Essays in Environmental and Resource Economics

by

Daniel Alexander Kannell

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael L. Anderson, Chair
Professor Brian Wright
Professor Solomon Hsiang

Fall 2022

Essays in Environmental and Resource Economics

by

Daniel Alexander Kannell

Abstract

Essays in Environmental and Resource Economics

by

Daniel Alexander Kannell

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Michael L. Anderson, Chair

The fields of environmental and resource economics provide us with the tools that allow us to better understand the world around us and the ways in which humans interact with it. People benefit from goods and resources provided by the environment, but can also harm the environment through externalities created by their economic activity. The utilization of environmental resources often requires government and society to implement policies in order to protect and preserve these resources for the future. This dissertation explores three topics in environmental and resource economics. The first measures an environmental externality caused by urban development and proposes a policy that balances continued development with the welfare of those harmed by the externality. The second explores how people adapt to changes in climate resulting from migration, which may inform us about adaptation to climate change and migration in the future. The third topic studies the effectiveness of policies aimed at protecting marine resources, at a time when marine protection is rapidly expanding.

The first chapter, titled "Shadow Prices: Measuring the Cost of Shadows from New Construction in New York City," studies one negative externality of urbanization – the blocking of sunlight by construction – and a policy that can be implemented to ensure a balance between urban development and the welfare of those who are harmed by their reduced access to sunlight. I begin by measuring the externality of urban shadows by estimating the impact of shadows created by new highrise construction on nearby housing prices. Making use of publicly available housing transactions data and building shapefiles for New York City, I create a shadow accumulation model to measure the amount of shadow created by highrises that are cast on residential units in Manhattan in each year for highrises constructed between 2005 and 2014. I then use a differenced regression model with spatial-time fixed effects to estimate the effect of a change in shadows on nearby residential unit sale prices. I estimate that a 10 percentage point increase in average daily shadow received by a unit (e.g. 1 hour of additional shadow in a day with 10 hours of sunlight) is associated with an approximately

3.78% decrease in unit price. Finally, I propose a policy that incorporates this estimate in regulating building height.

In the second chapter, titled "Estimating the Relationship Between Inter-Climate Migration and Air Conditioning Adoption," which I coauthored with Léopold Biardeau, we measure the relative increase in residential air-conditioning (AC) adoption rates in states that see higher levels of migration from relatively cooler states, within the contiguous United States. We consider how the average percent increase in Cooling Degree Days (CDDs) by migrants to a given destination state increases the average rate of residential AC adoption in the four decades spanning from the 1960s to the end of the 20^{th} century. We find significant positive effects of the percent increase in CDDs experienced by migrants on rates of air conditioning adoption. To confirm the validity of this relationship, we rely on an instrumental variables approach using origin state determinants of emigration, along with the distance between states, to provide evidence that this result is not the byproduct of a reverse causality relationship in which higher residential AC-adoption levels would be responsible for increased immigration from relatively cooler states. These results provide some insights regarding the expected impact relative temperature changes may have on Climate Change adaptation. In particular, we might expect that an increase in population displacements from warmer countries to relatively cooler ones lead to a lower increase in energy demand for cooling purposes that what would have been anticipated.

In the third and final chapter, titled "Evaluating the Effectiveness of Very Large Marine Protected Areas at Deterring Fishing Effort," which I coauthored with Léopold Biardeau and David Zilberman, we study the extent to which Very Large Marine Protected Areas (VLMPAs) have been successful at deterring fishing effort. The last decade has witnessed a considerable increase in the designation of VLMPAs, Marine Protected Areas spanning at least 100,000 km^2. On paper, these protected areas offer conservation benefits not seen in smaller ones. Yet, their large sizes may constitute a challenge for enforcement. Using on satellite-based data that tracks vessel fishing hours, we find evidence that VLMPAs have, on average, been able to deter fishing effort, although a case-by-case analysis reveals varying levels of success. To better understand the nature of possible illegal fishing in these VLMPAs, we investigate the characteristics of the vessels infringing on the fishing bans in these VLMPAs and find that most of the infractions can be traced back to a few industrialized countries.

