highly centralised. Yet, at the local level there are various planning and mapping initiatives in, for example, Central, West and North Kalimantan, Aceh, South Sumatra and Riau. Central Kalimantan has a jurisdictional REDD+ initiative. There is, however, a lack of coordination of these sub-national efforts and little clarity on how REDD+ funds can be used in support of such efforts, as well as investment in the capacity of sub-national governments to, e.g. enforce the Moratorium.

In 2021, the Indonesian government formally and unexpectedly terminated its REDD+ partnership with Norway ([Mongabay, 2021](#)). One reason cited is a delay in the delivery of the payment from Norway, although there are also reports of disagreements between the two countries regarding the conditions required by Norway for the payment to be disbursed (e.g. [Climate Home News, 2021](#)).

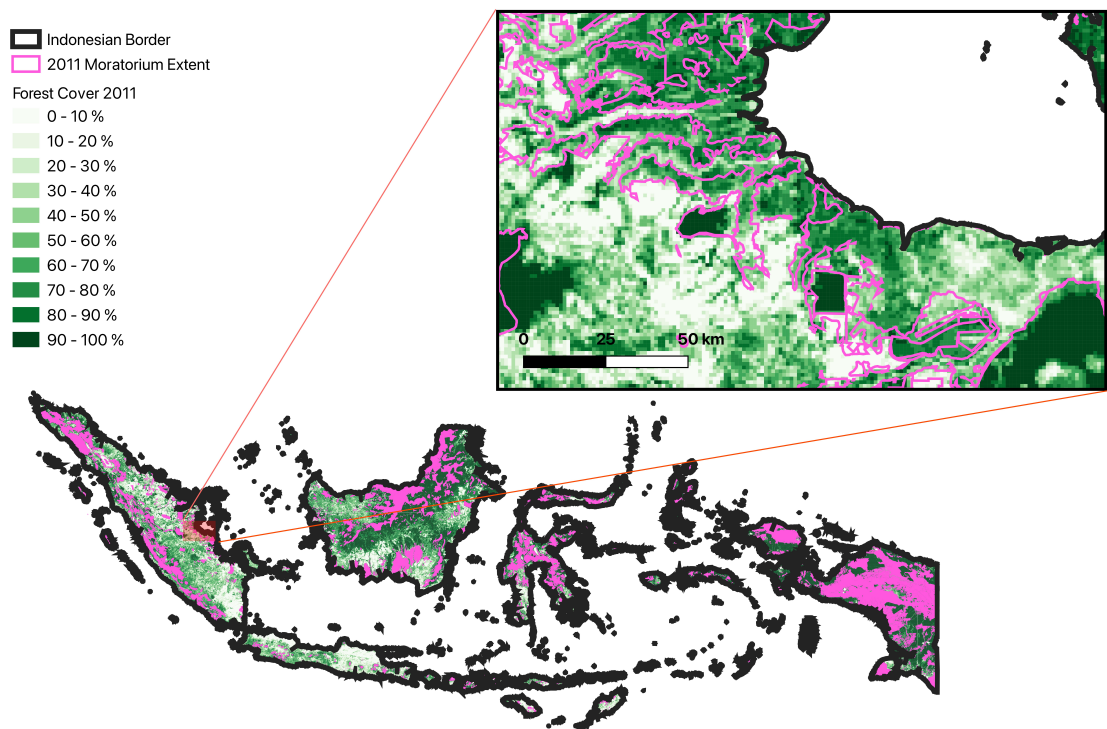
#### **4.A.2 The 2011 Moratorium on palm oil, timber and logging concessions**

The centrepiece of the Indonesia-Norway REDD+ partnership was a Moratorium on new palm oil, timber and logging concession licences, instigated in 2011. Such licences were issued by district governments, the third administrative tier (after provincial governments), and conferred certain rights to firms. Palm oil concessions refer to areas where industrial-scale oil palm plantations are established. A logging concession refers to a forest area for managing and extracting timber, and are distinct from timber concessions, where plantation forests are established for the production of pulp and paper. After the timber has been extracted from logging concessions, many of these concessions are re-designated as palm oil or timber concessions ([Forest Watch Indonesia, 2015](#)). [Figure 4.A.1](#) shows the district-level distribution of all pre-2011 concessions, by share of the total area of the country's land area designated as concessions before 2011.



**Figure 4.A.1:** Share of pre-2011 palm oil, timber and logging concessions (combined) by district.

Around 3% of land designated as concessions pre-2011 is located in forestland within the Moratorium's 2011 boundaries, which are shown in [Figure 4.A.2](#). Covering a substantial proportion of Indonesia's primary forest estate, the 2011 Moratorium was renewed every two years before being made permanent in 2019 (Presidential Instruction No. 10/2011, extended under No. 6/2013, then No 8/2015, then No. 6/2017 regarding a Moratorium on the Granting of New Licences and Improvement of Natural Primary Forest and Peatland Governance). The Moratorium boundaries established in 2011 have been periodically updated since. Initially covering 69 million hectares, the total area under the Moratorium has fluctuated since 2011, declining to about 64 million ha before rising to 66 million by 2018. These fluctuations were due to new spatial mapping inputs and surveys that changed the classification or designation of land ([LTS International, 2018](#)).



**Figure 4.A.2:** Extent of Moratorium in Indonesia, 2011. The 2011 Moratorium boundaries denote the treatment areas for our empirical analysis, and are assumed constant throughout our treatment periods. Sources: Authors' elaboration on forest cover data (Hansen *et al.*, 2013). Moratorium shapefile from Ministry of Forestry of the Republic of Indonesia (2011).

Some of the fluctuations in the Moratorium's boundaries were likely to have occurred due to changes in land classification that removed forestland from Moratorium areas to the benefit of concessionaires seeking new concession licences. Although the Moratorium is administered by the national government, the Moratorium map is revised by district governments, which then notify the national government of any changes (Mongabay, 2019). Once approved, areas of forest protected by the Moratorium are, in effect, dropped from the map. In particular, it has been documented that forestland has been re-designated, by district governments, from "official forest" to "non-institutionally recognized" forest, or land with forest cover that is designated as land "for other use" (*Areal Penggunaan Lain*, APL) (Indrarto *et al.*, 2012; Enrici and Hubacek, 2016). Areas of forest re-designated as APL have been shown to be highly susceptible to higher deforestation and degradation pressures (Margono *et al.*, 2012), because designation as APL land makes it easier to obtain concession licenses in such areas (Enrici and Hubacek, 2016).

According to the FAO, the total amount of APL forest land was around 8.5-8.6 million ha in 2009 and 2012 (FAO, 2015). It is difficult to precisely assess the extent to which the stock of APL forest land changed due to the addition of newly-designated APL land net of APL land being re-designated as new concessions, although Greenpeace has reported that, on the basis of an analysis of satellite imagery, up to 4.5 million ha of forest have been removed from the 2011 Moratorium map since 2011, of which at least 1.6 million ha was subsequently licenced out to concessionaires (Mongabay, 2019).

Concession area data, collected by Greenpeace circa 2010 and reported in Abood *et al.* (2015), suggest that logging, timber and palm oil concessions accounted for approximately 24, 11 and 12 million ha of Indonesia's land, respectively. When land is allocated to a concessionaire, it is important to distinguish between the area of land designated as a concession and the area of land that has been cleared and, in the context of palm oil and timber concessions, planted. Indonesian government data, reported by Forest Watch Indonesia to the European Union in 2015 (Forest Watch Indonesia, 2015), suggests that total timber concession areas increased from around 8-9 million ha in 2009-10 to around 10 million ha in 2011-13, of which 50% was planted. Palm oil concession areas also increased, from around 8 million ha in 2009-10 to 9-10 million ha in 2011-12, before apparently more than doubling to over 22 million ha by 2018 (Steinweg *et al.*, 2019). How much of this area was planted was unreported. Forest clearance in concessions has been well-documented, accounting for around half of Indonesia's aggregate forest loss (Abood *et al.*, 2015; Austin *et al.*, 2019).

In our empirical analysis, we retain the 2011 Moratorium boundaries for estimation of the policy treatment effect and ignore subsequent changes to these boundaries because of evidence that the Moratorium may have driven incentives to re-designate land protected by the Moratorium into APL land before being re-designated as concessions. Although not technically breaking the rules of the Moratorium, we argue that this behaviour was not in the spirit of the Moratorium and indeed occurred



in response to the Moratorium. Thus, any cells containing forestland within these 2011 boundaries that dropped out of the Moratorium map after 2011 are retained in our treatment group.

Our treatment period, on the other hand, reflects changes in the sectors affected by the Moratorium. Forest fires during the El Niño of 2015 damaged around 2.6 million ha of forest, releasing a billion tonnes of carbon dioxide equivalent (Glauber *et al.*, 2016). This led to new and enhanced regulations on peatland, a new Peatland Restoration Agency (BRG) and increased policy priority for fire prevention, and the expansion of the Moratorium to peatland forest to the whole country (Alisjahbana and Busch, 2017). All three types of concession were similarly affected. Thus, the treatment effect of the Moratorium on peatland forest, based on the 2011 boundaries, is estimated for all periods between 2011 and 2017 while the effect on dryland forest is also estimated for the period 2011-18. In September 2018, a new three-year Moratorium targeting new palm oil plantations across the whole country began but this is not considered in our analysis due to coming into effect at the end of our study period.

#### 4.A.3 How the Indonesia-Norway partnership estimated forest-based emissions reductions in 2017

In 2017, Indonesia reportedly reduced emissions from deforestation and forest degradation by 11.2 MtCO<sub>2</sub>-eq, the first reduction since the start of the Moratorium (Mongabay, 2020). Norway subsequently announced that it would pay Indonesia US\$ 56.2 million. In this sub-section, we summarise how these emissions reductions were estimated, drawing on a report produced by Indonesia's Ministry of Environment and Forestry (MoEF, 2019). This report was submitted to Norway's government and provides the evidence base for estimating payments from Norway's government to Indonesia's government in the event of reductions of GHG emissions from deforestation and forest degradation, for the period 2017-2020.

The reference emission level adopted as a baseline for estimating changes in emissions (in tons of carbon dioxide equivalent per year) and hence, Indonesia's REDD+ performance under the framework of Norway-Indonesia Partnership, is the Result Based Payment (RBP) baseline. To be updated periodically every five years, the RBP baseline has a different reference period (2006-2016) from Indonesia's Forest Reference Emission Level (FREL) that was submitted to the UNFCCC in 2016 (1990-2012).

The RBP baseline covers all "natural forests" in Indonesia, defined as "a land area of more than 6.25 hectares with trees higher than 5 meters at maturity and a canopy cover of more than 30 percent". These forests are classified according to land cover maps produced by the National Forest Monitoring System (NFMS) and used by the Ministry of Environment and Forestry. There are six natural forest classes that are used to estimate the RBP: primary dryland forest, secondary dryland forest, primary swamp forest, secondary swamp forest, primary mangrove forest, and secondary mangrove forest. Map data for the years 2006, 2009, 2011, 2012, 2013, 2014, 2015 and 2016 are used to analyse historical land cover changes and calculate changes in

emissions. Forest data were sourced from Landsat 5 Thematic Mapper (TM) and Landsat 7 Enhanced Thematic Mapper Plus (ETM+).

Not all natural forests in Indonesia are included in the estimation of the RSB (or the FREL) and although it is clear that land-use designation is a critical factor determining which forest areas are included or excluded, there is little information on precisely how designation has affected these estimates. In 2016, a forest estate of 95.2 million ha was designated according to “permanent forest” (87.7 million ha) and “forest area designated for conversion to other land uses” (7.5 million ha; APL - see above) (Government of Indonesia, 2018). Permanent forest includes “protection forest” (24.0 million ha), “conservation/fully protected forest and recreation forests” (17.4 million ha), and “production forest” (46.2 million ha). The last, the biggest forest category by area, contains forests designated for the exploitation of timber and other forest products. According to Indonesia’s FREL submission to the UNFCCC in 2016, up to 20 million ha of permanent forest (specifically that which, by law, is slated for conversion) and APL forest is excluded from the estimation of Indonesia’s FREL (Government of Indonesia, 2016). Further inference is not possible due to the focus of estimates of the RSB and the FREL on the levels of deforestation rather than on levels of forest cover.

For both the RSB and FREL, annual forest cover change was estimated by overlaying land cover maps of two subsequent periods. Deforestation is defined as the change of natural forests from a forest class to a non-forest class at a given location in the period 2006-2016, while degradation occurs when a change from a primary to a secondary forest class is identified. Calculation of the RBP baseline used emission factors for each forest class and are identical to those used in Indonesia’s first FREL. These emission factors were estimated using data primarily collected by the National Forest Inventory (NFI). Average annual emissions from deforestation and forest degradation in the period 2006-2016 were estimated to be 236.9 MtCO<sub>2</sub>-eq/yr and 41.6 MtCO<sub>2</sub>-eq/yr, respectively. Actual emissions in 2017 were estimated to be 228.3 MtCO<sub>2</sub>-eq/yr and 42.7 MtCO<sub>2</sub>-eq/yr from deforestation and forest degradation, respectively.

Against the 2006-2016 RBP baseline, Indonesia reduced emissions in 2017 by 7.4 MtCO<sub>2</sub>-eq/yr: 8.6 MtCO<sub>2</sub>-eq/yr from avoided deforestation (3.6% below the baseline) net of an increase in emissions from forest degradation, 1.2 MtCO<sub>2</sub>-eq/yr (2.9% above the baseline). Uncertainties related to the forest data and emissions factors used were calculated following IPCC guidelines (2006). To account for statistical uncertainty, the risk of carbon reversals and leakage, a “set aside factor” of 35% was applied to the estimates of emissions reductions, giving a net reduction of 4.8 MtCO<sub>2</sub>-eq/yr.

Also estimated were emissions from changes in peatland (swamp) forest, both from peat decomposition and peat fires. The former was defined as the changing process of peat form as a result of a decline in water levels caused by deforestation and forest degradation, and land utilization. When peat is drained and exposed, it oxidises thus

generating carbon dioxide. Emission factors for peat decomposition also followed IPCC guidelines, based on the assumption that all utilized areas are drained. The calculation of emissions from peat decomposition was otherwise identical to the method used for calculating emissions from deforestation and forest degradation but with the inclusion of ‘inherited’ emissions from peat decomposition. In 2017, 256.7 MtCO<sub>2</sub>-eq/yr was emitted from the decomposition of peat, which when compared to the RBP baseline of 260.6 MtCO<sub>2</sub>-eq/yr, generated an emissions reduction of 3.9 MtCO<sub>2</sub>-eq/yr.

Emissions from peat fires were estimated using data on burn scar areas. Fire incidence in peatland was 13,555 hectares in 2017. Combined with estimates of the mass of fuel available for combustion, to generate emission factors for burned peat, this amounted to 12.5 MtCO<sub>2</sub>-eq/yr. Against an annual emission baseline of 249 MtCO<sub>2</sub>-eq/yr between 2006 and 2016, there was an estimated emission reduction of 236.4 MtCO<sub>2</sub>-eq/yr, in 2017. In contrast to estimates of emissions from deforestation and forest degradation, these estimates of peat fire emissions were subject to considerably more uncertainty, which was not possible to quantify.

Perhaps due to this uncertainty, Indonesia’s government initially excluded emissions from peat fires in its calculations. However, Norway’s government factored in peat fires, coming up with a total emissions reduction estimate of 11.2 MtCO<sub>2</sub>-eq/yr for 2017 ([Mongabay, 2020](#)). Multiplied by a carbon price of US\$5 per ton, set by Norway, this generated an estimated payment of \$US 56.2 million to be paid by Norway to Indonesia for emissions reductions from deforestation and forest degradation in 2017.

## 4.B Main and Additional Results

### 4.B.1 ATT results: year-on-year

The ATT estimates in the following tables are our main results, which are used to derive estimates of avoided forest cover loss and carbon emissions reductions.

**Table 4.B.1:** DD dryland forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	0.137***	0.0380	3.641	0.000	160012	152118	7894
<b>2013</b>	0.237***	0.0510	4.633	0.000	160012	152118	7894
<b>2014</b>	0.318***	0.0670	4.762	0.000	160012	152118	7894
<b>2015</b>	0.472***	0.0780	6.037	0.000	160012	152118	7894
<b>2016</b>	0.718***	0.0940	7.674	0.000	160012	152118	7894
<b>2017</b>	0.813***	0.0990	8.186	0.000	160012	152118	7894
<b>2018</b>	0.938***	0.1050	8.892	0.000	160012	152118	7894

*Notes:* Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.2:** DDD dryland forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	0.108***	0.0390	2.761	0.006	160012	152118	7894
<b>2013</b>	0.178***	0.0570	3.146	0.002	160012	152118	7894
<b>2014</b>	0.229***	0.0760	2.998	0.003	160012	152118	7894
<b>2015</b>	0.354***	0.0940	3.780	0.000	160012	152118	7894
<b>2016</b>	0.571***	0.1140	4.991	0.000	160012	152118	7894
<b>2017</b>	0.637***	0.1280	4.984	0.000	160012	152118	7894
<b>2018</b>	0.732***	0.1420	5.150	0.000	160012	152118	7894

*Notes:* Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.3:** DD peatland forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	-0.016	0.0390	-0.414	0.679	95154	91024	4130
<b>2013</b>	-0.052	0.0530	-0.979	0.328	95154	91024	4130
<b>2014</b>	-0.153**	0.0710	-2.163	0.031	95154	91024	4130
<b>2015</b>	-0.073	0.0820	-0.883	0.377	95154	91024	4130
<b>2016</b>	-0.032	0.0960	-0.333	0.739	95154	91024	4130
<b>2017</b>	-0.047	0.1010	-0.461	0.645	95154	91024	4130

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.4:** DDD peatland forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	-0.011	0.0410	-0.260	0.795	95154	91024	4130
<b>2013</b>	-0.041	0.0600	-0.680	0.496	95154	91024	4130
<b>2014</b>	-0.136*	0.0830	-1.651	0.099	95154	91024	4130
<b>2015</b>	-0.05	0.1010	-0.497	0.619	95154	91024	4130
<b>2016</b>	-0.004	0.1220	-0.033	0.974	95154	91024	4130
<b>2017</b>	-0.013	0.1370	-0.096	0.923	95154	91024	4130

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

#### 4.B.2 Estimated avoided forest cover loss and carbon emissions reductions

The estimates of avoided forest cover loss and carbon emissions reductions in the following tables can be seen graphically in [Figure 4.A.2](#).

**Table 4.B.5:** DD dryland forest avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	12.723	21.967
2013	21.993	37.972
2014	29.441	50.830
2015	43.789	75.603
2016	66.550	114.901
2017	75.393	130.168
2018	86.931	150.089

**Table 4.B.6:** DDD dryland forest avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	9.990	17.248
2013	16.526	28.533
2014	21.240	36.672
2015	32.855	56.725
2016	52.883	91.303
2017	58.992	101.851
2018	67.797	117.053

**Table 4.B.7:** DD peatland forest avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	-0.677	-1.541
2013	-2.165	-4.929
2014	-6.403	-14.578
2015	-3.041	-6.924
2016	-1.336	-3.042
2017	-1.954	-4.449

**Table 4.B.8:** DDD peatland forest avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	-0.443	-1.009
2013	-1.698	-3.866
2014	-5.703	-12.983
2015	-2.107	-4.798
2016	-0.169	-0.384
2017	-0.553	-1.259

### 4.B.3 Robustness checks (2017 endpoint)

The following tables show the results from testing the robustness of our main results.

**Table 4.B.9:** DD dryland forest trimmed sample results (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	0.813***	0.0990	8.186	0.000	160012	152118	7894
<b>30%</b>	0.659***	0.1090	6.055	0.000	144501	136472	8029
<b>60%</b>	0.79***	0.1210	6.558	0.000	127375	119284	8091

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.10:** DDD dryland forest trimmed sample results (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	0.637***	0.1280	4.984	0.000	160012	152118	7894
<b>30%</b>	0.443***	0.1390	3.193	0.001	144501	136472	8029
<b>60%</b>	0.593***	0.1500	3.958	0.000	127375	119284	8091

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.11:** DD peatland forest trimmed sample results (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	-0.047	0.1010	-0.461	0.645	95154	91024	4130
<b>30%</b>	-0.097	0.1110	-0.876	0.381	86897	82832	4065
<b>60%</b>	-0.186	0.1220	-1.532	0.126	77777	73554	4223

*Notes:* Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.12:** DDD peatland forest trimmed sample results (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	-0.013	0.1370	-0.096	0.923	95154	91024	4130
<b>30%</b>	0.025	0.1480	0.169	0.866	86897	82832	4065
<b>60%</b>	-0.121	0.1580	-0.768	0.442	77777	73554	4223

*Notes:* Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .



**Table 4.B.13:** Coarsened Exact Matching dryland forest results (2017 endpoint)

	SATT	S.E.	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Model 1</b>	0.4868***	0.06	8.12	0.000	160012	125458	34554
<b>Model 2</b>	0.4786***	0.06	7.56	0.000	160012	141334	18678
<b>Model 3</b>	0.4664***	0.06	7.38	0.000	160012	141362	18650
<b>Model 4</b>	0.5221***	0.06	8.27	0.000	160012	141362	18650

Notes: \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

*Model 1:* 2005-2010 forest cover, slope, elevation and distance from roads included in the matching algorithm.

Distance from roads uses manually defined bins in lieu of the automated ones.

Distance from palm oil, timber and logging concessions are included in the linear outcome model.

*Model 2:* 2005-2010 forest cover, slope, elevation, 2005 and 2010 population included in the matching algorithm.

Distance from roads and carbon stocks are included in the linear outcome model.

*Model 3:* 2005-2010 forest cover, slope, elevation and distance from cities included in the matching algorithm.

Distance from palm oil, timber and logging concessions, distance from roads, carbon stocks and

2000, 2005 and 2010 population are included in the linear outcome model.

*Model 4:* 2005-2010 forest cover, slope, and elevation included in the matching algorithm.

Distance from roads, carbon stocks, and 2000, 2005 and 2010 population

are included in the linear outcome model.

**Table 4.B.14:** Coarsened Exact Matching peatland forest results (2017 endpoint)

	SATT	S.E.	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Model 1</b>	-0.012	0.08	-0.15	0.884	95154	71378	23776
<b>Model 2</b>	0.2843***	0.08	3.62	0.000	95154	78322	16832
<b>Model 3</b>	-0.2327*	0.09	-2.50	0.012	95154	56523	38631
<b>Model 4</b>	0.1185	0.08	1.50	0.135	95154	80365	14789

Notes: \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

*Model 1:* 2005-2010 forest cover, slope, elevation, 2005 and 2010 population. Moreover

Distance from roads, AGC and distance from palm oil use manually defined bins in lieu of the automated ones.

*Model 2:* 2005-2010 forest cover, slope, elevation, 2005 and 2010 population. Moreover

Distance from roads and AGC use manually defined bins in lieu of the automated ones.

*Model 3:* 2005-2010 forest cover, slope, elevation, 2005 and 2010 population. Moreover

Distance from roads and AGC use manually defined bins in lieu of the automated ones.

Distance from palm oil, timber and logging use manually defined bins in lieu of the automated ones.

*Model 4:* 2005-2010 forest cover, slope, elevation and distance from roads included in the matching algorithm.

Distance from palm oil, timber and logging concessions are included in the linear outcome model.

Distance from roads uses manually defined bins in lieu of the automated ones.

**Table 4.B.15:** DDD dryland forest ATT: caliper = 0.001

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	0.115**	0.0450	2.568	0.010	160012	159975	37
<b>2013</b>	0.176***	0.0650	2.724	0.006	160012	159975	37
<b>2014</b>	0.230***	0.0870	2.652	0.008	160012	159975	37
<b>2015</b>	0.373***	0.1070	3.492	0.000	160012	159975	37
<b>2016</b>	0.583***	0.1300	4.489	0.000	160012	159975	37
<b>2017</b>	0.646***	0.1450	4.454	0.000	160012	159975	37
<b>2018</b>	0.749***	0.1610	4.649	0.000	160012	159975	37

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < .01$ .

**Table 4.B.16:** DDD peatland forest ATT: caliper = 0.001

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	-0.008	0.0430	-0.194	0.846	95154	95004	150
<b>2013</b>	-0.035	0.0640	-0.550	0.582	95154	95004	150
<b>2014</b>	-0.121	0.0880	-1.374	0.170	95154	95004	150
<b>2015</b>	-0.047	0.1080	-0.440	0.660	95154	95004	150
<b>2016</b>	-0.013	0.1300	-0.100	0.920	95154	95004	150
<b>2017</b>	-0.021	0.1460	-0.147	0.883	95154	95004	150

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.17:** DDD dryland: robustness (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>1:2 Matching</b>	0.553***	0.1020	5.446	0.000	160012	138635	21377
<b>1:3 Matching</b>	0.54***	0.0850	6.346	0.000	160012	123432	36580
<b>1:5 Matching</b>	0.555***	0.0640	8.663	0.000	160012	95821	64191
<b>Above 1000 m</b>	0.223	0.1870	1.192	0.233	67693	49782	17911
<b>All elevation</b>	0.611***	0.1210	5.060	0.000	227705	224112	3593

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .


**Table 4.B.18:** DDD peatland: robustness (2017 endpoint)

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>1:2 Matching</b>	-0.022	0.1160	-0.191	0.849	95154	82953	12201
<b>1:3 Matching</b>	-0.052	0.1010	-0.509	0.611	95154	71764	23390
<b>1:5 Matching</b>	-0.176**	0.0800	-2.200	0.028	95154	49637	45517
<b>Above 1000 m</b>	1.696***	0.1690	10.044	0.000	1937	108	1829
<b>All elevation</b>	-0.165	0.1350	-1.216	0.224	97091	93210	3881

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .



**Figure 4.B.1:** Cumulative avoided forest loss ('000 ha) and avoided carbon dioxide emissions (MtCO<sub>2</sub>-eq): Dryland forest > 1,000m DDD, 2012-2018 (A); Peatland > 1,000m DDD, 2012-2017 (B). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI.

#### 4.B.4 “Intact primary” forest results

The following tables show the results from using a tighter definition of “forest” determined by extent of forest intactness and contiguity (Margono *et al.*, 2014; Turubanova *et al.*, 2018).

**Table 4.B.19:** DD dryland “intact primary” forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
2012	-0.018	0.0250	-0.743	0.457	160012	152252	7760
2013	0.043	0.0360	1.194	0.232	160012	152252	7760
2014	0.08*	0.0480	1.667	0.095	160012	152252	7760
2015	0.172***	0.0560	3.070	0.002	160012	152252	7760
2016	0.285***	0.0670	4.270	0.000	160012	152252	7760
2017	0.306***	0.0700	4.370	0.000	160012	152252	7760
2018	0.325***	0.0730	4.434	0.000	160012	152252	7760

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.20:** DDD dryland “intact primary” forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
2012	-0.039	0.0260	-1.505	0.132	160012	152252	7760
2013	0.003	0.0390	0.067	0.947	160012	152252	7760
2014	0.019	0.0540	0.359	0.719	160012	152252	7760
2015	0.091	0.0650	1.407	0.159	160012	152252	7760
2016	0.184**	0.0780	2.343	0.019	160012	152252	7760
2017	0.185**	0.0860	2.138	0.033	160012	152252	7760
2018	0.184*	0.0950	1.936	0.053	160012	152252	7760

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.21:** DD peatland “intact primary” forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	0.007	0.0280	0.265	0.791	95154	91302	3852
<b>2013</b>	-0.038	0.0390	-0.970	0.332	95154	91302	3852
<b>2014</b>	-0.132**	0.0530	-2.508	0.012	95154	91302	3852
<b>2015</b>	-0.104*	0.0610	-1.703	0.089	95154	91302	3852
<b>2016</b>	-0.104	0.0710	-1.459	0.145	95154	91302	3852
<b>2017</b>	-0.094	0.0740	-1.274	0.203	95154	91302	3852

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.22:** DDD peatland “intact primary” forest ATT

Endpoint	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2012</b>	0.002	0.0290	0.078	0.938	95154	91302	3852
<b>2013</b>	-0.048	0.0430	-1.116	0.264	95154	91302	3852
<b>2014</b>	-0.148**	0.0600	-2.474	0.013	95154	91302	3852
<b>2015</b>	-0.124*	0.0720	-1.714	0.087	95154	91302	3852
<b>2016</b>	-0.129	0.0870	-1.486	0.137	95154	91302	3852
<b>2017</b>	-0.125	0.0970	-1.297	0.195	95154	91302	3852

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.B.23:** DD dryland “intact primary” forest, avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	-1.711	-2.955
2013	3.987	6.883
2014	7.397	12.771
2015	15.956	27.549
2016	26.398	45.577
2017	28.347	48.942
2018	30.120	52.003

**Table 4.B.24:** DDD dryland “intact primary” forest, avoided forest cover loss and carbon emissions reductions

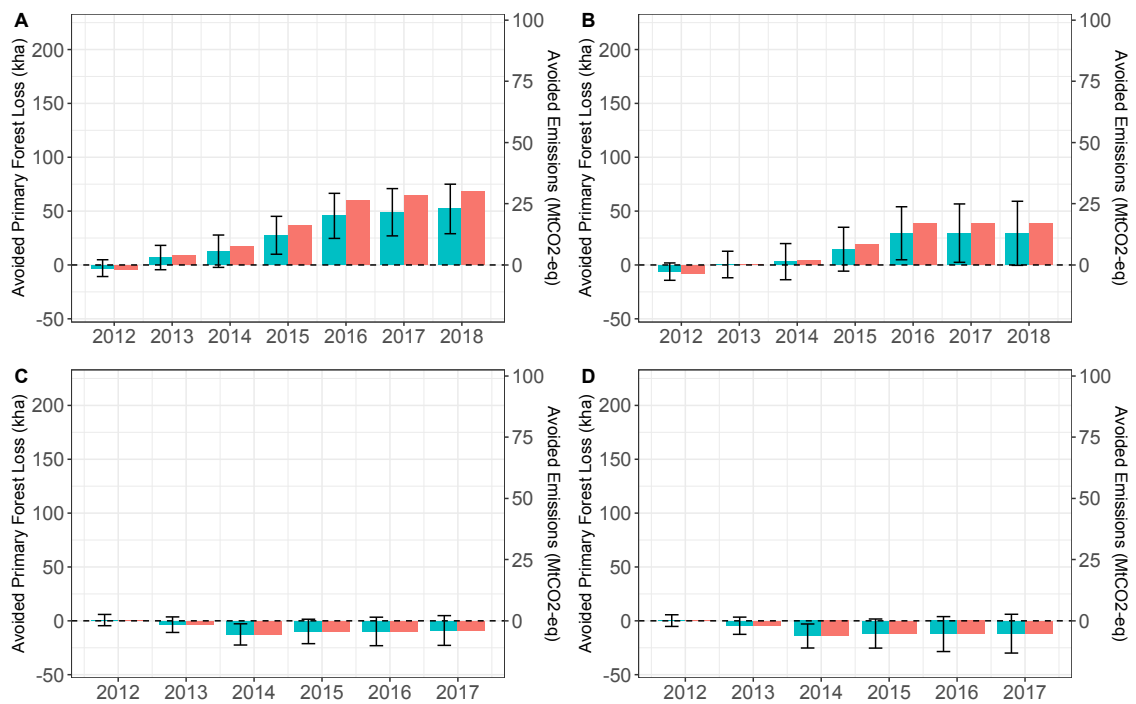
Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	-3.58	-6.186
2013	0.243	0.42
2014	1.782	3.076
2015	8.47	14.623
2016	17.04	29.42
2017	17.11	29.554
2018	17.01	29.383

**Table 4.B.25:** DD peatland “intact primary” forest, avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	0.310	0.706
2013	-1.569	-3.572
2014	-5.537	-12.607
2015	-4.328	-9.853
2016	-4.326	-9.848
2017	-3.942	-8.975

**Table 4.B.26:** DDD peatland “intact primary” forest, avoided forest cover loss and carbon emissions reductions

Endpoint	Emissions reduction (MtCO <sub>2</sub> -eq)	Avoided forest loss (’000 ha)
2012	0.095	0.216
2013	-2.00	-4.553
2014	-6.183	-14.078
2015	-5.189	-11.814
2016	-5.402	-12.299
2017	-5.234	-11.917

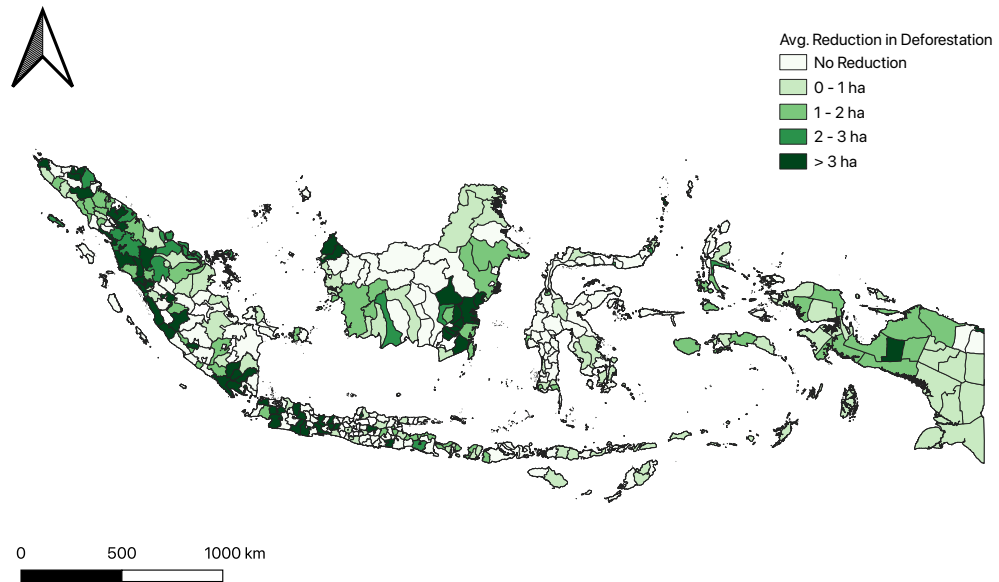


**Figure 4.B.2:** Cumulative avoided “intact primary” forest loss (’000 ha) and avoided carbon dioxide emissions (MtCO<sub>2</sub>-eq): dryland forest DD, 2012-2018 (A); dryland forest DDD, 2012-2018 (B); peatland forest DD, 2012-2017 (C); peatland DDD, 2012-2017 (D). The blue columns and left-hand y-axis in each panel shows the quantity of avoided forest loss while the red columns and right-hand y-axis shows the quantity of carbon emissions avoided. All quantities are aggregated up to the level of the whole Moratorium. Error bars denote the 95% CI.



#### 4.B.5 Spatially heterogeneous effects of the Moratorium, 2011-17

Figure 4.B.3 shows the heterogeneous effects of the Moratorium on dryland forest cover using district-level ATT generated from the aggregate ATT. Relatively large positive effects (shaded dark green) are revealed in the mountainous regions of Sumatra, districts in East and South Kalimantan, as well as parts of Java and Papua. Some of these districts are also the location of high shares of the total area of Indonesia's land designated as concessions pre-2011 (Figure 4.A.1) but this is not a consistent pattern across the country. For example, there is little or no effect in the lowlands of Sumatra, nor in parts of Central Kalimantan, where many palm oil and timber concessions are located.



**Figure 4.B.3:** Spatially heterogeneous effects of the Moratorium on dryland and peatland forests, 2011-17

## 4.C Summary Statistics

**Table 4.C.1:** Dryland forest non-concession grid cells, Moratorium

Statistic	Mean	S.D.	Min	Max	N <sub>obs</sub>
<b>AGC</b>	157.819	38.253	0.500	239.517	160012
<b>Distance cities</b>	1588.080	1340.752	2.000	6682.944	160012
<b>Distance logging</b>	87.257	172.887	0.00003	903.985	160012
<b>Distance palm oil</b>	102.348	141.246	0.00004	679.341	160012
<b>Distance timber</b>	123.843	143.885	0.00004	584.680	160012
<b>Distance roads</b>	35.148	32.350	0.056	220.476	160012
<b>Elevation</b>	470.037	271.519	-0.375	999.991	160012
<b>Peat depth</b>	0.000	0.000	0	0	160012
<b>Slope</b>	13.141	6.673	0.000	41.013	160012
<b>Population 2005</b>	17.298	135.354	0.000	10718.960	160012
<b>Population 2010</b>	14.925	126.867	0.000	10194.360	160012
<b>Forest cover 2005</b>	109.083	40.862	0.000	142.402	160012
<b>Forest cover 2011</b>	103.784	44.237	0.000	142.050	160012
<b>Forest cover 2016</b>	98.062	47.467	0.000	142.050	160012
<b>Forest cover 2017</b>	97.259	47.890	0.000	141.850	160012
<b>Forest cover 2018</b>	96.497	48.281	0.000	141.850	160012

**Table 4.C.2:** Dryland forest non-concession grid cells, non-Moratorium

Statistic	Mean	S.D.	Min	Max	N <sub>obs</sub>
<b>AGC</b>	92.168	52.597	0.500	239.803	407622
<b>Distance cities</b>	579.065	809.565	0.000	6155.591	407622
<b>Distance logging</b>	230.052	253.851	0.00000	911.465	407622
<b>Distance palm oil</b>	167.752	197.904	0.00002	685.338	407622
<b>Distance timber</b>	128.035	168.030	0.00003	584.152	407622
<b>Distance roads</b>	12.341	21.137	0.000	283.699	407622
<b>Elevation</b>	207.618	229.919	-3.000	999.982	407622
<b>Peat depth</b>	0.000	0.000	0	0	407622
<b>Slope</b>	6.146	5.456	0.000	39.215	407622
<b>Population 2005</b>	328.561	1595.480	0	107782	407622
<b>Population 2010</b>	328.886	1474.688	0	87358	407622
<b>Forest cover 2005</b>	95.494	44.674	0.000	143.305	407622
<b>Forest cover 2011</b>	87.948	46.892	0.000	143.215	407622
<b>Forest cover 2016</b>	79.604	49.271	0.000	143.215	407622
<b>Forest cover 2017</b>	78.393	49.590	0.000	143.215	407622
<b>Forest cover 2018</b>	77.283	49.902	0.000	143.215	407622

**Table 4.C.3:** Dryland forest concession grid cells, Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	156.486	37.495	6.432	225.012	8577
<b>Distance cities</b>	1534.415	1090.362	17	6389	8577
<b>Distance logging</b>	19.875	63.949	0.0001	754.726	8577
<b>Distance palm oil</b>	46.345	90.734	0.0005	634.679	8577
<b>Distance timber</b>	56.398	98.799	0.00002	535.576	8577
<b>Distance roads</b>	31.006	26.855	0.193	157.039	8577
<b>Elevation</b>	351.910	251.125	0.873	999.917	8577
<b>Peat depth</b>	0.000	0.000	0	0	8577
<b>Slope</b>	11.088	6.683	0.078	35.263	8577
<b>Population 2005</b>	7.406	50.199	0.000	2385.000	8577
<b>Population 2010</b>	6.910	50.805	0.000	2481.000	8577
<b>Forest cover 2005</b>	105.033	42.964	0.000	141.074	8577
<b>Forest cover 2011</b>	98.834	46.061	0.000	141.074	8577
<b>Forest cover 2016</b>	92.489	49.188	0.000	141.074	8577
<b>Forest cover 2017</b>	91.435	49.675	0.000	141.074	8577
<b>Forest cover 2018</b>	90.567	50.063	0.000	141.074	8577