These three chapters fall back on three important question in environmental and resource economics: What are the environmental externalities associated with economic activity and what policies can be implemented to compensate those harmed while balancing the interest of allowing continued economic activity? How do people adapt to changes in their environment? And finally, how successful are current policies that are designed to protect the environment and its resources? The research contained in this dissertation applies each of these questions in different contexts, and in doing so, helps us to better understand the economics of resources and the environment.

To Lorraine Rabinowitz

for your support during the first three years of my PhD program and throughout my life prior to graduate school.

and to Alexandra Kannell

for your continued support.

Contents

Contents **ii**

List of Figures **iv**

List of Tables **v**

1 Shadow Prices: Measuring the Cost of Shadows from New Construction in New York City **1**
1.1 Introduction . . . 1
1.2 Data . . . 3
1.3 Empirical Strategy . . . 9
1.4 Results . . . 14
1.5 Policy Implications . . . 20
1.6 Conclusion . . . 32

2 Estimating the Relationship Between Inter-Climate Migration and Air Conditioning Adoption **33**
2.1 Introduction . . . 33
2.2 Data . . . 36
2.3 Empirical Model . . . 37
2.4 Estimation Results . . . 41
2.5 Conclusion . . . 47

3 Evaluating the Effectiveness of Very Large Marine Protected Areas at Deterring Fishing Effort **51**
3.1 Introduction . . . 51
3.2 Data . . . 54
3.3 Estimating the Average Reduction in Fishing Effort . . . 59
3.4 Investigating Displacement of Fishing Effort . . . 64
3.5 Analyzing the Profile of Infringing Vessels . . . 69
3.6 Discussion . . . 69

Bibliography **72**

A Chapter 1 Appendix **79**
A.1 Additional tables and figures 79

B Chapter 3 Appendix **86**
B.1 Additional tables and figures 86
B.2 VLMPA Creation Timelines 93

List of Figures

1.1 Shadow Projection From the New York Federal Reserve 6
1.2 Throwing Shade: Shadowing Another Building 8
1.3 Vertical Shadow Measurement 9
1.4 Locations of new buildings over 50 feet tall constructed between 2005 and 2014 . 12
1.5 500m Directional Quadrants Centered on NY Federal Reserve 18
1.6 Central Park with 500 meter buffer 20
1.7 Graph of marginal net negative externality 23
1.8 Graph of marginal profits 24
1.9 Finding optimal height 25
1.10 FAR to the left of the optimal regulation 27
1.11 FAR to the right of the optimal regulation 28

2.1 Evolution of residential AC adoption rates in the contiguous US between 1960 and 2000 34
2.2 Distributions of coefficient estimates for migrant share of the population from 10,000 randomizations 48
2.3 Distributions of coefficient estimates for migration-induced relative change in CDDs from 10,000 randomizations 49

3.1 Evolution of tracked vessels and fishing hours and expansion of VLMPAs over time 53
3.2 Eight Very Large Marine Protected Areas (VLMPAs) in sample 54
3.3 Evolution of daily hours of fishing effort per 1,000 km^2 in each VLMPA 57
3.4 Evolution of average daily hours of fishing effort per 1,000 km^2 in all VLMPAs . 58
3.5 Displacement of fishing effort 65
3.6 Quartic polynomial curves of best fit and their associated 95% confidence intervals 67

A.1 Graph of downward-U shaped marginal profits 84
A.2 Finding optimal height with downward-U shaped marginal profits 85

B.1 Displacement of fishing effort, all vessels 91
B.2 Evolution of the number of vessels by country 92

List of Tables

1.1 Effect of changes in shadow on changes in housing price 15
1.2 Effect of changes in shadow relative to baseline shadow on changes in housing price 16
1.3 Height of residential buildings by direction from new buildings 19
1.4 Effect of changes in shadow on changes in housing price, excluding units within 500 meters of Central Park 21