**Table 4.C.4:** Dryland forest concession grid cells, non-Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	139.703	46.721	0.500	242.199	254107
<b>Distance cities</b>	1422.768	1131.570	0.778	6257.167	254107
<b>Distance logging</b>	29.753	66.249	0.00002	757.565	254107
<b>Distance palm oil</b>	45.676	87.319	0.0001	638.973	254107
<b>Distance timber</b>	53.801	108.858	0.0001	568.691	254107
<b>Distance roads</b>	30.790	29.945	0.089	184.405	254107
<b>Elevation</b>	193.861	202.703	0.507	999.970	254107
<b>Peat depth</b>	0.000	0.000	0	0	254107
<b>Slope</b>	6.500	5.413	0.000	40.154	254107
<b>Population 2005</b>	15.568	131.062	0.000	29963.160	254107
<b>Population 2010</b>	15.141	104.916	0.000	12425.760	254107
<b>Forest cover 2005</b>	104.498	42.969	0.000	143.123	254107
<b>Forest cover 2011</b>	97.803	46.314	0.000	143.123	254107
<b>Forest cover 2016</b>	90.610	49.492	0.000	143.123	254107
<b>Forest cover 2017</b>	89.656	49.873	0.000	143.123	254107
<b>Forest cover 2018</b>	88.787	50.232	0.000	143.123	254107

**Table 4.C.5:** Peatland forest non-concession grid cells, Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	119.680	52.600	0.500	248.805	95154
<b>Distance cities</b>	1738.694	1250.135	1.000	6541.667	95154
<b>Distance logging</b>	49.646	63.999	0.0001	862.399	95154
<b>Distance palm oil</b>	59.646	60.120	0.0001	672.619	95154
<b>Distance timber</b>	134.295	143.310	0.0001	585.525	95154
<b>Distance roads</b>	54.333	43.582	0.022	231.603	95154
<b>Elevation</b>	87.082	162.877	-1.715	999.828	95154
<b>Peat depth</b>	5.018	3.314	1	10	95154
<b>Slope</b>	3.292	5.535	0.00000	44.659	95154
<b>Population 2005</b>	21.879	190.073	0.000	14271.050	95154
<b>Population 2010</b>	20.223	199.281	0.000	18457.200	95154
<b>Forest cover 2005</b>	111.444	39.854	0.000	143.530	95154
<b>Forest cover 2011</b>	106.188	43.791	0.000	143.530	95154
<b>Forest cover 2016</b>	100.822	47.414	0.000	143.530	95154
<b>Forest cover 2017</b>	100.116	47.855	0.000	143.530	95154
<b>Forest cover 2018</b>	99.418	48.270	0.000	143.530	95154

**Table 4.C.6:** Peatland forest non-concession grid cells, non-Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	68.444	48.624	0.500	230.178	191668
<b>Distance cities</b>	783.989	987.146	0	6110	191668
<b>Distance logging</b>	125.511	170.883	0.00001	909.426	191668
<b>Distance palm oil</b>	77.761	130.750	0.0001	669.294	191668
<b>Distance timber</b>	89.125	138.527	0.00002	584.481	191668
<b>Distance roads</b>	24.147	36.776	-0.096	240.210	191668
<b>Elevation</b>	48.089	109.444	-6.848	999.334	191668
<b>Peat depth</b>	4.234	2.959	1	10	191668
<b>Slope</b>	1.985	3.056	0.000	36.552	191668
<b>Population 2005</b>	429.286	2444.776	0.000	106258.800	191668
<b>Population 2010</b>	435.098	2318.126	0.000	85667.030	191668
<b>Forest cover 2005</b>	98.455	43.815	0.000	142.203	191668
<b>Forest cover 2011</b>	89.987	46.968	0.000	141.715	191668
<b>Forest cover 2016</b>	81.242	49.863	0.000	141.715	191668
<b>Forest cover 2017</b>	80.026	50.222	0.000	141.715	191668
<b>Forest cover 2018</b>	78.896	50.573	0.000	141.715	191668

**Table 4.C.7:** Peatland forest concession grid cells, Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	108.714	52.756	0.897	226.424	6930
<b>Distance cities</b>	1219.530	984.166	13	5424	6930
<b>Distance logging</b>	46.441	53.984	0.0001	271.455	6930
<b>Distance palm oil</b>	17.699	36.409	0.0001	409.311	6930
<b>Distance timber</b>	63.798	91.195	0.00002	557.789	6930
<b>Distance roads</b>	34.717	35.464	0.146	189.795	6930
<b>Elevation</b>	54.809	100.569	1.411	890.440	6930
<b>Peat depth</b>	5.615	3.460	1.000	10.000	6930
<b>Slope</b>	2.586	4.481	0.016	32.867	6930
<b>Population 2005</b>	15.282	56.990	0.000	1632.000	6930
<b>Population 2010</b>	11.285	46.319	0.000	1945.800	6930
<b>Forest cover 2005</b>	100.037	45.993	0.000	141.297	6930
<b>Forest cover 2011</b>	93.515	48.552	0.000	141.297	6930
<b>Forest cover 2016</b>	86.579	51.578	0.000	141.297	6930
<b>Forest cover 2017</b>	85.614	51.923	0.000	141.297	6930
<b>Forest cover 2018</b>	84.581	52.327	0.000	141.297	6930

**Table 4.C.8:** Peatland forest concession grid cells, non-Moratorium

<b>Statistic</b>	<b>Mean</b>	<b>S.D.</b>	<b>Min</b>	<b>Max</b>	<b>N<sub>obs</sub></b>
<b>AGC</b>	112.349	55.941	0.595	240.324	113954
<b>Distance cities</b>	1519.586	1199.065	3.167	6291.000	113954
<b>Distance logging</b>	39.990	56.940	0.00002	732.571	113954
<b>Distance palm oil</b>	32.212	49.361	0.0001	636.106	113954
<b>Distance timber</b>	65.691	119.646	0.0003	571.354	113954
<b>Distance roads</b>	38.486	34.998	0.034	189.846	113954
<b>Elevation</b>	62.648	96.492	0.267	998.759	113954
<b>Peat depth</b>	4.837	3.439	1.000	10.000	113954
<b>Slope</b>	2.516	3.568	0.000	34.959	113954
<b>Population 2005</b>	19.275	99.717	0.000	9538.440	113954
<b>Population 2010</b>	17.854	87.030	0.000	9514.040	113954
<b>Forest cover 2005</b>	104.618	42.932	0.000	141.937	113954
<b>Forest cover 2011</b>	97.172	46.966	0.000	141.686	113954
<b>Forest cover 2016</b>	89.343	50.675	0.000	141.322	113954
<b>Forest cover 2017</b>	88.366	51.077	0.000	141.322	113954
<b>Forest cover 2018</b>	87.461	51.441	0.000	141.322	113954

## 4.D Methods

### 4.D.1 Summary of empirical approach

The impacts of the 2011 Moratorium on forest cover and carbon emissions are estimated for the period starting in 2011 and ending variously in years 2012-2018. Up until 2017, the estimates are undertaken separately for dryland and peatland forest but for the period 2011-2018 the estimates are for dryland forest only due to the additional restrictions on peatland implemented in 2017 (Alisjahbana and Busch, 2017). Pre- and post-Moratorium panel data on forest cover spanning the period 2000-2018 are used. Identification of the causal effect stems from a difference-in-differences (DD) research design, which allows for a comparison of the treated (Moratorium) grid cells with untreated, control (non-Moratorium) grid cells, while controlling for pre-existing differences and a secular counterfactual trend in forest cover. A matching approach is taken so that counterfactual cells have the same or similar characteristics as the treated cells. A key identification assumption is that of parallel trends in unobservable cell characteristics between Moratorium and non-Moratorium cells. Although impossible to test directly, we undertake several placebo tests in order to check this assumption. The main analysis uses pre-treatment data from 2004. The placebo tests require data from before the time horizon considered in the main analysis and so also use data from 2000 to 2004. DD can be implemented in a number of ways, and decisions have to be taken about the appropriate DD estimator to deploy. We select among a number of different parametric and non-parametric (matched) DD and triple difference (DDD) estimators through a process of empirical testing of typical DD identification assumptions, and a sensitivity analysis of matching estimators and

standard parametric (linear) models. This process of model selection culminates with a preference for a matched triple differences estimator following [Chabé-Ferret and Subervie \(2013\)](#) and [Chabé-Ferret \(2017\)](#). After outlining the identification problem, the following sub-sections provide: an account of the estimators; our process of model selection via empirical testing of their identification assumptions, including placebo tests; and, the robustness checks applied to our main results, including sensitivity to matching procedures and a test of the Standard Unit Treatment Value Assumption (SUTVA), another identification assumption for DD.

### The identification problem

The core empirical problem with evaluating the causal impact of the Moratorium is one of selection bias: the forest areas *treated* by the Moratorium differ in their observable and unobservable characteristics. Any imbalance in these characteristics implies that a direct comparison of forest cover in Moratorium and untreated non-Moratorium areas will capture pre-existing imbalances in, for example, their suitability for palm oil or timber production, level of forest cover or susceptibility to deforestation. These factors would then confound the estimate of the treatment effect. This generic programme evaluation problem, and the associated imbalances in underlying characteristics, is common in the analysis of area-based conservation policies ([Albers and Ferraro, 2006](#); [Andam \*et al.\*, 2008](#); [Joppa and Pfaff, 2010b](#)). Protected areas, for instance, are often established in economically-marginal land with “few profitable alternative uses” ([Albers and Ferraro, 2006](#); [Sloan \*et al.\*, 2012](#)). Forest areas covered by the Moratorium are similarly distinctive from those not covered by the Moratorium.

[Figure 4.2.1](#) shows the average level of forest cover for grid cells outside concessions and cells within concessions, for both Moratorium and non-Moratorium areas. Areas under the Moratorium start with higher levels of forest cover in 2000, and this remains true throughout the period of observation and for both peatland and dryland forest, and for all concession types (palm oil, logging and timber). In all but one within-concession case, there are clear differences in forest levels between the treated and control cells. [Figure 4.2.1](#) also illustrates forest cover trends over our observation period. The general picture is one of steady decline, with forest cover declining by at least 10 percentage points over the 18-year observation period. These trends are steeper in concession cells than in non-concession cells irrespective of whether forest is part of the Moratorium or not. [Tables 4.C.1 - 4.C.8](#) present descriptive statistics illustrating additional differences between the Moratorium and non-Moratorium forest areas in terms of the key variables that determine the likelihood of forest areas becoming new concessions.

The essence of the DD design is depicted in [Figure 4.2.1](#). We have pre-and post-Moratorium data, and spatially well-defined treatment (Moratorium) and control (non-Moratorium) groups. A simple parametric DD would identify the impact of the Moratorium by estimating the difference in forest cover before and after the Moratorium was implemented in areas covered by the Moratorium, and takes from this the

difference in forest cover in the non-Moratorium areas. Under certain assumptions DD approaches identify causal impacts, removing both the pre-existing differences between Moratorium and non-Moratorium areas, which would confound a simple comparison of treatment and control, and the trend in non-Moratorium areas under the assumption that this trend would have continued in the post-implementation period in Moratorium areas. Two key assumptions for the control area trend to act as a counterfactual for the treatment and DD estimates to be interpreted as causal are: (1) the unobservable trends in both Moratorium and non-Moratorium areas are the same on average; and, (2) the Stable Unit Treatment Value Assumption (SUTVA), that is, there are no spillovers ('leakage') from Moratorium areas to non-Moratorium areas which would confound the estimated treatment effect. In what follows we explain how we address these identification issues.

Two further issues are noteworthy for our empirical approach. First, the Moratorium mandates that district governments stop issuing new concession licenses in forest areas covered by the Moratorium. Yet, after 2011 licenses continued to be issued in forest areas that dropped out of the Moratorium determined by boundaries established in 2011. Our measure of impact is based on a treatment group with boundaries that were established in 2011, and we assume that any concession-driven deforestation observed within these boundaries after 2011 violated the Moratorium. Second, our empirical strategy must overcome the fact that after 2011 we do not observe where new oil palm, timber and logging concessions would have located in the absence of the Moratorium, in forest areas covered by the Moratorium. For this reason, the empirical strategy must control for characteristics of grid cells that determine the likelihood becoming a new concession, and ensure that these characteristics are the same or similar between Moratorium and non-Moratorium areas. The hope is that by *balancing* the observable characteristics, the confounding effect of unobservable characteristics will also be balanced or controlled for (Imbens, 2014; Wooldridge, 2010; Abadie and Imbens, 2008).

We take several steps to ensure that the grid cells compared in the Moratorium and non-Moratorium forest areas are similar in terms of their observable characteristics. First, we exclude cells which are part of the Indonesian protected area network, both within and outside the Moratorium, as conversion in these cells is already strictly prohibited. Second, we remove all cells outside of concessions with an elevation of 1,000 metres or more above sea level. The likelihood of these cells being a realistic proposition for a concession in either Moratorium or non-Moratorium areas is close to zero because above 1,000 metres land is unsuitable for palm oil cultivation (Austin *et al.*, 2015) and for *Acacia mangium*, the main tree species employed for the production of wood pulp and paper (Krisnawati *et al.*, 2011). Robustness to this choice of sample is discussed below, in the final sub-section. Third, as noted earlier, we undertake separate analyses for dryland and peatland areas. These three steps can be thought of as a type of "matching" on elevation and land types, whereby Moratorium and non-Moratorium grid cells are constrained to be similar in these dimensions. Lastly, we use a suite of matching approaches to ensure the analysis takes place on a sample whose Moratorium grid cells can be matched to non-Moratorium



grid-cells on the basis of a set of more specific grid-level characteristics. The main analysis uses the sample arising from a one-to-one propensity score caliper matching procedure. Other matching procedures are used to check robustness to this particular procedure. In the propensity score estimates, the sample ensures common support in the propensity score (Imbens, 2014). The matched dataset for the main analysis is used for parametric and non-parametric estimators. For the analysis of the impact of the Moratorium across forest types (peatland, dryland) and land uses (concession, non-concession), matching approaches are applied separately. For instance, the propensity score is estimated separately for each forest type, and the matched dataset upon which the impact estimated differs from one forest type to another.

Our process of model selection involves a sequence of placebo and falsification tests, which provide insights on the likelihood that the identification assumption of DD, parallel trends, holds. Finally, we test the SUTVA assumption. This process leads us to select a triple differences (DDD) estimator, which provides a correction for non-parallel trends. The estimators and the process of model selection are now explained in detail.

### Estimators

We use a number of estimators to estimate the causal impact of the Moratorium. We estimate the Average Treatment on the Treated ( $ATT_{DD}$ ) using a DD design, using both parametric and a non-parametric propensity score matching estimator.

The observed data on the outcome variable,  $Y_{it}$ , for individual grid-cell  $i$  at time  $t$ , are generated via a switching process governed by the occurrence of the Moratorium, which is indicated by the variable  $D_{it}$  :

$$Y_{it} = Y_{0it} + D_{it} (Y_{1it} - Y_{0it})$$

where  $D_{it}$  is an indicator variable that interacts a variable for pre- and post- treatment periods,  $r_t$ , with an indicator for treatment status after the treatment period,  $d_i$  :  $D_{it} = r_t d_i$  (Lee, 2005, p. 100-101).  $Y_{0it}$  and  $Y_{1it}$  are the potential outcomes for grid cell  $i$  in the treated state (1) and untreated state (0) at time  $t$ . Potential outcomes can be separated into their expected values,  $\mu_{1it}$  and  $\mu_{0it}$ , and a mean zero random component,  $\varepsilon_{kit}$ :

$$Y_{0it} = \mu_{0it} + \varepsilon_{0it}$$

$$Y_{1it} = \mu_{1it} + \varepsilon_{1it}$$

where  $E[\varepsilon_{0it}] = E[\varepsilon_{1it}] = 0$ . Suppose that the  $\varepsilon_{kit}$  are comprised of an individual fixed effect,  $\lambda_i$ , and a transitory component,  $u_{kit}$ :  $\varepsilon_{kit} = \lambda_i + u_{kit}$ . If the pre-treatment period is  $T_0$  and the post-treatment period is  $T_1$ , then the difference-in-differences estimator uses the observed data from the switching equation to identify the Average Treatment on the Treated ( $ATT_{DD}$ ) defined in terms of the potential outcomes as follows:

$$DD = E[\Delta_{T_1, T_0} Y_{it} | d_i = 1] - E[\Delta_{T_1, T_0} Y_{it} | d_i = 0] \quad (4.8)$$

$$= \{E[\Delta_{T_1, T_0} \mu_{1it} | d_i = 1] - E[\Delta_{T_1, T_0} \mu_{0it} | d_i = 1]\} \\ + \{E[\Delta_{T_1, T_0} u_{1it} | d_i = 1] - E[\Delta_{T_1, T_0} u_{0it} | d_i = 0]\} \quad (4.9)$$

$$= E[\Delta_{T_1, T_0} Y_{1it} | d_i = 1] - E[\Delta_{T_1, T_0} Y_{0it} | d_i = 1] \quad (4.10)$$

$$= E[Y_{1iT_1} - Y_{0iT_1} | d_i = 1] = ATTD_D \quad (4.11)$$

where Eq. 4.8 uses the observed data, Eq. 4.9 reflects the potential outcomes, and Eq. 4.11 is simply a definition. Eq. 4.9 follows because the individual fixed effects,  $\lambda_i$ , drop out when first differences are taken. Eq. 4.10 follows under the assumption of parallel trends in unobservables, that is if:

$$E[\Delta_{T_1, T_0} u_{1it} | d_i = 1] - E[\Delta_{T_1, T_0} u_{0it} | d_i = 0] = 0$$

and Eq. 4.11 is the definition of the Average Treatment Effect and follows because by definition:

$$E[Y_{1iT_0} - Y_{0iT_0} | d_i = 1] = 0$$

meaning that potential outcomes pre-treatment are identical for the treated group. Identification of  $ATTD_D$  using the DD estimator above does not rely on parametric assumptions, and is in principle a non-parametric estimator. With the assumption of a separable disturbance term:  $\varepsilon_{kit} = \lambda_i + u_{kit}$ , with individual heterogeneity,  $\lambda_i$ , fixed over time and between potential outcomes, the existence of panel data allows us to take differences within each individual grid cell,  $i$ , remove these fixed effects and effectively control for any endogeneity that might be associated with individual heterogeneity of this type. A parametric fixed effects regression would use a linear model to control for individual heterogeneity in a similar way. We use a range of parametric and non-parametric matching DD estimators to estimate the treatment effect of the Moratorium. The parametric DD estimator is specified as:

$$Y_{it} = \alpha + \beta_1 D_{it} + \sum_{k=2}^n \beta_k X_{kit} + \lambda_i + \theta_t + \varepsilon_{it} \quad (4.12)$$

where  $Y_{it}$  is forest cover in (non-concession) grid cell  $i$  in year  $t$ ,  $D_{it}$  is the time-varying Moratorium treatment indicator,  $X_{kit}$  are  $n$  time-varying control variables, which could include including climatic controls (minimum and maximum temperature, precipitation). Other controls, such as slope, elevation, and distance from cities and roads, are unnecessary for this particular estimation strategy because grid-cell fixed effects adjust for time-invariant controls.  $D_{it}$  is the “policy” variable taking the value 1 after 2010 for Moratorium areas, otherwise 0 and  $\beta_1$  is the ATT. This basic model controls for time-invariant characteristics via the individual grid-level fixed effects,  $\lambda_i$ , and time fixed effects,  $\theta_t$ , which capture shocks common to all grid cells such as weather shocks. The estimate of  $\beta_1$  in Eq. 4.12 provides an estimate of ATT that reflects the average of the annual impacts over the post-treatment period (from period  $T_1$  to  $T$ ). Also interesting is to estimate the impact year-by-year, which requires the following specification (see e.g. Wooldridge, 2021):

$$Y_{it} = \alpha + \sum_{s=\tau}^T \beta_{1s} D_{sit} + \sum_{k=2}^n \beta_k X_{it} + \lambda_i + \theta_t + \varepsilon_{it} \quad (4.13)$$

In Eq. 4.13, the  $\beta_{1s}$  coefficients represent the DD estimates of ATT for each year  $s$ . The coefficients are comparable in interpretation to the estimates from our alternative estimators, which belong to the non-parametric, matched DD family of estimators.