2.1 Effect of migrant share and migration-induced relative change in CDDs on AC adoption 42
2.2 Granger causality test for effect of AC adoption on the migration-induced relative change in CDDs 44
2.3 Effect of migrant share and migration-induced relative change in CDDs and HDDs on AC adoption 45

3.1 Characteristics of the eight VLMPAs under consideration. 56
3.2 Breakdown of fishing effort in each of the 8 VLMPAs 60
3.3 Effect of VLMPA announcement and implementation on average daily fishing . 63
3.4 Area under the 4^{th}-order polynomial curve estimating the difference in annualized fishing hours 68
3.5 Ranking of the countries associated with prohibited fishing inside of the established VLMPAs 70

A.1 Effect of changes in shadow on changes in housing price with 500 meter fixed effects and cluster grid 79
A.2 Effect of changes in shadow relative to baseline shadow on changes in housing price with 500 meter fixed effects and cluster grid 80
A.3 Effect of changes in shadow on changes in housing price with 500 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park . . . 81
A.4 Effect of changes in shadow relative to baseline shadow on changes in housing price with 250 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park 82
A.5 Effect of changes in shadow relative to baseline shadow on changes in housing price with 500 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park 83

B.1 Effect of VLMPA announcement and implementation on average daily fishing, all donut sizes, no controls 87
B.2 Effect of VLMPA announcement and implementation on average daily fishing, all donut sizes, with controls 88
B.3 Areas under three distinct order polynomial curves estimating the difference in annualized fishing hours 89

Acknowledgments

I would like to start by acknowledging those who helped me in my work and research in graduate school. First, I would like to thank my co-authors who worked with me on our shared research: Léopold Biardeau and David Zilberman. I would also like to thank those professors who helped me in preparing my research, through advising, serving on my committees, or attending my oral exam back when my shadows project was still in its infancy: Professors Michael Anderson, Brian Wright, Solomon Hsiang, Meredith Fowlie, Maximilian Auffhammer, Jeremy Magruder, and Peter Berck. I would next like to thank those professors with whom I had the opportunity to work with as a research assistant in graduate school, and from whom I learned so much: Professors Peter Berck, Edward Miguel, and Meredith Fowlie. Next, I would like to thank those professors and instructors with whom I worked as a teaching assistant or co-instructor: Martha Olney, Brian Wright, Larry Karp, Meredith Fowlie, Joseph Shapiro, and Hal Gordon. And of course, thank you to everyone who attended my talks, provided comments and feedback on my work, or assisted me as a co-GSI. I also acknowledge the generous support from the Sacheti Family Fellowship.

I would also like to acknowledge my ARE family and friends in Berkeley. First, Alana, Carmen, and Diana, whose assistance and support throughout my time in graduate school was tremendously valuable. Second, my entire first-year cohort: Léo, Hannah, Shelley, Luna, Wei, Robert, Jay, James, Molly, Carly, and Katie; and to other close friends in California: Karl, Salil, Devin, and Helen. From travels to Tanzania or Tahiti, trips to Tahoe where it snowed too much to ski, and hikes in the Bay and Yosemite, the support of my friends was invaluable. Finally, a thank you to Jerry, my undergraduate alumni mentor who continued to provide both mentorship and friendship even during my time in graduate school.

Finally, I would like to acknowledge my family and pets, who gave me someone to call, and vacations and holidays to look forward to. Even my family's cat who tried to keep me from writing my dissertation by sitting in front of my computer at every opportunity deserves a thank you.

Chapter 1

Shadow Prices: Measuring the Cost of Shadows from New Construction in New York City

1.1 Introduction

The reduction in access to sunlight due to shadows created by new construction is a classic externality problem. As cities grow taller and denser, more people will be shadowed from the sun. The first citywide zoning ordinance in the US, New York's 1916 Zoning Resolution, was a direct response to the rise of skyscrapers which were perceived as preventing sunlight and air from reaching the street. The zoning ordinance required that buildings be further set back from the street as they rose past certain heights. This law contributed directly to the development of Art Deco-style skyscrapers in NY, with some of New York's most iconic buildings incorporating these setbacks into their design (Bliss, 2016).