Matching DD methods have some advantages over parametric DD in that they do not rely on parametric assumptions to control for differences in observable characteristics (Imbens, 2014) and hence, are more likely to obtain like-for-like comparisons between treated grid cells and untreated ones. Matching estimators allow for arbitrary heterogeneity in the treatment effect compared to parametric approaches where the heterogeneity must be specified explicitly (Caliendo and Kopeinig, 2008). As discussed in the following sub-section, while differencing accounts for unobservable fixed effects in DD estimators, under the identifying assumptions matching ensures that the heterogeneous trend in the transitory component,  $\varepsilon_{it}$ , is balanced between treatment and control groups (see e.g. Chabé-Ferret and Subervie, 2013).

Matching methods have been used to evaluate other area-based conservation policies due to these properties (Andam *et al.*, 2008; Gaveau *et al.*, 2009; Joppa and Pfaff, 2010b; Nelson and Chomitz, 2011). One typical matching approach matches on the propensity score: the probability of a unit being in the treated or untreated group. We use caliper one-to-one propensity score matching so that each Moratorium cell is matched with a similar non-Moratorium cell on the basis of the likelihood of being in the Moratorium. While propensity score matching does rely on parametric assumptions in estimating the propensity score, the estimation of the treatment effect is otherwise non-parametric, so we maintain the distinction between parametric and non-parametric (matching) DD approaches in what follows.

Propensity score matching DD conditions the differences in forest cover on the propensity score  $P(X)$  to remove bias in levels and trends from confounding variables (Imbens, 2014; Chabé-Ferret and Subervie, 2013; Heckman *et al.*, 1997) under the assumption of selection on observables:  $d_i$  is independent of  $(Y_0, Y_1)$  when outcomes are conditioned on  $P(X)$ . Suppose that  $T$  is our final post treatment period, then together with the parallel trends assumption:

$$\begin{aligned} ATT_{DD,T} &= E_{P(X)} [E [Y_{1,T} - Y_{1,T_1} | d_i = 1, P(X_{T_m})] \\ &\quad - E [Y_{0,T} - Y_{0,T_1} | d_i = 1, P(X_{T_m})] \\ &= E [Y_{1T} - Y_{0T_1} | d_i = 1] \end{aligned} \quad (4.14)$$

where  $T$  ranges from 2012-2018, reflecting the different time-horizons over which the Moratorium's impact is estimated,  $T_1 = 2011$  is the treatment year, and  $T_m < T_1$  is the pre-treatment year(s) from which the matching variables' data are taken. To

estimate  $ATT_{DD,T}$  we use one-to-one caliper propensity score matching on data  $Y_{it}^M$  from the Moratorium grid cell  $i$ , matched with data  $Y_{it}^{j,NM}$  in cell  $j$  from the non-Moratorium grid-cells:

$$\widehat{ATT}_{DD,T} = \frac{1}{N_M} \sum_i I_0 \left[ [Y_{i,T}^M - Y_{i,T_1}^M] - [Y_{i,T}^{j,NM} - Y_{i,T_1}^{j,NM}] \right] \quad (4.15)$$

where  $I_0$  is an indicator variable that is equal to 1 if a grid cell  $i$  in the Moratorium has a counterfactual grid cell  $j$  in the non-Moratorium area whose propensity scores  $p_i$  and  $p_j$  fall within the caliper:

$$|p_i - p_j| < \epsilon$$

where  $\epsilon$  is a predetermined distance in propensity score space. The caliper defines the set of one-to-one matches from the non-Moratorium area,  $C(i)$ , such that  $j \in C(i)$ . Therefore,  $I_0 = I(\min |p_i - p_j| : |p_i - p_j| < \epsilon)$ . Our large dataset allows us to overcome the chief concern in the literature on optimal caliper choice: the loss of sample size with increased precision (Austin, 2009). We use a precise caliper of 0.0001 which leads to a small percentage of observations being dropped where Moratorium cells do not have matches within this caliper. These are relatively low-quality matches and do not appear in  $C(i)$ . In the process of model selection, described in the following sub-section, we show the results are in any event insensitive to the choice of a wider caliper of 0.01. The remaining matched sample is of size  $N_M$ . The average of the counterfactual matched outcome is taken in case of ties in the non-Moratorium group. Our process of model selection suggests that the parallel trends assumption, which underpins the identification of the treatment effect when using the DD estimator, may not hold. Triple Difference estimators (DDD) are argued to have weaker identification assumptions than DD estimators (e.g. Olden and Møen, 2020, p.11), and can be used as a robustness check or corrective measure (Olden and Møen, 2020; Chabé-Ferret, 2017). We use DDD as corrective measure following Chabé-Ferret and Subervie (2013). Our DDD estimator takes the form:

$$\begin{aligned} \widehat{ATT}_{DDD,T} = & \frac{1}{N_M} \sum_i I_0 \left\{ \left[ [Y_{i,T}^M - Y_{i,T_1}^M] - [Y_{i,T}^{j,NM} - Y_{i,T_1}^{j,NM}] \right] \right. \\ & \left. - \left( \left[ [Y_{i,T_1'}^M - Y_{i,T_0}^M] - [Y_{i,T_1'}^{j,NM} - Y_{i,T_0}^{j,NM}] \right] \right) \left( \frac{T - T_1}{T_1' - T_0} \right) \right\} \quad (4.16) \end{aligned}$$

where the second line reflects the correction for the non-parallel pre-treatment trends between  $T_0$  and  $T_1'$  with a correction  $\left(\frac{T-T_1}{T_1'-T_0}\right)$  to adjust the trend-correction for potentially different pre- and post-treatment time horizons.

The additional adjustment in the DDD estimator identifies the treatment effect in the event that the expected divergence of trends between the Moratorium and non-Moratorium areas is linear in time and the conditioning variables include pre-treatment outcomes including the year immediately prior to the treatment, in our case 2010. To demonstrate, leaving the conditioning on observable characteristics

implicit, in a typical DD framework the  $ATT_{DD}$  is non-parametrically identified as follows:

$$\begin{aligned} & E[\Delta_{T,T_1} Y_{it} | D = 1] - E[\Delta_{T,T_1} Y_{it} | D = 0] \\ &= ATT_{DD,T,T_1} + E[\Delta_{T,T_1} u_{1it} | d_i = 1] - E[\Delta_{T,T_1} u_{0it} | d_i = 0] \end{aligned} \quad (4.17)$$

where  $Y_{it}$  are the observable data,  $\Delta_{T,T_1}$  represents the change operator between the period  $T$  and treatment period  $T_1$ ,  $ATT_{DD,T,T_1} = E[\Delta_{T,T_1} Y_{1it} - \Delta_{T,T_1} Y_{0it}]$  (and as above  $Y_{ikt}$  represent the potential outcomes for treated:  $k = 1$ , and untreated:  $k = 0$  grid cells), and the last two terms in Eq. 4.17 are the trends in the unobservable characteristics in the treated and untreated states, which sum to zero if the parallel trends assumption holds.

Placebo tests are necessary to test whether or not the pre-trends are parallel in all cases for the simple DD estimator even when matching on several pre-treatment outcomes. If they are not parallel then this would imply that the DD estimator over the time horizon  $T_1$  to  $T$  is potentially contaminated by a diverging trend in unobservables: the difference between the last two terms of Eq. 4.17. To try and resolve this issue, which could be problematic over long-time horizons, the DDD approach (proposed by [Chabé-Ferret and Subervie \(2013\)](#)) could be adopted. The main identification assumptions are twofold. First, the non-parallel trend is linear:  $E[\Delta_t u_{1it} | d_i = 1] - E[\Delta_t u_{0it} | d_i = 0] = \gamma t$ . Second, the trend holds for all  $t$  and is unaffected by the treatment. With these assumptions, the confounding term in Eq. 4.17 becomes  $\gamma(T - T_1)$ . For pre-treatment time periods  $T_0$  to  $T'_1$  the DD estimator in Eq. 4.17 becomes:

$$\begin{aligned} & E[\Delta_{T'_1,T_0} Y_{it} | d_i = 1] - E[\Delta_{T'_1,T_0} Y_{it} | d_i = 0] \\ &= ATT_{DD,T'_1,T_0} + E[\Delta_{T'_1,T_0} u_{1it} | d_i = 1] - E[\Delta_{T'_1,T_0} u_{0it} | d_i = 0] \\ &= \gamma(T'_1 - T_0) \end{aligned} \quad (4.18)$$

because  $ATT_{DD,T'_1,T_0} = 0$  in the pre-treatment phase. This expression is estimated using propensity score matching DD in the pre-treatment period for Moratorium and non-Moratorium areas. Having conditioned on pre-treatment outcomes in the matching routine, a trend correction is required for the period  $T_1$  to  $T$ , which requires a conversion factor:  $\left(\frac{T-T_1}{T'_1-T_0}\right)$  to relate the estimate in Eq. 4.18 to the required treatment effect in Eq. 4.16. This is a minor augmentation of the DDD approach proposed in [Chabé-Ferret \(2017\)](#). The years used in the main analysis were  $T_0 = 2004$ ,  $T'_1 = 2010$ ,  $T_1 = 2011$  and  $T = \text{endpoint year: } 2012\text{-}2018$ . For the placebo tests the placebo treatment year is 2005 and the pre- and post-treatment periods considered are, respectively, 2000-2004 and 2005-2010, with sensitivity using 2005-2011.

#### 4.D.2 Model selection and balance tests

We now explain our empirical approach and the motivation for the choice of estimator. Our *model selection* process has four steps: (1) a comparison of non-parametric and parametric estimators; (2) placebo and falsification tests in time; (3) the response to the placebo tests in time; and, (4) placebo and falsification tests in space, as a final check on the selected estimator for parallel trends and the comparability of treatment and control areas.

##### Step 1: Parametric versus non-parametric estimators

Although  $ATT_{DD}$  is non-parametrically *identified* above, both parametric and non-parametric estimators can be used to *estimate*  $ATT_{DD}$ . Each controls for time invariant unobservables and each shares the following identification assumptions: (i) selection on observables; and, (ii) parallel trends in unobservable characteristics. Parametric estimators are typically less well-suited to dealing with individual heterogeneity in estimating treatment effects (Heckman *et al.*, 1998). In the difference-in-differences framework, the parallel trends assumption is often argued to be more likely to hold when treated units are closely matched to non-treated (Heckman *et al.*, 1998) and experimental-empirical evidence often supports this view (e.g. Ryan *et al.*, 2018). Furthermore, parametric approaches to conditioning on observable characteristics often place restrictive assumptions on functional forms of the relationship between conditioning variables and outcome variables, which are absent from non-parametric approaches and their estimation of treatment effects. Since these shortcomings can introduce bias in the estimates, we have an *a priori* preference for non-parametric estimates, among which we include propensity score matching. Nevertheless, matching estimators can also be subject to biases of their own (see e.g. Chabé-Ferret and Subervie, 2013; King and Nielsen, 2019).

To transparently investigate potential differences between estimators, we undertake some exploratory sensitivity analysis of different estimators. We estimate parametric DD models of the form described in Eq. 4.12 using a fixed-effects estimator. We compare these estimates to a propensity score, one-to-one caliper, matched DD estimator (NP-DD). To rule out the possibility that the balance in the sample of observable characteristics between Moratorium and non-Moratorium areas causes differences between the parametric and non-parametric estimates, we use the same matched sample for the parametric DD estimation as is used for matched DD estimators (see Imbens, 2014). We then undertake some basic sensitivity analysis on the non-parametric estimators by relaxing the precision of the matching procedure in two ways: (i) widening the caliper; and, (ii) sampling matches without replacement. We then compare the results with the parametric estimates and draw tentative conclusions.

## Results

Table 4.D.1 presents the results (the ATT for 2017) of this comparative analysis. The results do not undermine, and rather strengthen, our *a priori* preference for non-parametric matching estimators. The parametric estimators (P-1 to Conley-HAC 4) are based on Eq. 4.13. Table 4.D.1 reports the results of these specifications. P-1 controls for grid-cell ( $\lambda_i$ ) and year ( $\theta_t$ ) fixed effects. Model P-2 also controls for district-by-year fixed effects. Standard errors are clustered at the district level. These interaction effects capture heterogeneous trends by district. Models Conley 1 to Conley 4 are based on the same estimating equation as models P-1 and P-2, but employ Conley standard errors with respectively 10, 20, 50 and 100km thresholds for spatial autocorrelation decay. Specifications labelled Conley-HAC 1 to Conley-HAC 4 augment these latter models with Newey-West corrections for heteroskedasticity and temporal autocorrelations (10 lags).

Comparing estimators P-1 and P-2, we find that the parametric DD is sensitive to whether or not district-by-year interaction effects are included in the estimation. The estimated ATT for P-1 is 3.76, whereas with district-by-year fixed effects the estimate drops by almost 80% to 0.79. Models Conley 1 to Conley-HAC 4 differ from P-2 exclusively with respect to the standard errors around the point estimates. We find that for thresholds of spatial autocorrelation decay below 20Km, Conley and Conley-HAC standard errors are smaller than those clustered at the district level; these are similar in magnitude to Conley and Conley-HAC standard errors using a 50km radius around observational units, while employing a 100km radius results in larger standard errors and slightly lower statistical significance. Nonetheless, the parametric results remain strongly significant for any standard error specification we deploy. The sensitivity of the ATT's magnitude to the inclusion of district-by-year interactions most likely reflects that concessions, and land uses in general, are licensed and administered at the district level and these administrative processes may differ across districts, causing trends to diverge in a way that is correlated with the location of the Moratorium. This suggests that district-level trends should be controlled for when estimating treatment effects.

Rather than control for this possibility in a parametric model, we use a matching estimator which matches on grid-level characteristics. The matching variables explicitly capture important differences between the Moratorium and non-Moratorium grid cells, their dynamics and suitability for future concessions. We use pre-treatment values of: distance to concessions (palm oil, timber and logging), distance to roads and cities, population (2005 and 2010), forest cover for each year from 2005 to 2010, elevation, slope, peat depth and above-ground carbon stock in the year 2000. With regard to the heterogeneous trends that were shown to be an important factor in the parametric estimator, we control for heterogeneous pre-treatment trends by matching on pre-treatment trends in forest cover and population for several pre-treatment years: 2005-2010 for forest cover, 2005 and 2010 for population. We prefer this approach to the simple inclusion of district-by-year fixed-effects since the approach in principle controls for more sources of heterogeneous trends due to conditioning on other matching variables.

Comparing model P-2 with the propensity score matched result from the main analysis, NP-DD (caliper = 0.0001), we find that the NP-DD estimates are almost identical: NP-DD = 0.81 and P-2 = 0.79. Table 4.D.9 shows the extensive list of matching variables that underpin the NP-DD estimate, which help to control for heterogeneity.

When the matches are made less precise by obtaining matches without replacement (NP-1), the matching ATT decreases significantly: NP-1 = 0.36. NP-1, however, drops ten times the number of observations (> 70,000) in the process, and so is less favourable. Widening the caliper from 0.0001 to 0.01 (NP-2) drops fewer observations (2 in total), but this has little effect on the estimates: NP-2 = 0.81, as does removing the caliper altogether (NP-3), which is equivalent to nearest neighbour matching: NP-3 = 0.81. With no caliper and no replacement (NP-4), estimates increase towards the potentially problematic parametric estimates that do not control for district-by-year effects: NP-4 = 1.6. From these comparisons we conclude that the matching estimates are sensitive only to extreme reductions in precision of the matching (no replacement or no caliper), and there are some reasons to prefer the NP-DD estimator.

While these comparisons of estimators are exploratory, we prefer the propensity score matching estimator because it allows arbitrary heterogeneity in the treatment effect and closely matches treatment and control cells on the basis of pre-treatment trends, thereby accommodating heterogeneous pre-trends at district or indeed other levels, without imposing any additional structure on the estimation process. It is possible that controlling for district-by-year fixed effects in the parametric model may attenuate the estimate of the treatment effect by controlling for district level post-treatment responses, which would be a legitimate channel for the Moratorium being successful. Furthermore, linear parametric approaches to causal estimation also have been shown to be problematic in some multi-period cases (Imai and Kim, 2021). We therefore pursue a variety of non-parametric matching estimators to estimate the impact of the Moratorium while noting that the appropriate parametric approaches (P-2 to Conley-HAC 4) provide estimates of essentially the same magnitude, albeit slightly lower. We then subject our preferred estimator to rigorous robustness testing of the matching approach algorithm and identification assumptions.


## Step 2: Placebo tests in time

Our model selection approach turns to the identification assumptions of the DD estimator, namely the parallel trends assumption. Using the preferred non-parametric matched DD estimates of the treatment effect of the Moratorium for peatland and dryland forest, we attempt to empirically falsify the parallel trends assumption. To do this we undertake two distinct temporal placebo tests. The relevant time periods for the placebo tests can be seen in Figure 4.D.1, which reproduces the trends for non-concession dryland grid cells for the period 2000-2018, from Figure 4.2.1.

The first placebo test, Placebo Test 1, maintains the one-to-one matches between



the Moratorium grid cells and the non-Moratorium grid cells that were used in the original DD estimates shown in Figure 4.2.1 of the main text. In this case, matching variables are drawn from the period between  $T_0$  to  $T_1'$  in Figure 4.D.1. These cells are then subjected to a placebo Moratorium in 2005 rather than the actual Moratorium in 2011. We use the difference in forest cover post- and pre-placebo treatment, respectively, the period between 2000 and  $T_0$ , and  $T_0$  to  $T_1'$ , in Figure 4.D.1, and estimate the placebo treatment effect  $ATT_{DDP}$ . This tests how successful the strategy of matching on pre-treatment outcomes in 2005-2010 is in ensuring, on average, parallel trends between Moratorium and non-Moratorium grid cells. The null hypothesis is that these pre-treatment controls ensure that the parallel trends assumption cannot be falsified and hence, we ought not to observe a significant effect of the placebo treatment. The alternative hypothesis is that the parallel trends assumption can be falsified despite conditioning on outcomes between 2005-2010. Placebo Test 1 is motivated by Chabé-Ferret and Subervie (2013) and Chabé-Ferret (2017), who showed that matching on pre-treatment outcomes as a means of obtaining parallel pre-trends can sometimes fail and actually introduce bias. We match on several pre-treatment lags of the outcome variable, and so our expectation was that the test would fail to reject the null and hence, fail to falsify the parallel trends assumption.



**Figure 4.D.1:** Forest cover trends inside and outside the Moratorium, 2005-2018: non-concession dryland grid cells.

The second placebo test in time, Placebo Test 2, is a more traditional placebo test which shifts the entire analysis in time. That is, matching takes place on pre-placebo data (2000 to  $T_0$ ) rather than on pre-(real) treatment data ( $T_0$  to  $T_1$ ) as in Placebo Test 1. Again, the null hypothesis is that the parallel trends assumption of the matching estimator is not falsified. Both placebo tests are undertaken on each forest type. Together, these placebo tests can also be seen in light of the concerns raised by [Daw and Hatfield \(2018\)](#) about regression to the mean when matching on pre-treatment outcomes.

## Results

The placebo/falsification tests indicate that the parallel trends assumption fails in some cases for the DD estimator. Tables 4.D.2-4.D.7 report the results of all temporal placebo tests for dryland and peatland. For dryland areas, Placebo Test 1 yields a positive treatment effect in some cases (See [Table 4.D.2](#)). For stringency we undertake sensitivity on the definition of forested land, at the grid-cell scale, to be included in the analysis ranging from all land (untrimmed) to 30% forest cover in the base year, 2005, to 60% forest cover in 2005. Our central results use the unrestricted definition. In all cases but the 60% definition, we reject the null of Placebo Test 1. On the other hand, Placebo Test 2 is unable to falsify the parallel trends assumption, with well-defined zero impacts estimated across all three of our definitions of forest cover.

For peatland areas, Placebo Test 1 fails to reject the null hypothesis and falsify the parallel trends assumption, suggesting that matching on pre-treatment outcomes may be sufficient for identification. However, Placebo Test 2 rejects the null, with a statistically significant and negative effect of the placebo in 2005. This suggests that the parallel trends assumption may not hold in general for peatland. The failure of either of the placebo tests for dryland and peatland suggests remedial action is necessary.