In an interview on buildingtheskyline.org, economist Edward Glaeser of Harvard University said "I think that the starting point [in the debate on taxing or regulating the construction of tall luxury apartment buildings] has to be a serious quantification of the externalities [from height]. Almost none of the land use regulations that cities have ardently adopted for over a century have been justified by any serious quantitative work. The existence of even the smallest externality, such as a shadow, has been seen as justification for massive interventions in the housing market" (Barr, 2018). In 2017, the City of Berkeley, California was sued over its denying a permit to a contractor that planned to demolish a home and replace it with three smaller units on a single lot. A focal point that led to the permit's denial was a shadow that would be cast on the neighboring lot's vegetable garden. Sophie Hahn, at the time a member of Berkeley's Zoning Adjustments Board, said with respect to the project, "When you completely shadow all of the open space, you really impact the

ability for anybody to possibly grow food in this community" (Dougherty, 2017).

There are a few channels through which access to sunlight may impact housing prices. Fist, sunlight is an amenity which may directly affect people's utility of living in a certain location, for example, by directly impacting people's health. An et al. (2016) find that exposure to sunlight was a dominant predictor of anxiety, depressed mood, job satisfaction, and organizational commitment. Lambert et al. (2002) find that the duration of bright sunlight relates directly to serotonin production in the brain. Beaulieu and Dumont (2007) find that sunlight influences the secretion of melatonin which regulates the sleep-wake cycle and can be used to treat some mood and sleep disorders. Holick (2004) finds that exposure to sunlight can help with vitamin D sufficiency.

Second, sunlight may affect energy use. For example, there may be a positive benefit of summer shadows and a negative effect of winter shadows due to decreased cooling costs in summer and increased heating costs in winter. Donovan and Butry (2009) find in Sacramento that shade trees located to the south or west of a house reduced summertime energy use by 5.2%, a potential positive benefit of shadows.

Some papers that have attempted to answer the question of what is the cost of the externality of shadows. Fleming et al. (2018) estimated the value of sunlight for housing sales in Wellington, New Zealand using a hedonic regression and found that an extra hour of average daily sunlight is associated with a 2.6% increase in house sale price. Li (2022), however, finds that while new highrises decrease nearby rents, there does not appear to be any affect of new building height on nearby rental prices, a finding which the author suggests makes it implausible that shadows can explain the negative effect of new highrises .

There is now a desire to use new sources of "big" data at fine geographic scales to answer questions in urban economics (Glaeser et al., 2018). Methods to estimate specific directionality of shadows from buildings have recently been developed using highly detailed 3D models of New York City's buildings. Miranda et al. (2019) use the properties of sun movement and ray tracing to develop a shadow accrual map for all of New York City.

In this paper, I estimate the external cost of shadow from the construction of new buildings in Manhattan between 2005 and 2014. I first computationally generate data on shadows cast by the stock of buildings in each year. I next measure the amount of additional shadow received by shaded buildings from each shadow-casting building. Then, making use of year-by-year changes in shadow on residential properties from the construction of new high-rises, I measure the effect of additional average shadow on housing prices.

My empirical strategy differs from that of previous papers estimating the impact of shadows on housing prices in that this paper's methodology makes use of quasi-experimental changes in shadows. I regress changes in housing prices on changes in average daily shadow, control-

ling for a quadratic function of distance to the newly constructed building, which is designed to control for most non-shadow externalities of the new building, including views. In order to control for neighborhood-level changes in housing prices over time, I use spatial-time fixed effects, which control for average housing price changes within the neighborhood of a given apartment unit over time. This method exploits variation in shadow position, in that some units nearby to new construction may receive more shadow than others at a similar radius from the new building based on their direction from the new building, resulting from relative solar position. While this does not completely rule out the effect of views, as long as there is not a consistent direction towards which views are favored that aligns with solar position, views can be ruled out as the primary driver of my results.