### Step 3: Response to placebo tests in time

We now argue that NP-DDD estimator shown in [Eq. 4.16](#), which makes a correction in such cases, is therefore preferred. Rejection of the null hypothesis of the Placebo Test 1 is a sign that controlling for pre-treatment variables may have failed to ensure parallel pre-trends between Moratorium and non-Moratorium grid cells. Trends in unobservables diverge post-placebo and therefore are likely to do so post-treatment in the main analysis, when the real treatment happens in 2011 ( $T_1$ ). Rejection of Placebo Test 2 has a similar implication: falsification of the parallel trends assumption. While often framed as a *test* of the parallel trends assumption, these tests are only indicative of unobservable phenomena. Failure to reject the null is an encouraging sign, rejection is suggestive of a need for remedial measures.

Following [Chabé-Ferret and Subervie \(2013\)](#) and [Chabé-Ferret \(2017\)](#), falsification

can be ameliorated by a triple differences (DDD) approach that removes the pre-treatment difference in trends from the NP-DD estimate. Formally, this estimator (NP-DDD) is given by Eq. 4.16, where the last terms in square brackets represents the difference in the pre-trend between Moratorium and non-Moratorium areas between 2005 and 2010 at the matched grid-cell level. In Eq. 4.18, the assumption of a linear trend in the divergence of parallel trends is assumed, and the effect of this trend is corrected for the time horizon of the treatment effect being estimated (the choice of  $T$  in Figure 4.D.1). This specification is an adaptation to the proposal by Chabé-Ferret (2017) and its application to deforestation in Chabé-Ferret and Subervie (2013).

## Results

In practice, the NP-DDD estimator in Eq. 4.16 requires estimating and removing an adjusted (for time horizon) version the estimated ATT in the placebo test. For instance, in Table 4.D.2 we see that 0.181 of the estimated effect in the main NP-DD estimate (untrimmed) of effect of the Moratorium may well arise because of non-parallel trends. The NP-DDD estimator simply removes this estimated effect, adjusting for the different time-horizon, from the NP-DD estimate for each year  $T$  in which it is estimated.

Finally, we perform a third placebo test in time, Placebo Test 3, which aims at establishing whether the DDD methodology identifies a positive effect of the Moratorium on dryland cells when it is applied over the 2000-2010 (2011) placebo interval, with placebo treatment set to start in 2005. For the DDD estimator, we fail to reject the null of parallel trends in the cases where DD rejected the null. Our main results are derived using the NP-DDD estimator implemented in this way, and Figure 4.2.2 compares the results from NP-DD and NP-DDD estimators graphically to illustrate the differences.

### Step 4: Placebo tests in space

Finally, as a final check of the robustness of the NP-DDD estimator, we undertake another placebo-falsification test. We run spatial placebo tests using untreated cells in the Moratorium (treated) area as the placebo treatment cells. This spatial placebo test exploits a loophole in the Moratorium rules, which allows for continued deforestation in pre-2011 (pre-Moratorium) concession areas irrespective of whether the concessions are in the Moratorium or not. Since grid cells in concessions in the Moratorium area should be unaffected by the Moratorium, they are obvious candidates for Moratorium placebo treatment cells.

The null hypothesis of the spatial placebo test is that the NP-DDD estimator is well-identified: the parallel trends assumption cannot be falsified and there is no impact of the Moratorium on Moratorium concession cells compared to non-Moratorium concession cells. The alternative hypothesis is that there remain some residual unobservable differences between the Moratorium and non-Moratorium areas that

confound the estimate of the impact of the Moratorium. This would be the implication of a non-zero estimate of the Moratorium in the spatial placebo test. If the null hypothesis is rejected, it can be interpreted as a falsification of the parallel trends assumption. In principle, rejection of the null can inform a spatial triple differences DDD design to remove any deviation of parallel trends from the central NP-DDD estimate (Lee, 2005; Ravallion *et al.*, 2005; Gruber, 1994; Wooldridge, 2010). We run the spatial placebo test for all concession types, and also investigate the placebo for individual concession and forest types using the non-parametric (Chabé-Ferret, 2017) NP-DDD matching estimator of the treatment effect in 2011 ( $T_1$ ). The empirical model can be thought of as Eq. 4.16, with the exception that the treatment and control data come from the concession cells in the Moratorium and concession cells in the non-Moratorium areas, instead of Moratorium and non-Moratorium non-concession forested areas.

### Results.

Tables 4.D.7-4.D.8 show the results of these spatial placebo tests, all of which apply the NP-DDD estimator. Table 4.D.7 shows that for dryland areas, there are no differences in forest cover between Moratorium and non-Moratorium concession cells. Logging concessions show a negative effect, which is only statistically significant at the 10% level. For peatland areas, a significantly positive effect is detected for logging concessions only (Table 4.D.8), while for all concessions we fail to reject the null of the placebo test. We take these results to mean that on average, the NP-DDD is not susceptible to differences in trends due to unobserved differences between Moratorium and non-Moratorium areas.

With the propensity score, one-to-one, caliper matching NP-DDD estimator shown to satisfy the spatial placebo test, we undertake some additional checks and robustness analyses. First, we illustrate the success of our propensity score matching procedure at balancing the characteristics between Moratorium and non-Moratorium areas. We then test the robustness of the results to different types of propensity score matching approaches and a non-propensity score matching procedure, a series of Coarsened Exact Matching estimates. Finally we test the Stable Unit Treatment Value Assumption (SUTVA).

**Table 4.D.1:** Results of model selection (endpoint 2017)

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Main Result</b>	0.813***	0.0990	8.186	0.000	160012	152118	7894
NP-1	0.357***	0.0856	4.170	0.000	160012	88120	71892
NP-2	0.81***	0.1130	7.166	0.000	160012	160010	2
NP-3	0.81***	0.1130	7.166	0.000	160012	160012	0
NP-4	1.58***	0.0629	25.126	0.000	160012	160012	0
P-1	3.760***	0.2509	14.987	0.000	160012	152118	7894
P-2	0.7854***	0.2612	3.007	0.003	160012	152118	7894
Conley 1	0.7854***	0.1769	4.440	0.000	160012	152118	7894
Conley 2	0.7854***	0.2106	3.729	0.000	160012	152118	7894
Conley 3	0.7854***	0.2630	2.986	0.003	160012	152118	7894
Conley 4	0.7854**	0.3211	2.446	0.014	160012	152118	7894
Conley-HAC 1	0.7854***	0.1803	4.356	0.000	160012	152118	7894
Conley-HAC 2	0.7854***	0.2135	3.679	0.000	160012	152118	7894
Conley-HAC 3	0.7854***	0.2653	2.960	0.003	160012	152118	7894
Conley-HAC 4	0.7854**	0.3230	2.432	0.015	160012	152118	7894

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors for all non-parametric specifications

(NP-1) No replacement, caliper width: 0.0001. 1-to-1 matching keeping ties.

(NP-2) Wider caliper width: 0.01. 1-to-1 matching with replacement, keeping ties.

(NP-3) No caliper, 1-to-1 matching with replacement, keeping ties.

(NP-4) No caliper, no replacement, 1-to-1 matching, keeping ties.

(P-1) Parametric DD, on matched sample from NP-1, year fixed effects, SEs clustered at the district level.

(P-2) Parametric DD, on matched sample from NP-1, year and district\*year fixed effects, SEs clustered at the district level.

(Conley 1) Same as P-2, Conley SEs with 10km radius for spatial autocorrelation decay.

(Conley 2) Same as P-2, Conley SEs with 20km radius for spatial autocorrelation decay.

(Conley 3) Same as P-2, Conley SEs with 50km radius for spatial autocorrelation decay.

(Conley 4) Same as P-2, Conley SEs with 100km radius for spatial autocorrelation decay.

(Conley-HAC 1) Same as Conley 1, with Newey-West heteroskedasticity and autocovariance (HAC) correction (10 lags).

(Conley-HAC 2) Same as Conley 2, with Newey-West heteroskedasticity and autocovariance (HAC) correction (10 lags).

(Conley-HAC 3) Same as Conley 3, with Newey-West heteroskedasticity and autocovariance (HAC) correction (10 lags).

(Conley-HAC 4) Same as Conley 4, with Newey-West heteroskedasticity and autocovariance (HAC) correction (10 lags).

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.2:** Placebo test in time 1: DD dryland, 2005-2010 outcome, same matching

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	0.181**	0.0840	2.158	0.031	160012	152118	7894
<b>30%</b>	0.203**	0.0930	2.175	0.030	144501	136472	8029
<b>60%</b>	0.149	0.1010	1.468	0.142	127375	119284	8091

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.3:** Placebo test in time 2: DD dryland, 2005-2010 outcome, matching up to 2004

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	-0.048	0.0820	-0.580	0.562	160012	152408	7604
<b>30%</b>	-0.038	0.0920	-0.409	0.682	144501	136304	8197
<b>60%</b>	-0.064	0.1000	-0.640	0.522	127375	119677	7698

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.4:** Placebo test in time 1: DD peatland, 2005-2010 outcome, same matching

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	0.05	0.0930	0.533	0.594	95154	91024	4130
<b>30%</b>	-0.06	0.1010	-0.593	0.553	86897	82832	4065
<b>60%</b>	-0.047	0.1090	-0.433	0.665	77777	73554	4223

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.5:** Placebo test in time 2: DD peatland, 2005-2010 outcome, matching up to 2004

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>Untrimmed</b>	-0.39***	0.0890	-4.367	0.000	95154	91132	4022
<b>30%</b>	-0.486***	0.0970	-4.998	0.000	86897	82852	4045
<b>60%</b>	-0.496***	0.1050	-4.741	0.000	77777	73430	4347

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.6:** Placebo test in time 3: DDD dryland, placebo treatment in 2005, same matching

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>2010</b>	-0.018	0.0920	-0.191	0.848	160012	152118	7894
<b>2011</b>	0.049	0.1060	0.465	0.642	160012	152118	7894

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.7:** Spatial placebo: Dryland, 2011-2017 outcome

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>All concessions</b>	-0.33	0.3000	-1.099	0.272	8577	7702	875
<b>Palm Oil</b>	0.909	0.7750	1.173	0.241	2318	2249	69
<b>Timber</b>	0.86	0.9650	0.890	0.373	1183	1176	7
<b>Logging</b>	-0.68*	0.4000	-1.697	0.090	5193	5188	5

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Table 4.D.8:** Spatial placebo: Peatland, 2011-2017 outcome

	ATT	AI SE	t-stat	p-value	Treated	Matched	Dropped (treated)
<b>All concessions</b>	0.485	0.4090	1.185	0.236	6930	6604	326
<b>Palm Oil</b>	0.156	0.5960	0.262	0.793	4553	4553	0
<b>Timber</b>	0.437	1.7480	0.250	0.803	643	640	3
<b>Logging</b>	1.067*	0.5490	1.943	0.052	1765	1762	3

Notes: Abadie-Imbens (Abadie and Imbens, 2006) standard errors.

Maximum caliper width: 0.0001. 1-to-1 matching with replacement, keeping ties.

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .



### Balance tests for one-to-one caliper matching

The descriptive statistics in Tables 4.C.1-4.C.8 show that in the raw data there are large differences between the Moratorium and non-Moratorium areas. Matching variables are selected on the basis of these differences, shown in Table 4.D.9-4.D.20 by forest type and land use. For example, Table 4.D.9 describes the successful results of the matching procedure for dryland forest areas, and the success in balancing the treatment and control grid cells is similar across all of the remaining matching tables. In Table 4.D.9, forest cover in the six years prior to the treatment (2011) in the matched sample has a mean difference of no more than 0.5 ha compared to approximately 15 ha before matching. The standardised difference is no more than 0.55. These small mean differences are also complemented by two comparisons of the distributions of these characteristics: the Kolmogorov Smirnov test and a test of differences in the cumulative distribution function (eCDF). Both reveal highly balanced control variables in the matched sample, with small differences in the CDFs of the variables. Our conclusion is that the matching procedure performs well.

**Table 4.D.9:** Dryland forest, untrimmed, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	109.08	95.49	33.26	0.18	0.11	108.37	108.56	-0.44	0.03	0.02
<b>Forest 2006</b>	108.15	94.25	33.54	0.19	0.11	107.42	107.51	-0.22	0.04	0.02
<b>Forest 2007</b>	107.33	93.12	33.85	0.19	0.11	106.60	106.66	-0.16	0.03	0.02
<b>Forest 2008</b>	106.56	91.92	34.49	0.19	0.11	105.82	105.90	-0.19	0.04	0.02
<b>Forest 2009</b>	105.40	90.23	35.12	0.19	0.12	104.64	104.69	-0.11	0.04	0.02
<b>Forest 2010</b>	104.68	89.23	35.39	0.19	0.12	103.92	103.92	-0.00	0.04	0.02
<b>Elevation</b>	470.04	207.62	96.65	0.47	0.28	457.29	459.78	-0.93	0.07	0.03
<b>Slope</b>	13.14	6.15	104.81	0.47	0.29	12.85	13.12	-4.16	0.08	0.03
<b>Distance roads</b>	35.15	12.34	70.50	0.42	0.28	33.67	35.51	-5.79	0.09	0.05
<b>Distance cities</b>	1588.08	579.06	75.26	0.43	0.26	1512.51	1492.76	1.53	0.14	0.04
<b>Distance palm oil</b>	102.35	167.75	-46.31	0.20	0.07	103.94	102.26	1.16	0.06	0.04
<b>Distance timber</b>	123.84	128.03	-2.91	0.10	0.05	121.11	103.41	12.39	0.07	0.04
<b>Distance logging</b>	87.26	230.05	-82.59	0.32	0.20	90.53	92.81	-1.29	0.06	0.03
<b>AGC</b>	157.82	92.17	171.62	0.57	0.33	157.02	155.03	5.11	0.02	0.01
<b>Population 2005</b>	17.30	328.56	-229.96	0.49	0.28	18.06	21.38	-2.39	0.07	0.04
<b>Population 2010</b>	14.92	328.89	-247.47	0.50	0.29	15.60	18.91	-2.55	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.10:** Dryland forest, untrimmed, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2000</b>	111.83	99.34	31.42	0.17	0.09	111.20	111.46	-0.65	0.03	0.01
<b>Forest 2001</b>	111.41	98.71	31.84	0.18	0.10	110.78	110.98	-0.51	0.03	0.01
<b>Forest 2002</b>	110.87	97.99	32.19	0.18	0.10	110.23	110.41	-0.46	0.03	0.01
<b>Forest 2003</b>	110.55	97.55	32.40	0.18	0.10	109.89	110.05	-0.39	0.03	0.02
<b>Forest 2004</b>	109.78	96.46	32.91	0.18	0.11	109.11	109.21	-0.25	0.03	0.02
<b>Elevation</b>	470.04	207.62	96.65	0.47	0.28	458.01	460.68	-0.99	0.07	0.03
<b>Slope</b>	13.14	6.15	104.81	0.47	0.29	12.86	13.14	-4.22	0.07	0.03
<b>Distance roads</b>	35.15	12.34	70.50	0.42	0.28	33.73	35.50	-5.56	0.09	0.05
<b>Distance cities</b>	1588.08	579.06	75.26	0.43	0.26	1514.88	1501.71	1.02	0.14	0.04
<b>Distance palm oil</b>	102.35	167.75	-46.31	0.20	0.07	103.89	102.83	0.74	0.07	0.04
<b>Distance timber</b>	123.84	128.03	-2.91	0.10	0.05	120.99	104.20	11.76	0.07	0.04
<b>Distance logging</b>	87.26	230.05	-82.59	0.32	0.20	90.38	93.30	-1.66	0.06	0.04
<b>AGC</b>	157.82	92.17	171.62	0.57	0.33	157.04	154.93	5.41	0.02	0.01
<b>Population 2005</b>	17.30	328.56	-229.96	0.49	0.28	18.04	21.40	-2.43	0.07	0.04
<b>Population 2010</b>	14.92	328.89	-247.47	0.50	0.29	15.58	19.08	-2.70	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.11:** Dryland forest, > 30% forest cover in 2005, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	118.92	110.41	29.53	0.17	0.10	118.46	118.82	-1.22	0.03	0.01
<b>Forest 2006</b>	117.94	109.01	29.87	0.18	0.10	117.46	117.76	-1.01	0.03	0.01
<b>Forest 2007</b>	117.06	107.72	30.24	0.18	0.10	116.57	116.84	-0.86	0.03	0.02
<b>Forest 2008</b>	116.22	106.32	31.14	0.18	0.10	115.73	115.99	-0.82	0.03	0.02
<b>Forest 2009</b>	114.97	104.38	31.96	0.18	0.11	114.47	114.68	-0.63	0.03	0.02
<b>Forest 2010</b>	114.20	103.21	32.28	0.18	0.11	113.68	113.83	-0.44	0.03	0.02
<b>Elevation</b>	471.68	208.82	96.79	0.47	0.28	457.97	461.10	-1.16	0.06	0.03
<b>Slope</b>	13.12	6.21	103.49	0.46	0.29	12.80	13.07	-4.11	0.07	0.02
<b>Distance roads</b>	36.34	12.81	72.16	0.42	0.28	34.81	36.66	-5.76	0.08	0.04
<b>Distance cities</b>	1648.19	611.43	76.28	0.43	0.26	1568.10	1552.27	1.20	0.13	0.04
<b>Distance palm oil</b>	98.81	161.36	-46.28	0.20	0.07	100.44	98.33	1.53	0.06	0.03
<b>Distance timber</b>	127.91	123.30	3.17	0.12	0.05	124.45	106.56	12.44	0.06	0.03
<b>Distance logging</b>	81.38	221.39	-84.92	0.31	0.19	84.70	87.03	-1.38	0.06	0.03
<b>AGC</b>	159.12	94.00	174.17	0.56	0.33	158.25	156.31	5.08	0.03	0.01
<b>Population 2005</b>	16.50	318.27	-221.71	0.48	0.28	17.31	20.49	-2.27	0.07	0.04
<b>Population 2010</b>	14.30	319.18	-240.29	0.50	0.29	15.03	18.28	-2.49	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.12:** Dryland forest, > 30% forest cover in 2005, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2000</b>	121.18	113.93	26.19	0.16	0.07	120.79	121.12	-1.20	0.03	0.01
<b>Forest 2001</b>	120.81	113.34	26.89	0.16	0.08	120.41	120.71	-1.05	0.03	0.01
<b>Forest 2002</b>	120.34	112.64	27.49	0.17	0.09	119.92	120.20	-0.98	0.03	0.01
<b>Forest 2003</b>	120.06	112.23	27.82	0.17	0.09	119.63	119.89	-0.91	0.03	0.01
<b>Forest 2004</b>	119.47	111.26	28.84	0.17	0.09	119.03	119.23	-0.72	0.03	0.02
<b>Elevation</b>	471.68	208.82	96.79	0.47	0.28	457.34	462.50	-1.92	0.06	0.03
<b>Slope</b>	13.12	6.21	103.49	0.46	0.29	12.78	13.04	-3.90	0.07	0.02
<b>Distance roads</b>	36.34	12.81	72.16	0.42	0.28	34.70	36.79	-6.53	0.08	0.04
<b>Distance cities</b>	1648.19	611.43	76.28	0.43	0.26	1566.34	1551.51	1.13	0.13	0.04
<b>Distance palm oil</b>	98.81	161.36	-46.28	0.20	0.07	100.47	98.99	1.07	0.06	0.03
<b>Distance timber</b>	127.91	123.30	3.17	0.12	0.05	124.30	106.51	12.38	0.06	0.04
<b>Distance logging</b>	81.38	221.39	-84.92	0.31	0.19	84.79	87.60	-1.67	0.06	0.03
<b>AGC</b>	159.12	94.00	174.17	0.56	0.33	158.25	156.07	5.71	0.03	0.01
<b>Population 2005</b>	16.50	318.27	-221.71	0.48	0.28	17.33	20.61	-2.34	0.07	0.04
<b>Population 2010</b>	14.30	319.18	-240.29	0.50	0.29	15.04	18.43	-2.60	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.13:** Dryland forest, > 60% forest cover in 2005, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	126.61	121.58	25.40	0.16	0.09	126.37	126.41	-0.21	0.03	0.02
<b>Forest 2006</b>	125.74	120.27	25.92	0.16	0.09	125.48	125.45	0.12	0.03	0.02
<b>Forest 2007</b>	124.90	118.98	26.29	0.16	0.10	124.62	124.58	0.18	0.03	0.02
<b>Forest 2008</b>	124.08	117.55	27.39	0.16	0.10	123.78	123.71	0.28	0.03	0.02
<b>Forest 2009</b>	122.86	115.57	28.34	0.17	0.10	122.53	122.45	0.28	0.03	0.02
<b>Forest 2010</b>	122.09	114.34	28.80	0.17	0.10	121.74	121.64	0.40	0.03	0.02
<b>Elevation</b>	473.06	210.05	97.03	0.47	0.28	458.02	464.71	-2.49	0.05	0.02
<b>Slope</b>	13.11	6.28	102.41	0.46	0.29	12.76	13.00	-3.64	0.05	0.02
<b>Distance roads</b>	37.66	13.45	73.75	0.43	0.28	35.97	38.13	-6.68	0.07	0.04
<b>Distance cities</b>	1717.08	656.78	77.15	0.43	0.26	1634.30	1627.18	0.53	0.10	0.03
<b>Distance palm oil</b>	96.25	158.63	-47.90	0.20	0.07	97.88	95.02	2.14	0.05	0.02
<b>Distance timber</b>	133.23	120.82	8.43	0.13	0.06	129.22	110.48	12.86	0.05	0.02
<b>Distance logging</b>	76.03	214.49	-87.45	0.30	0.19	79.47	82.90	-2.11	0.05	0.02
<b>AGC</b>	160.50	95.94	176.78	0.56	0.33	159.58	157.33	6.03	0.03	0.01
<b>Population 2005</b>	15.77	303.70	-209.51	0.48	0.27	16.67	20.04	-2.37	0.07	0.04
<b>Population 2010</b>	13.75	305.33	-228.24	0.49	0.28	14.56	18.06	-2.66	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.14:** Dryland forest, > 60% forest cover in 2005, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
Forest 2000	128.21	124.16	21.27	0.15	0.07	128.01	128.27	-1.38	0.02	0.01
Forest 2001	127.93	123.70	22.11	0.15	0.08	127.72	127.96	-1.25	0.02	0.01
Forest 2002	127.57	123.17	22.79	0.15	0.08	127.35	127.57	-1.15	0.02	0.01
Forest 2003	127.36	122.87	23.17	0.15	0.08	127.14	127.35	-1.08	0.03	0.01
Forest 2004	126.95	122.17	24.41	0.15	0.09	126.72	126.90	-0.94	0.03	0.01
Elevation	473.06	210.05	97.03	0.47	0.28	458.49	464.89	-2.39	0.05	0.02
Slope	13.11	6.28	102.41	0.46	0.29	12.76	13.05	-4.38	0.05	0.02
Distance roads	37.66	13.45	73.75	0.43	0.28	36.07	38.10	-6.26	0.07	0.04
Distance cities	1717.08	656.78	77.15	0.43	0.26	1639.50	1622.20	1.29	0.10	0.03
Distance palm oil	96.25	158.63	-47.90	0.20	0.07	97.79	94.64	2.36	0.05	0.02
Distance timber	133.23	120.82	8.43	0.13	0.06	129.49	108.89	14.13	0.05	0.02
Distance logging	76.03	214.49	-87.45	0.30	0.19	79.26	81.72	-1.52	0.05	0.02
AGC	160.50	95.94	176.78	0.56	0.33	159.64	157.50	5.73	0.03	0.01
Population 2005	15.77	303.70	-209.51	0.48	0.27	16.62	19.81	-2.26	0.07	0.04
Population 2010	13.75	305.33	-228.24	0.49	0.28	14.51	17.84	-2.53	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.15:** Peatland forest, untrimmed, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
Forest 2005	111.44	98.45	32.59	0.19	0.11	111.00	110.89	0.28	0.03	0.01
Forest 2006	110.66	97.15	33.45	0.19	0.11	110.20	110.09	0.27	0.03	0.02
Forest 2007	109.83	95.79	34.25	0.20	0.12	109.35	109.23	0.27	0.03	0.01
Forest 2008	108.98	94.38	35.08	0.20	0.12	108.47	108.37	0.25	0.03	0.01
Forest 2009	107.95	92.58	36.22	0.20	0.12	107.41	107.28	0.32	0.04	0.02
Forest 2010	107.09	91.37	36.48	0.20	0.12	106.53	106.37	0.38	0.04	0.02
Elevation	87.08	48.09	23.94	0.19	0.11	83.31	76.20	4.41	0.18	0.10
Slope	3.29	1.98	23.62	0.08	0.03	3.20	3.03	3.07	0.19	0.09
Distance roads	54.33	24.15	69.26	0.40	0.26	53.10	54.26	-2.67	0.15	0.09
Distance cities	1738.69	783.99	76.37	0.41	0.25	1693.34	1743.92	-4.08	0.13	0.04
Distance palm oil	59.65	77.76	-30.13	0.17	0.10	59.10	57.01	3.43	0.14	0.08
Distance timber	134.29	89.12	31.52	0.24	0.12	127.53	117.72	7.04	0.14	0.08
Distance logging	49.65	125.51	-118.54	0.24	0.13	50.68	49.37	2.00	0.12	0.06
AGC	119.68	68.44	97.41	0.43	0.26	117.36	118.00	-1.23	0.08	0.02
Population 2005	21.88	429.29	-214.34	0.41	0.25	22.79	26.85	-2.09	0.05	0.03
Population 2010	20.22	435.10	-208.19	0.42	0.26	21.07	25.34	-2.10	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.16:** Peatland forest, untrimmed, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2000</b>	114.26	102.69	30.14	0.17	0.09	113.86	113.73	0.34	0.05	0.02
<b>Forest 2001</b>	113.81	102.00	30.60	0.18	0.10	113.40	113.26	0.35	0.05	0.02
<b>Forest 2002</b>	113.30	101.26	31.02	0.18	0.10	112.87	112.73	0.37	0.05	0.02
<b>Forest 2003</b>	113.00	100.80	31.32	0.18	0.10	112.57	112.40	0.44	0.06	0.02
<b>Forest 2004</b>	112.22	99.60	32.06	0.19	0.11	111.77	111.63	0.36	0.05	0.02
<b>Elevation</b>	87.08	48.09	23.94	0.19	0.11	83.50	76.31	4.44	0.18	0.10
<b>Slope</b>	3.29	1.98	23.62	0.08	0.03	3.21	3.07	2.43	0.19	0.10
<b>Distance roads</b>	54.33	24.15	69.26	0.40	0.26	53.05	54.03	-2.25	0.15	0.09
<b>Distance cities</b>	1738.69	783.99	76.37	0.41	0.25	1695.90	1744.58	-3.92	0.13	0.04
<b>Distance palm oil</b>	59.65	77.76	-30.13	0.17	0.10	59.12	56.79	3.84	0.13	0.08
<b>Distance timber</b>	134.29	89.12	31.52	0.24	0.12	127.75	118.00	6.99	0.14	0.08
<b>Distance logging</b>	49.65	125.51	-118.54	0.24	0.13	50.62	49.34	1.97	0.12	0.06
<b>AGC</b>	119.68	68.44	97.41	0.43	0.26	117.45	118.25	-1.53	0.08	0.02
<b>Population 2005</b>	21.88	429.29	-214.34	0.41	0.25	22.77	26.83	-2.09	0.05	0.03
<b>Population 2010</b>	20.22	435.10	-208.19	0.42	0.26	21.04	25.30	-2.09	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.17:** Peatland forest, > 30% forest cover in 2005, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	120.39	112.14	29.15	0.18	0.10	120.12	119.96	0.55	0.05	0.02
<b>Forest 2006</b>	119.58	110.69	30.42	0.18	0.10	119.29	119.19	0.33	0.05	0.02
<b>Forest 2007</b>	118.71	109.15	31.58	0.19	0.11	118.39	118.31	0.27	0.05	0.02
<b>Forest 2008</b>	117.81	107.55	32.72	0.19	0.11	117.46	117.39	0.25	0.05	0.02
<b>Forest 2009</b>	116.72	105.52	34.27	0.19	0.12	116.34	116.25	0.28	0.05	0.02
<b>Forest 2010</b>	115.81	104.13	34.54	0.20	0.12	115.40	115.30	0.28	0.05	0.02
<b>Elevation</b>	90.29	48.98	24.89	0.20	0.11	86.65	80.08	3.99	0.17	0.10
<b>Slope</b>	3.38	2.02	24.20	0.09	0.04	3.29	3.16	2.38	0.19	0.10
<b>Distance roads</b>	56.35	25.50	69.90	0.41	0.26	55.09	55.87	-1.78	0.18	0.10
<b>Distance cities</b>	1800.73	830.62	77.52	0.41	0.25	1757.17	1798.41	-3.32	0.15	0.04
<b>Distance palm oil</b>	61.92	75.10	-22.25	0.19	0.11	61.42	58.80	4.37	0.13	0.07
<b>Distance timber</b>	141.70	88.66	36.58	0.26	0.13	134.88	126.11	6.20	0.13	0.08
<b>Distance logging</b>	48.35	120.26	-115.38	0.23	0.12	49.35	48.12	1.94	0.11	0.05
<b>AGC</b>	122.20	70.04	100.25	0.44	0.26	119.89	120.86	-1.86	0.08	0.02
<b>Population 2005</b>	19.57	414.06	-220.31	0.42	0.25	20.44	24.18	-2.04	0.06	0.03
<b>Population 2010</b>	18.33	419.06	-210.11	0.42	0.26	19.15	23.11	-2.03	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.18:** Peatland forest, > 30% forest cover in 2005, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
Forest 2000	122.51	115.81	24.79	0.16	0.08	122.28	122.34	-0.22	0.05	0.02
Forest 2001	122.19	115.19	25.78	0.17	0.09	121.95	122.01	-0.22	0.05	0.02
Forest 2002	121.78	114.52	26.55	0.17	0.09	121.53	121.59	-0.22	0.06	0.02
Forest 2003	121.55	114.11	27.05	0.17	0.09	121.29	121.35	-0.20	0.06	0.02
Forest 2004	120.99	113.14	28.18	0.18	0.10	120.71	120.77	-0.20	0.06	0.02
Elevation	90.29	48.98	24.89	0.20	0.11	86.33	77.35	5.46	0.16	0.09
Slope	3.38	2.02	24.20	0.09	0.04	3.28	3.07	3.77	0.19	0.09
Distance roads	56.35	25.50	69.90	0.41	0.26	55.10	56.22	-2.53	0.18	0.10
Distance cities	1800.73	830.62	77.52	0.41	0.25	1757.06	1816.16	-4.75	0.15	0.04
Distance palm oil	61.92	75.10	-22.25	0.19	0.11	61.45	58.83	4.37	0.13	0.07
Distance timber	141.70	88.66	36.58	0.26	0.13	134.73	125.05	6.85	0.13	0.08
Distance logging	48.35	120.26	-115.38	0.23	0.12	49.33	48.09	1.95	0.11	0.05
AGC	122.20	70.04	100.25	0.44	0.26	119.86	120.56	-1.35	0.08	0.02
Population 2005	19.57	414.06	-220.31	0.42	0.25	20.44	24.22	-2.06	0.05	0.03
Population 2010	18.33	419.06	-210.11	0.42	0.26	19.14	23.18	-2.07	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.19:** Peatland forest, > 60% forest cover in 2005, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
Forest 2005	127.17	122.55	22.74	0.16	0.09	127.03	126.93	0.51	0.04	0.02
Forest 2006	126.46	121.21	24.59	0.16	0.10	126.30	126.22	0.37	0.05	0.02
Forest 2007	125.67	119.68	26.44	0.17	0.10	125.48	125.45	0.13	0.04	0.02
Forest 2008	124.83	118.05	28.16	0.17	0.10	124.61	124.60	0.03	0.04	0.02
Forest 2009	123.81	116.00	30.34	0.18	0.11	123.55	123.56	-0.03	0.04	0.02
Forest 2010	122.93	114.56	30.80	0.18	0.11	122.64	122.58	0.20	0.05	0.02
Elevation	92.96	49.41	25.89	0.21	0.12	88.79	82.32	3.88	0.17	0.09
Slope	3.44	2.06	24.49	0.09	0.04	3.34	3.21	2.35	0.19	0.09
Distance roads	58.53	27.42	69.53	0.41	0.25	57.24	58.04	-1.77	0.18	0.09
Distance cities	1873.79	898.77	78.12	0.41	0.25	1829.01	1867.37	-3.09	0.15	0.04
Distance palm oil	63.94	73.52	-16.40	0.20	0.12	63.45	60.67	4.71	0.12	0.07
Distance timber	150.56	89.51	41.61	0.29	0.15	142.98	133.61	6.55	0.13	0.07
Distance logging	46.85	114.19	-111.88	0.22	0.12	47.89	46.62	2.06	0.11	0.05
AGC	124.77	72.15	102.93	0.45	0.26	122.32	123.65	-2.60	0.09	0.02
Population 2005	17.09	397.29	-229.70	0.41	0.25	17.97	21.50	-2.07	0.07	0.03
Population 2010	16.27	399.93	-212.99	0.42	0.26	17.10	21.07	-2.14	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.20:** Peatland forest, > 60% forest cover in 2005, matching up to 2004

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2000</b>	128.57	125.12	17.59	0.14	0.07	128.43	128.56	-0.63	0.03	0.01
<b>Forest 2001</b>	128.34	124.66	18.68	0.14	0.08	128.20	128.32	-0.64	0.03	0.01
<b>Forest 2002</b>	128.05	124.16	19.62	0.15	0.08	127.90	128.03	-0.63	0.03	0.01
<b>Forest 2003</b>	127.89	123.88	20.18	0.15	0.08	127.74	127.86	-0.59	0.04	0.01
<b>Forest 2004</b>	127.54	123.21	21.55	0.15	0.09	127.37	127.50	-0.65	0.04	0.02
<b>Elevation</b>	92.96	49.41	25.89	0.21	0.12	88.56	80.98	4.56	0.17	0.09
<b>Slope</b>	3.44	2.06	24.49	0.09	0.04	3.33	3.16	3.07	0.19	0.09
<b>Distance roads</b>	58.53	27.42	69.53	0.41	0.25	57.21	57.96	-1.67	0.18	0.09
<b>Distance cities</b>	1873.79	898.77	78.12	0.41	0.25	1829.60	1870.73	-3.31	0.15	0.04
<b>Distance palm oil</b>	63.94	73.52	-16.40	0.20	0.12	63.42	60.84	4.36	0.12	0.07
<b>Distance timber</b>	150.56	89.51	41.61	0.29	0.15	142.77	134.00	6.14	0.13	0.07
<b>Distance logging</b>	46.85	114.19	-111.88	0.22	0.12	47.96	46.17	2.91	0.11	0.05
<b>AGC</b>	124.77	72.15	102.93	0.45	0.26	122.21	123.40	-2.33	0.09	0.02
<b>Population 2005</b>	17.09	397.29	-229.70	0.41	0.25	17.99	21.55	-2.09	0.06	0.03
<b>Population 2010</b>	16.27	399.93	-212.99	0.42	0.26	17.13	21.08	-2.13	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.21:** Dryland forest, cells above 1000m, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	109.30	103.29	14.64	0.10	0.06	109.44	108.79	1.59	0.02	0.01
<b>Forest 2006</b>	108.58	102.23	15.32	0.10	0.06	108.70	108.02	1.64	0.02	0.01
<b>Forest 2007</b>	108.02	101.31	16.06	0.11	0.06	108.13	107.49	1.54	0.02	0.01
<b>Forest 2008</b>	107.43	100.34	16.85	0.11	0.06	107.54	106.90	1.53	0.02	0.01
<b>Forest 2009</b>	106.43	98.69	18.11	0.11	0.07	106.52	105.99	1.24	0.02	0.01
<b>Forest 2010</b>	105.83	97.69	18.86	0.11	0.07	105.91	105.37	1.26	0.02	0.01
<b>Elevation</b>	1634.42	1421.82	36.25	0.21	0.12	1584.76	1644.43	-10.84	0.06	0.02
<b>Slope</b>	19.75	13.38	93.64	0.36	0.24	19.41	19.21	3.17	0.03	0.02
<b>Distance roads</b>	34.24	19.95	52.28	0.32	0.20	34.08	36.55	-9.12	0.06	0.02
<b>Distance cities</b>	1915.67	1120.87	51.68	0.41	0.21	1883.70	2111.55	-15.13	0.12	0.07
<b>Distance palm oil</b>	95.77	135.35	-37.14	0.14	0.07	92.14	91.42	0.69	0.07	0.03
<b>Distance timber</b>	154.93	147.36	5.04	0.10	0.06	144.37	132.79	7.93	0.07	0.04
<b>Distance logging</b>	76.62	157.52	-57.13	0.16	0.08	72.46	65.27	5.25	0.08	0.03
<b>AGC</b>	159.03	110.94	167.21	0.49	0.29	159.03	159.07	-0.13	0.04	0.02
<b>Population 2005</b>	14.66	180.94	-160.89	0.39	0.27	13.36	15.88	-2.82	0.08	0.04
<b>Population 2010</b>	12.61	192.48	-153.10	0.40	0.28	11.50	14.08	-2.43	0.08	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.22:** Peatland forest, cells above 1000m, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	115.92	118.13	-5.92	0.05	0.02	116.89	116.68	0.55	0.09	0.02
<b>Forest 2006</b>	115.65	117.77	-5.67	0.05	0.02	116.68	116.34	0.88	0.10	0.03
<b>Forest 2007</b>	115.43	117.44	-5.34	0.05	0.02	116.59	115.67	2.40	0.09	0.02
<b>Forest 2008</b>	115.27	117.09	-4.82	0.05	0.02	116.49	114.99	3.89	0.10	0.03
<b>Forest 2009</b>	114.96	116.58	-4.29	0.05	0.02	116.04	114.68	3.53	0.09	0.03
<b>Forest 2010</b>	114.65	116.27	-4.23	0.04	0.02	115.94	114.43	3.92	0.09	0.03
<b>Elevation</b>	1556.01	1568.82	-2.70	0.13	0.05	1505.52	1591.79	-17.44	0.19	0.09
<b>Slope</b>	22.45	13.72	99.12	0.37	0.24	20.92	21.39	-5.43	0.18	0.05
<b>Distance roads</b>	41.38	33.58	28.68	0.19	0.09	38.52	43.15	-16.53	0.19	0.05
<b>Distance cities</b>	2700.91	1808.45	49.44	0.30	0.16	2580.48	3082.23	-31.72	0.24	0.08
<b>Distance palm oil</b>	94.65	106.80	-19.80	0.13	0.04	97.98	97.90	0.13	0.19	0.07
<b>Distance timber</b>	251.93	246.14	4.68	0.17	0.06	257.39	242.91	11.92	0.16	0.06
<b>Distance logging</b>	45.27	71.69	-61.49	0.21	0.12	43.01	53.34	-37.60	0.30	0.14
<b>AGC</b>	150.48	109.44	116.33	0.47	0.29	149.33	144.63	12.80	0.22	0.09
<b>Population 2005</b>	9.07	88.07	-127.77	0.38	0.25	5.90	7.56	-12.73	0.17	0.07
<b>Population 2010</b>	8.90	116.32	-126.65	0.39	0.26	5.84	6.49	-4.78	0.12	0.05

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.23:** Dryland forest, all-elevation cells, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	109.15	95.90	32.36	0.18	0.11	109.01	109.67	-1.60	0.01	0.01
<b>Forest 2006</b>	108.28	94.67	32.83	0.19	0.11	108.14	108.77	-1.53	0.02	0.01
<b>Forest 2007</b>	107.53	93.56	33.35	0.19	0.11	107.39	108.00	-1.47	0.02	0.01
<b>Forest 2008</b>	106.82	92.36	34.14	0.19	0.11	106.67	107.31	-1.50	0.02	0.01
<b>Forest 2009</b>	105.70	90.68	34.91	0.19	0.12	105.55	106.17	-1.46	0.02	0.01
<b>Forest 2010</b>	105.02	89.67	35.27	0.19	0.12	104.86	105.46	-1.38	0.02	0.01
<b>Elevation</b>	816.19	271.64	82.35	0.50	0.31	804.01	810.41	-0.98	0.06	0.03
<b>Slope</b>	15.10	6.53	116.55	0.51	0.32	14.99	15.00	-0.07	0.04	0.01
<b>Distance roads</b>	34.88	12.74	71.53	0.43	0.28	34.65	37.11	-7.95	0.07	0.04
<b>Distance cities</b>	1685.47	607.63	76.43	0.45	0.26	1668.69	1728.52	-4.27	0.09	0.03
<b>Distance palm oil</b>	100.39	166.04	-49.76	0.21	0.07	100.59	100.86	-0.21	0.04	0.02
<b>Distance timber</b>	133.08	129.05	2.75	0.12	0.06	132.00	115.03	11.61	0.05	0.02
<b>Distance logging</b>	84.10	226.23	-86.52	0.31	0.18	84.68	84.27	0.25	0.04	0.02
<b>AGC</b>	158.18	93.16	182.13	0.58	0.33	157.97	155.70	6.34	0.04	0.02
<b>Population 2005</b>	16.51	320.78	-240.16	0.47	0.28	16.67	18.72	-1.60	0.07	0.04
<b>Population 2010</b>	14.24	321.69	-247.64	0.49	0.29	14.39	16.65	-1.81	0.07	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.



**Table 4.D.24:** Peatland forest, all-elevation cells, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Forest 2005</b>	111.53	98.58	32.54	0.19	0.11	111.13	111.17	-0.10	0.03	0.01
<b>Forest 2006</b>	110.76	97.28	33.41	0.19	0.11	110.34	110.40	-0.16	0.03	0.02
<b>Forest 2007</b>	109.94	95.93	34.23	0.20	0.12	109.50	109.55	-0.13	0.03	0.01
<b>Forest 2008</b>	109.11	94.53	35.09	0.20	0.12	108.64	108.70	-0.12	0.03	0.01
<b>Forest 2009</b>	108.09	92.74	36.25	0.20	0.12	107.60	107.64	-0.08	0.04	0.02
<b>Forest 2010</b>	107.24	91.53	36.52	0.20	0.12	106.73	106.73	0.01	0.04	0.02
<b>Elevation</b>	116.39	57.84	21.72	0.19	0.11	112.14	101.69	3.92	0.18	0.10
<b>Slope</b>	3.67	2.06	25.93	0.10	0.04	3.56	3.31	4.04	0.19	0.09
<b>Distance roads</b>	54.08	24.21	68.89	0.40	0.26	52.93	54.08	-2.65	0.15	0.09
<b>Distance cities</b>	1757.89	790.56	76.12	0.41	0.25	1713.34	1763.31	-3.98	0.13	0.04
<b>Distance palm oil</b>	60.34	77.95	-29.17	0.17	0.10	59.81	57.08	4.48	0.13	0.08
<b>Distance timber</b>	136.64	90.13	32.32	0.24	0.12	130.31	120.73	6.83	0.14	0.08
<b>Distance logging</b>	49.56	125.17	-118.79	0.24	0.13	50.47	49.44	1.59	0.12	0.06
<b>AGC</b>	120.29	68.71	98.29	0.44	0.26	118.15	118.71	-1.07	0.08	0.02
<b>Population 2005</b>	21.62	427.10	-215.25	0.41	0.25	22.44	26.42	-2.07	0.05	0.03
<b>Population 2010</b>	20.00	433.05	-208.98	0.42	0.26	20.75	24.99	-2.10	0.06	0.03

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.25:** Dryland “intact primary” forest, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Primary forest 2005</b>	80.40	52.40	43.99	0.20	0.19	79.00	78.99	0.03	0.04	0.03
<b>Primary forest 2006</b>	80.09	52.06	44.02	0.20	0.19	78.68	78.64	0.07	0.04	0.03
<b>Primary forest 2007</b>	79.77	51.71	44.05	0.21	0.19	78.37	78.30	0.10	0.04	0.03
<b>Primary forest 2008</b>	79.49	51.37	44.14	0.21	0.19	78.08	78.00	0.13	0.03	0.03
<b>Primary forest 2009</b>	79.03	50.85	44.19	0.21	0.19	77.63	77.52	0.17	0.03	0.03
<b>Primary forest 2010</b>	78.73	50.48	44.28	0.21	0.19	77.32	77.18	0.22	0.03	0.03
<b>Elevation</b>	470.04	207.62	96.65	0.47	0.28	457.63	459.76	-0.79	0.07	0.03
<b>Slope</b>	13.14	6.15	104.81	0.47	0.29	12.85	13.14	-4.32	0.07	0.02
<b>Distance roads</b>	35.15	12.34	70.50	0.42	0.28	33.70	35.46	-5.56	0.09	0.05
<b>Distance cities</b>	1588.08	579.06	75.26	0.43	0.26	1512.67	1496.02	1.29	0.13	0.04
<b>Distance palm oil</b>	102.35	167.75	-46.31	0.20	0.07	103.96	102.84	0.78	0.07	0.04
<b>Distance timber</b>	123.84	128.03	-2.91	0.10	0.05	121.04	103.89	12.01	0.07	0.04
<b>Distance logging</b>	87.26	230.05	-82.59	0.32	0.20	90.49	93.82	-1.89	0.06	0.04
<b>AGC</b>	157.82	92.17	171.62	0.57	0.33	157.02	154.96	5.31	0.02	0.01
<b>Population 2005</b>	17.30	328.56	-229.96	0.49	0.28	18.05	21.23	-2.29	0.07	0.04
<b>Population 2010</b>	14.92	328.89	-247.47	0.50	0.29	15.59	18.89	-2.54	0.06	0.04

Notes: 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

**Table 4.D.26:** Peatland “intact primary” forest, matching up to 2010

	Unmatched					Matched				
	T	C	Diff	KS	eCDF	T	C	Diff	KS	eCDF
<b>Primary forest 2005</b>	87.07	57.25	47.97	0.22	0.20	86.12	85.76	0.59	0.02	0.01
<b>Primary forest 2006</b>	86.80	56.86	48.11	0.22	0.20	85.85	85.49	0.57	0.02	0.01
<b>Primary forest 2007</b>	86.46	56.36	48.33	0.22	0.20	85.50	85.15	0.57	0.02	0.01
<b>Primary forest 2008</b>	86.12	55.91	48.46	0.22	0.20	85.16	84.81	0.56	0.02	0.01
<b>Primary forest 2009</b>	85.67	55.32	48.63	0.23	0.20	84.70	84.33	0.60	0.02	0.01
<b>Primary forest 2010</b>	85.22	54.79	48.70	0.23	0.20	84.24	83.84	0.65	0.02	0.01
<b>Elevation</b>	87.08	48.09	23.94	0.19	0.11	83.72	75.37	5.17	0.18	0.10
<b>Slope</b>	3.29	1.98	23.62	0.08	0.03	3.22	3.03	3.44	0.19	0.09
<b>Distance roads</b>	54.33	24.15	69.26	0.40	0.26	53.10	54.31	-2.77	0.15	0.09
<b>Distance cities</b>	1738.69	783.99	76.37	0.41	0.25	1696.50	1741.37	-3.62	0.13	0.04
<b>Distance palm oil</b>	59.65	77.76	-30.13	0.17	0.10	59.18	56.66	4.15	0.13	0.08
<b>Distance timber</b>	134.29	89.12	31.52	0.24	0.12	127.92	118.59	6.69	0.14	0.09
<b>Distance logging</b>	49.65	125.51	-118.54	0.24	0.13	50.60	49.77	1.27	0.12	0.06
<b>AGC</b>	119.68	68.44	97.41	0.43	0.26	117.55	118.10	-1.05	0.08	0.02
<b>Population 2005</b>	21.88	429.29	-214.34	0.41	0.25	22.72	26.74	-2.07	0.05	0.03
<b>Population 2010</b>	20.22	435.10	-208.19	0.42	0.26	21.00	25.39	-2.16	0.06	0.03

*Notes:* 1000 bootstrap samples. T: treated group. C: control group. Diff: absolute standardised difference.

KS: Kolmogorov-Smirnov test statistic. eCDF: mean difference in empirical CDFs.

## 4.E Robustness

### 4.E.1 Sensitivity to matching procedures

We conduct additional sensitivity analyses on the assumptions used in propensity score matching estimates and a robustness analysis of the central NP-DDD on alternative approaches to matching procedures that do not rely on the propensity score. In part these test the possibility that propensity score matching may introduce biases due to the way in which it reduces the dimensionality of the matching problem to matching on a single dimension: the propensity score (e.g. King and Nielsen, 2019).

On the sensitivity of the propensity score estimates to the parameter choices in implementation, the analyses are: (1)  $k > 1$  nearest neighbours analysis in which we test the sensitivity of our results to matching on  $k = 2, 3, 5$  nearest neighbours; and, (2) caliper choices, where we test calipers of 0.01 and 0.001 in one-to-one nearest neighbour matching. With regard to alternative, non-propensity score, matching estimators we undertake Coarsened Exact Matching (CEM) (King and Zeng, 2006; Iacus *et al.*, 2012). The use of CEM means that the sample sizes are slightly more sensitive to particular choices of matching variables and the coarseness of the matching. For this reason we undertake four separate CEM approaches with various degrees of coarseness to understand these sensitivities. We use the STATA 16 *cem* routine and default coarseness settings unless otherwise stated (See table notes).

Tables 4.B.13 and 4.B.14 shows the results of the sensitivity analyses on dryland and peatland forests. The results show that the main results are not being driven by the choice of the propensity score matching estimator over CEM. Tables 4.B.13 and 4.B.14 show that the CEM estimates do not differ to any great extent for peatland and dryland forest. Furthermore, Tables 4.B.17 and 4.B.18 show that the main results are robust to the number of nearest neighbours chosen for 2,3 and 5 neighbours. Finally, our choice of sample removed areas above 1,000 m for agronomic reasons. Tables 4.B.17 and 4.B.18 also show that this selection of the sample is not driving the results either. Overall, these results are strongly supportive of our estimation approach leading to our main results.

### Testing the SUTVA assumption using Regression Discontinuity: ‘Leakage’?

A final identification assumption of the DD and DDD estimators is the Stable Unit Treatment Value Assumption (SUTVA): the treatment causes leakage to untreated forest areas via spatial, behavioural or general equilibrium effects, a common confounder in the evaluation of area-based policies (Gaveau *et al.*, 2009; Nelson and Chomitz, 2011; Andam *et al.*, 2008; Joppa and Pfaff, 2010a; Joppa and Pfaff, 2010b). If protection via the Moratorium induces the displacement of forest clearing to outside the Moratorium’s boundaries, deforestation rates inside and outside these boundaries are contemporaneously affected in opposite directions resulting in treatment effects of a higher magnitude than the “true” effects. To check for leakage, we use a regression discontinuity design (Lee and Lemieux, 2010; Calonico *et al.*, 2014) with a sharp

cut-off at the Moratorium’s boundaries. Regression discontinuity estimates the LATE of the Moratorium in the proximity of its boundaries with non-Moratorium land. In particular, we estimate the following equation (following Burgess *et al.* (2019)):

$$D_i = \alpha + \beta_1 \text{Moratorium}_i + \beta_2 X_i + f(\text{Distance}_i) + \epsilon_i$$

where  $D_i$  is deforestation (in pixels) in grid cell  $i$ ;  $\text{Moratorium}_i$  is a binary indicator representing treatment assignment,  $X_i$  is a vector of cross-sectional control variables, namely slope, elevation, distance from cities and roads, and above-ground biomass content, and

$$f(\text{Distance}_i) = \text{Moratorium}_i * f_{\text{Moratorium}}(\text{Distance}_i) + \\ + (1 - \text{Moratorium}_i) f_{\text{Non-Moratorium}}(\text{Distance}_i)$$

is a polynomial term which weighs observations with respect to their distance from the Moratorium’s boundaries.

We specify separate linear polynomials on both sides of the boundary cut-off (following Gelman and Imbens (2019) and Burgess *et al.* (2019)), and estimate treatment effects via OLS regressions with robust standard errors clustered at the district level. Our preferred results are obtained via separate linear polynomials and optimal bandwidth selection through the Calonico *et al.* (2014) method. In Fig. 4.E.1, we also present results using the Imbens and Kalyanaraman (2012) bandwidth selection algorithm (panel B).


In Fig. 4.E.2-4.E.6, we visualise the full dataset employed in the spatial regression discontinuity designs, letting the bandwidth size vary according to either the Calonico *et al.* (2014) or the Imbens and Kalyanaraman (2012) method. First, we report two plots which relate the forcing variable, distance from the Moratorium border, to forest cover in pixels, for 2004 and 2010 (Fig. 4.E.2). We treat 2004 as the “base year” since we estimate regression discontinuity LATEs for 2005-2018. As shown in Fig. 4.E.2 (left panel), forest cover is not distributed smoothly around the threshold at the baseline, with Moratorium grid cells discontinuously more forested than the surrounding areas. Fig. 4.E.2 (right panel) confirms the trend for 2010, the actual baseline year prior to the implementation of the Moratorium. These descriptive results are a confirmation of the economic marginality of Moratorium grid cells, which tended to be cleared at a slower rate than non-Moratorium grid cells in the immediate vicinity, even before the Moratorium was implemented.

In Fig. 4.E.3 and 4.E.4, we report the discontinuity in deforestation rates around Moratorium boundaries, restricted to the bandwidth level calculated with our preferred Calonico *et al.* (2014) procedure. Fig. 4.E.3 examines pre-treatment years while Fig. 4.E.4 visualises data for 2011-2018. Clearly, deforestation is persistently higher outside of Moratorium boundaries, starting from 2005 and up until 2018, with only slight changes in the underlying data distribution. In Fig. 4.E.5 and 4.E.6, we report analogous regression discontinuity plots using the Imbens and Kalyanaraman (2012) optimal bandwidth selection procedure, which yield similar results: deforestation rates are always higher outside the Moratorium, enhancing our confidence in

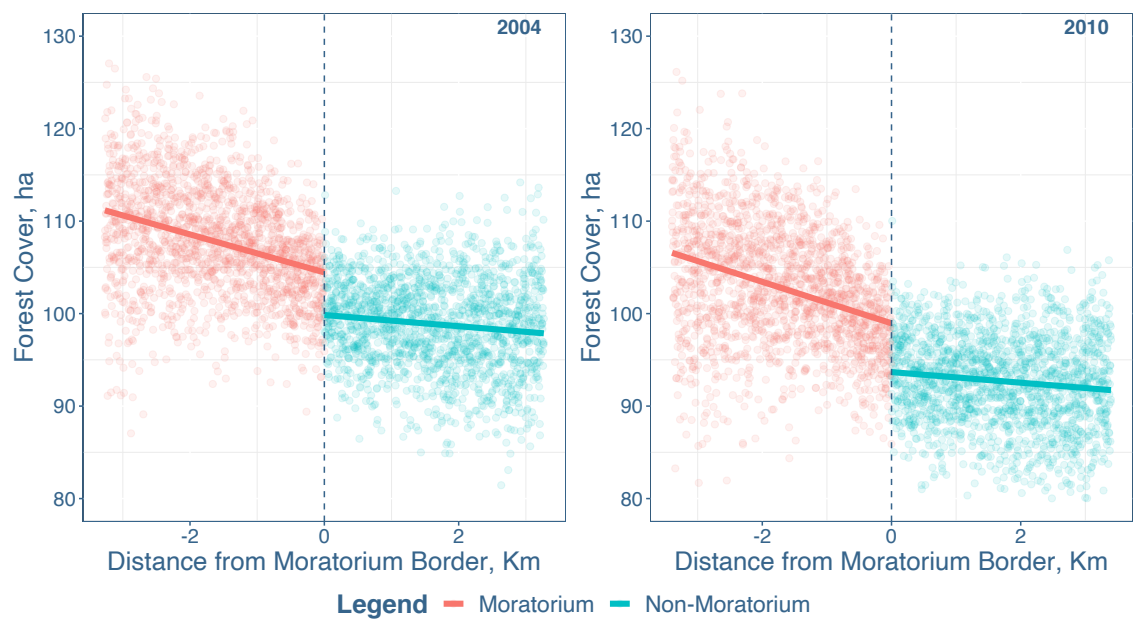
ruling out the occurrence of spatial leakage.

We perform several robustness checks to assess the validity of our regression discontinuity estimates, which rule out the potential role of leakage in our results. In particular, we first run regressions with alternative bandwidths, namely: a 5 km fixed bandwidth; a 10 km fixed bandwidth; and, a 20 km fixed bandwidth (reported in the top row of Fig. 4.E.7). Given the spatial configuration of the Moratorium (Fig. 4.2.2), our bandwidths cover a large extent of non-Moratorium cells in which leakage could feasibly occur. One reason why we might not observe a leakage effect is because the Moratorium's boundaries have changed after forest areas dropped out of the Moratorium. This can be ruled out, at least partially, by the application of these different bandwidths.


The sensitivity of our results is tested with respect to alternative polynomial specifications, employing a quadratic function of distance on both sides of the cut-off using: the [Imbens and Kalyanaraman \(2012\)](#) optimal bandwidth; a 10 km fixed bandwidth; and, the [Calonico \*et al.\* \(2014\)](#) optimal bandwidth. These results are reported in the bottom row of Fig. 4.E.7. Notably, the 95% confidence interval for the regression discontinuity LATE always includes zero, thereby reassuring us about the stability of our estimates with respect to alternative specifications. We are therefore able to rule out the role of leakage in affecting our main results.



**Figure 4.E.1:** Regression discontinuity LATE with [Calonico \*et al.\* \(2014\)](#) bandwidth (A); regression discontinuity LATE with [Imbens and Kalyanaraman \(2012\)](#) bandwidth (B).

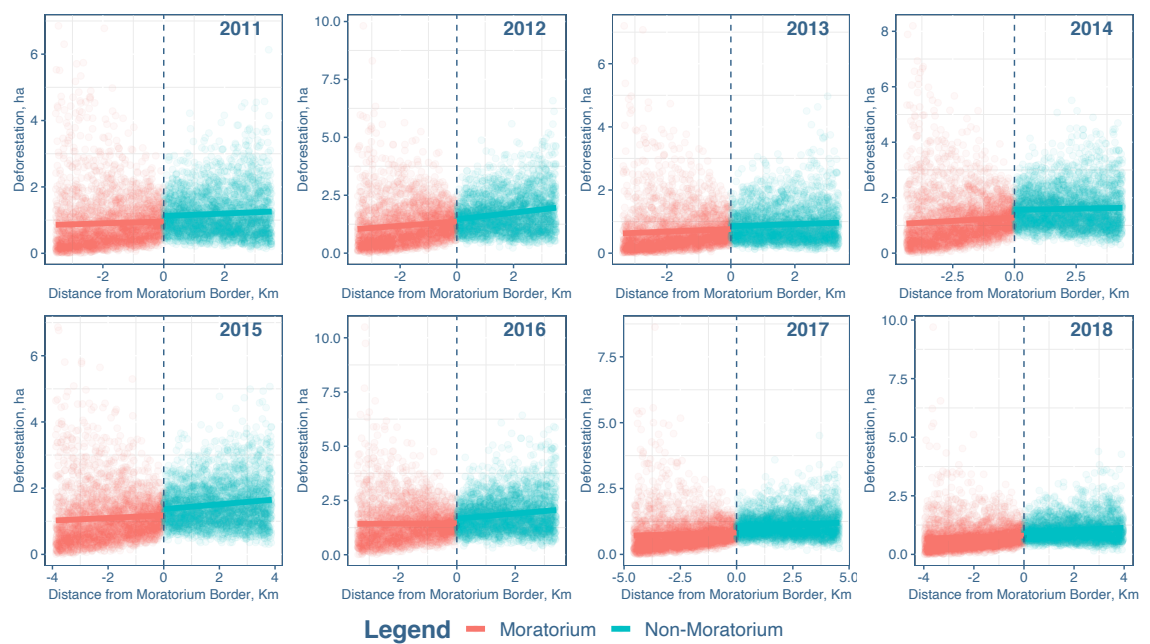


**Figure 4.E.2:** Forest cover in (A) 2004 and (B) 2010, with respect to distance from Moratorium boundaries

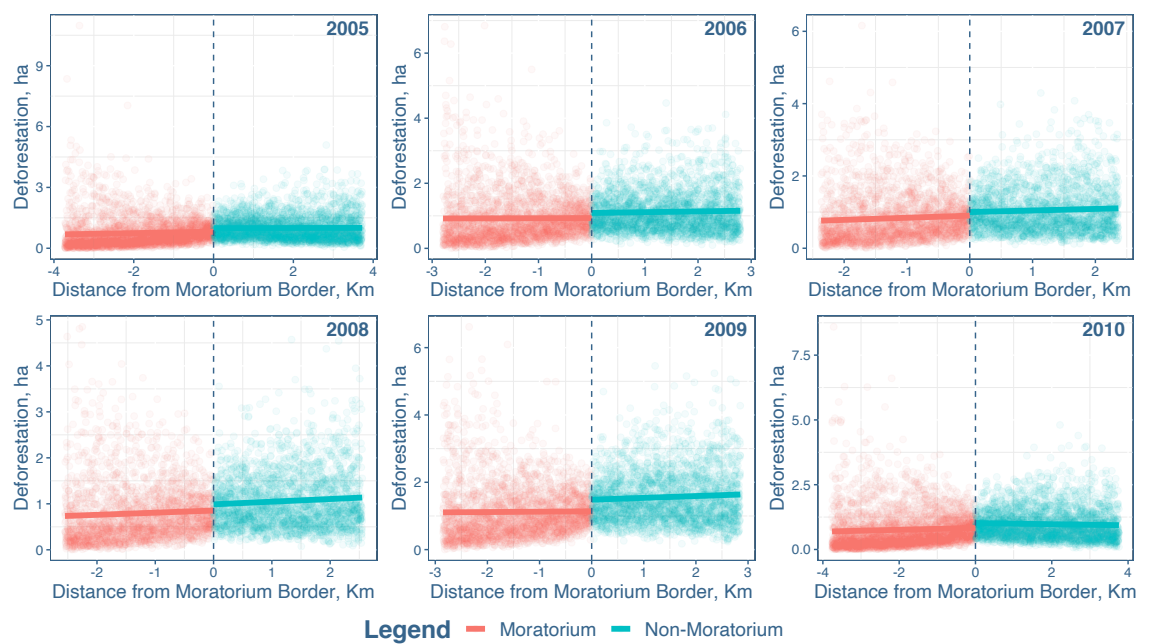


**Figure 4.E.3:** Deforestation in the proximity of Moratorium boundaries, *Calonico et al. (2014)* bandwidth, pre-treatment years.

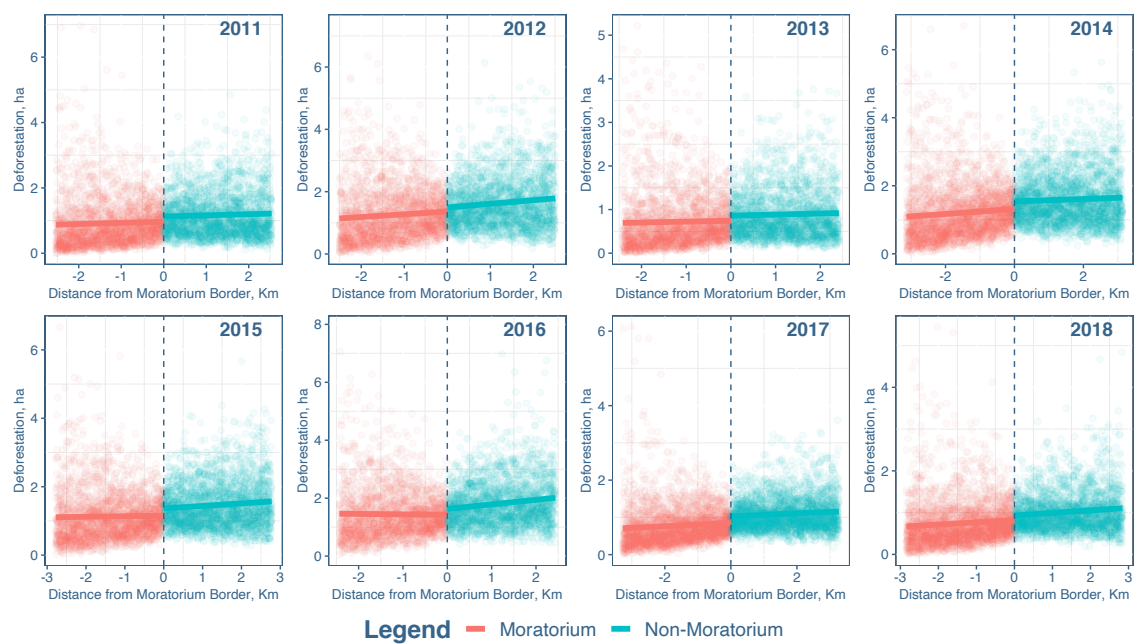




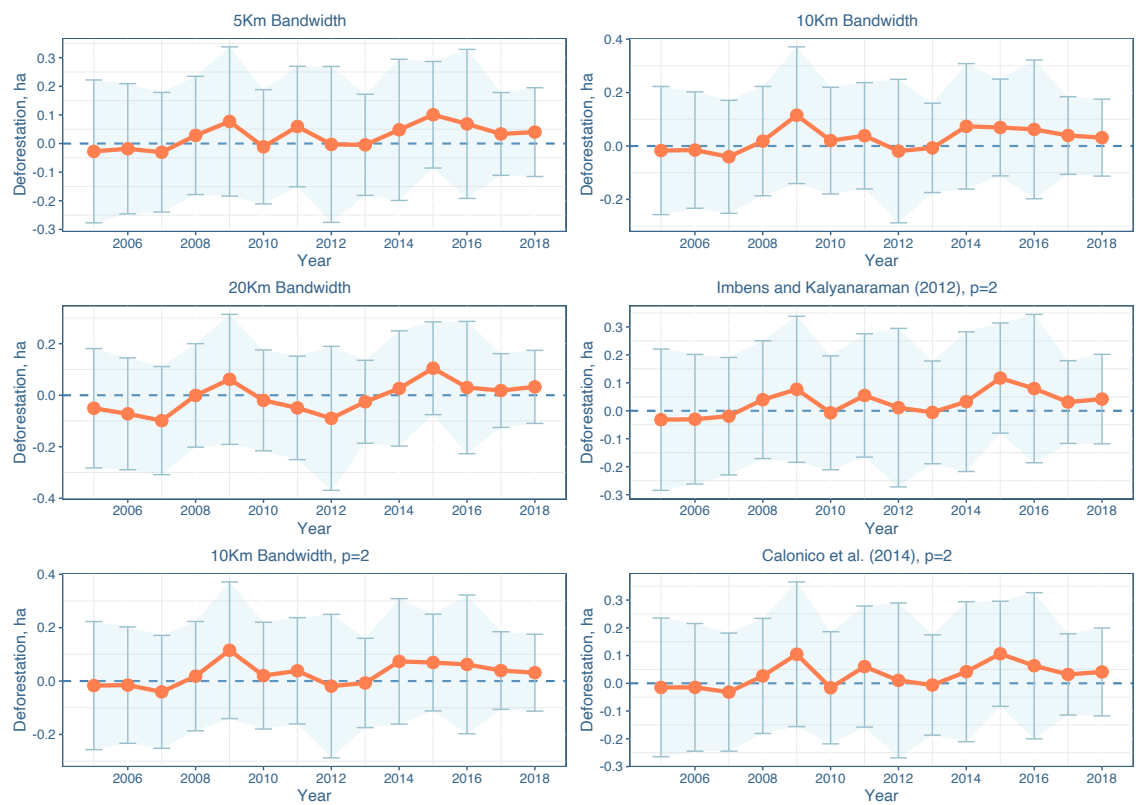
**Figure 4.E.4:** Deforestation in the proximity of Moratorium boundaries, *Calonico et al. (2014)* bandwidth, post-treatment years.



**Figure 4.E.5:** Deforestation in the proximity of Moratorium boundaries, [Imbens and Kalyanaraman \(2012\)](#) bandwidth, pre-treatment years.



**Figure 4.E.6:** Deforestation in the proximity of Moratorium boundaries, [Imbens and Kalyanaraman \(2012\)](#) bandwidth, post-treatment years.



**Figure 4.E.7:** Regression discontinuity robustness checks

# References

- Abadie, Alberto and Imbens, Guido W (2006). “Large sample properties of matching estimators”. *Econometrica* 74.1, pp. 235–267.
- Abadie, Alberto and Imbens, Guido W (2008). “Notes and comments on the failure of the bootstrap for matching estimators”. *Econometrica* 76.6, pp. 1537–1557.
- Abood, Sinan A, Ser, Janice, Lee, Huay, Burivalova, Zuzana, Garcia-ulloa, John, and Koh, Lian Pin (2015). “Relative contributions of the logging, fiber, oil palm, and mining industries to forest loss in Indonesia”. *Conservation Letters* 8, pp. 58–67.
- Albers, H. J. and Ferraro, P. (2006). “The Economics of Terrestrial Biodiversity Conservation in Developing Nations”. *Economic Development and Environmental Sustainability*. Ed. by R. Lòpez and M. Toman. Oxford: Oxford University Press.
- Alisjahbana, Armida S. and Busch, Jonah M. (2017). “Forestry, Forest Fires, and Climate Change in Indonesia”. *Bulletin of Indonesian Economic Studies* 53, pp. 111–136.
- Andam, Kwaw S., Ferraro, Paul J., Pfaff, Alexander, Sanchez-Azofeifa, G. Arturo, and Robalino, Juan A. (2008). “Measuring the effectiveness of protected area networks in reducing deforestation”. *Proceedings of the National Academy of Sciences* 105, pp. 16089–16094.
- Austin, Kemen G., Kasibhatla, Prasad S., Urban, Dean L., Stolle, Fred, and Vincent, Jeffrey (2015). “Reconciling oil palm expansion and climate change mitigation in Kalimantan, Indonesia”. *PLOS ONE* 10.5, pp. 1–17.
- Austin, Kemen G., Schwantes, Amanda, Gu, Yaofeng, and Kasibhatla, Prasad S. (2019). “What causes deforestation in Indonesia?” *Environmental Research Letters* 14.2, p. 024007.
- Austin, Peter C. (2009). “Some methods of propensity-score Matching had superior performance to others: results of an empirical investigation and Monte Carlo simulations”. *Biometrical Journal* 51.1, pp. 171–184.
- Burgess, Robin, Costa, Francisco, and Olken, Ben (Aug. 2019). *The Brazilian Amazon’s double reversal of fortune*. Working Paper.
- Caliendo, Marco and Kopeinig, Sabine (2008). “Some practical guidance for the implementation of propensity score matching”. *Journal of Economic Surveys* 22.1, pp. 31–72.
- Calonico, Sebastian, Cattaneo, Matias D., and Titiunik, Rocio (2014). “Robust non-parametric confidence intervals for regression-discontinuity designs”. *Econometrica* 82.6, pp. 2295–2326.
- Chabé-Ferret, Sylvain (2017). *Should We Combine Difference In Differences with Conditioning on Pre-Treatment Outcomes?* Tech. rep. Toulouse: Toulouse School of Economics, INRA and IAST.

- Chabé-Ferret, Sylvain and Subervie, Julie (2013). “How much green for the buck? Estimating additional and windfall effects of French agro-environmental schemes by DID-matching”. *Journal of Environmental Economics and Management* 65.1, pp. 12–27.
- Climate Home News (2021). *Indonesia ends forest protection deal with Norway, raising deforestation fears*. URL: <https://www.climatechangenews.com/2021/09/20/indonesia-ends-forest-protection-deal-norway-raising-deforestation-fears/>.
- Daw, Jamie R. and Hatfield, Laura A. (2018). “Matching in difference-in-differences: between a rock and a hard place”. *Health Services Research* 53.6, pp. 4111–4117.
- Enrici, Ashley and Hubacek, Klaus (2016). “Business as usual in Indonesia: governance factors effecting the acceleration of the deforestation rate after the introduction of REDD+”. *Energy, Ecology and Environment* 1, pp. 183–196.
- FAO (2015). *Global Forest Resources Assessment 2015 Country Report: Indonesia*. Country Report. Food and Agriculture Organization of the United Nations, Rome.
- Forest Watch Indonesia (2015). *The State of Forest Report in Indonesia*. Brussels presentation, May 2015. Forest Watch Indonesia, Bogor, Indonesia.
- Gaveau, David L.A., Epting, Justin, Lyne, Owen, Linkie, Matthew, Kumara, Indra, Kanninen, Markku, and Leader-Williams, Nigel (2009). “Evaluating whether protected areas reduce tropical deforestation in Sumatra”. *Journal of Biogeography* 36.11, pp. 2165–2175.
- Gelman, Andrew and Imbens, Guido (2019). “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs”. *Journal of Business & Economic Statistics* 37.3, pp. 447–456.
- Glauber, Ann Jeannette, Moyer, Sarah, Adriani, Magda, and Gunawan, Iwan (2016). *The Cost of Fire : An Economic Analysis of Indonesia’s 2015 Fire Crisis*. Tech. rep. 1. World Bank, Jakarta, Indonesia.
- Government of Indonesia (2018). *Indonesia Second Biennial Update Report Under the United Nations Framework Convention on Climate Change*. Directorate General of Climate Change, Ministry of Environment and Forestry, Government of Indonesia report.
- Government of Indonesia (2016). *National Forest Reference Emission Level for Deforestation and Forest Degradation: Submission by Indonesia*. Post Technical Assessment by UNFCCC. Directorate General of Climate Change. The Ministry of Environment and Forestry, Government of Indonesia.
- Gruber, Jonathan (1994). “The Incidence of Mandated Maternity Benefits”. *The American Economic Review* 84.3, pp. 622–641.
- Hansen, M. C. *et al.* (2013). “High-Resolution Global Maps of 21st-Century Forest Cover Change”. *Science* 342.6160, pp. 850–853.
- Heckman, James J., Ichimura, Hidehiko, and Todd, Petra (1998). “Matching As An Econometric Evaluation Estimator”. *Review of Economic Studies* 65.2, pp. 261–294.
- Heckman, James J., Ichimura, Hidehiko, and Todd, Petra (1997). “Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme”. *Review of Economic Studies* 64, pp. 605–654.

- Iacus, Stefano M, King, Gary, and Porro, Giuseppe (2012). “Causal inference without balance checking: Coarsened Exact Matching”. *Political analysis* 20.1, pp. 1–24.
- Imai, Kosuke and Kim, In Song (2021). “On the Use of Two-Way Fixed Effects Regression Models for Causal Inference with Panel Data”. eng. *Political analysis* 29.3, pp. 405–415. ISSN: 1047-1987.
- Imbens, Guido (2014). *Matching Methods in Practice: Three Examples*. Working Paper 19959. National Bureau of Economic Research.
- Imbens, Guido and Kalyanaraman, Karthik (2012). “Optimal Bandwidth Choice for the Regression Discontinuity Estimator”. *The Review of Economic Studies* 79.3, pp. 933–959.
- Indrarto, G.B., Murharjanti, P., Khatarina, J., Pulungan, I., Ivalerina, F., Rahman, J., Prana, M. N., Resosudarmo, I. A. P., Muharrom, E., and Justitia, R. (2012). *The context of REDD+ in Indonesia: drivers, agents and institutions*. 92. Center for International Forestry Research (CIFOR), Bogor, Indonesia.
- Joppa, Lucas N and Pfaff, Alexander (2010a). “Global protected area impacts”. *Proceedings of the Royal Society B* 278, 1633–1638.
- Joppa, Lucas and Pfaff, Alexander (2010b). “Reassessing the forest impacts of protection”. *Annals of the New York Academy of Sciences* 1185.1, pp. 135–149.
- King, Gary and Nielsen, Richard (2019). “Why propensity scores should not be used for matching”. *Political Analysis* 27.4, pp. 435–454.
- King, Gary and Zeng, Langche (2006). “The dangers of extreme counterfactuals”. *Political Analysis* 14.2, pp. 131–159.
- Krisnawati, Haruni, Kallio, Maarit, and Kanninen, Markku (2011). *Acacia Mangium Willd.: ecology, silviculture and productivity*. Tech. rep. Bogor, Indonesia: CIFOR.
- LTS International (2018). *Third independent review of the Indonesia-Norway cooperation on reducing greenhouse gas emissions from REDD+*.
- Lee, David S. and Lemieux, Thomas (2010). “Regression discontinuity designs in economics”. *Journal of Economic Literature* 48.2, pp. 281–355.
- Lee, Myoung-jae (2005). *Micro-Econometrics for Policy, Program and Treatment Effects*. Advanced Texts in Econometrics. Oxford: Oxford University Press. ISBN: 9780199267699.
- Margono, B. A., Turubanova, S., Zhuravleva, I., Potapov, P., Tyukavina, A., and Baccini, A (2012). “Mapping and monitoring deforestation and forest degradation in Sumatra (Indonesia) using Landsat time series data sets from 1990 to 2010”. *Environmental Research Letters* 7.3, p. 34010.
- Margono, Belinda Arunarwati, Potapov, Peter V., Turubanova, Svetlana, Stolle, Fred, and Hansen, Matthew C. (2014). “Primary forest cover loss in Indonesia over 2000–2012”. *Nature Climate Change* 4, pp. 730–735.
- Ministry of Forestry of the Republic of Indonesia (2011). *Indicative Map of New Revised Permit Delays* 5. URL: <http://appgis.dephut.go.id/appgis/download.aspx>.
- MoEF (2019). *Emissions Reduction Report for the Indonesia-Norway Partnership*. government report. Directorate General of Climate Change, Ministry of Environment and Forestry, Republic of Indonesia.

- Mongabay (2020). *Experts question integrity of Indonesia’s claim of avoided deforestation*. URL: <https://news.mongabay.com/2020/09/green-climate-fund-indonesia-redd-deforestation>.
- Mongabay (2019). *Indonesia forest-clearing ban is made permanent, but labeled “propaganda”*. URL: <https://news.mongabay.com/2019/08/indonesia-forest-clearing-ban-is-made-permanent-but-labeled-propaganda/>.
- Mongabay (2021). *Indonesia terminates agreement with Norway on \$1b REDD+ scheme*. URL: <https://news.mongabay.com/2021/09/indonesia-terminates-agreement-with-norway-on-1b-redd-scheme/>.
- Nelson, Andrew and Chomitz, Kenneth M. (2011). “Effectiveness of strict vs. multiple use protected areas in reducing tropical forest fires: A global analysis using matching methods”. *PLoS ONE* 6.8.
- Olden, Andreas and Møen, Jarle (2020). *The Triple Difference Estimator*. Discussion Papers 2020/1. Norwegian School of Economics, Department of Business and Management Science.
- Ravallion, Martin, Galasso, Emanuela, Lazo, Teodoro, and Philipp, Ernesto (2005). “What Can Ex-Participants Reveal about a Program’s Impact?” *Journal of Human Resources* 40.1.
- Ryan, Andrew M., Evangelos, Kontopantelis, Ariel, Linden, and Burgess, James F. (2018). “Now trending: Coping with non-parallel trends in difference-in-differences analysis”. *Statistical Methods in Medical Research* 28.12, pp. 3697–3711.
- Sloan, Sean, Edwards, David P., and Laurance, William F. (2012). “Does Indonesia’s REDD+ moratorium on new concessions spare imminently threatened forests?” *Conservation Letters* 5.3, pp. 222–231.
- Steinweg, T., Kuepper, B, and Piotrowski, M. (2019). *28 Percent of Indonesia’s Palm Oil Landbank Is Stranded*. Working Paper. Chain Reaction Research, Washington D.C.
- Turubanova, S., Popapov, P. V., Tyukavina, A., and Hansen, M. C. (2018). “Ongoing primary forest loss in Brazil, Democratic Republic of the Congo, and Indonesia”. *Environmental Research Letters* 13, p. 074028.
- Wooldridge, J.M. (2021). *Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Event Study Estimators*. eng.
- Wooldridge, Jeffrey M. (2010). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.