In the results section, I demonstrate that a 10 percentage point increase in average daily shadow for a unit is associated with an approximately 3.78% decrease in unit price. I further test whether this percentage varies based on existing shadow. I find that the more existing shadow there already is, the larger will be the percentage decrease in price resulting from a 10 percentage point increase in average daily shadow, although this effect is not significant across specifications. A subsection considers robustness checks for the potential impact of views on my results.

Finally, I describe a potential policy that can result from this research. In cities with existing height restrictions, my proposal is to give developers the ability to build past the height restriction, in exchange for paying a tax on the shadow externality that they create. This tax is proposed specifically based on the proportion of the building surpassing existing height restrictions, as the tax does not simultaneously subsidize positive externalities (e.g. due to job or amenity effects). I work through the details of such a policy, and present a case study of what the tax would look like for a current building under development, and how we could determine the optimal height of the building under the tax which maximizes profits for the developer.

The paper proceeds as follows. In Section 2, I describe my sources of data and explain how I generate the shadow data. In Section 3, I describe my empirical strategy. In Section 4, I describe the results of my estimation, including a subsection for robustness to views. In Section 5, I describe a policy that could be implemented based on these results, including a subsection with a real-world application. Section 6 concludes.

1.2 Data

My paper relies on three primary sources of data: transactions data, which provide the transaction prices of residential units in NYC; buildings data, which provide the shape data for each building in NYC and the year of its construction; and shadow data, which I generate myself for each observation in the transactions data in Manhattan between 2004 and 2014

based on the shadows generated by buildings in the buildings data. (Note that 2004 is one year prior to the first year that I measure changes in shadow from the construction of new buildings, as baseline shadows need to exist prior to a new building's construction in order to estimate a change in shadow for a given residential unit.) I additionally use solar position data from the Pysolar Python package to assist in generating the shadow data.

Transactions data uses the City of New York's Annualized Sales Update data.[1] These data list all property sales in New York City since 2003. These data include variables for building class (e.g. walkup apartment), address, apartment number (if applicable), number of residential units, number of commercial units, square feet of land, square feet of property, year built, sale price, and sale date, among others. I use the Google Maps API to geocode street addresses into latitude and longitude (lat-lon) coordinates, which allows for merging with other sources of data. I use these transactions to estimate the prices of transactions of NYC housing units, restricting the observations to transactions of individual residential units.

Buildings data uses the City of New York's Building Footprints data.[2] These data contain the shapes of each NYC building's footprint in latitude and longitude coordinates, year of construction of the building, and height of the building. I use these data both to determine the shape of the building that houses each residential unit, and also the shape of all other buildings in Manhattan for determining the height and shape of the shadows that they cast.

I generate shadow data computationally for each residential transaction in the transactions data occurring in Manhattan between 2004 and 2014, based on the shapes of their building in the building footprints data. I begin by matching the location of each residential address to a building from my buildings data. Using the lat-lon coordinates for a given residential building using Google Maps for geocoding addresses, I see which building footprint in the buildings data contains each residential unit's geocoded location. This gives me the shape and height of the building containing each residential unit.

Once I determine each building that hosts a residential unit in my data, I estimate the shadow received by that building in a given year. As shadows vary both across times within a day, and across days within a year, to limit my computation time, I calculate shadows on three specific dates at 5 minute intervals: on December 21 (winter solstice) between 9am and 2:55pm EST (UTC–05:00), on June 20 (summer solstice) between 7:30am and 4:25pm EST, and September 22 (fall equinox) between 6am and 5:55pm EST, for each year between 2004 and 2014. (Note the use of EST for all dates, which centers the measurement for each date on approximately solar noon.) For each of these dates, I use solar position based on the 2016 solar position regardless of year of observation, both for simplicity and consistency in comparing times of day with equal solar positions across years. The date of the spring

[1] https://www1.nyc.gov/site/finance/taxes/property-annualized-sales-update.page

[2] https://data.cityofnewyork.us/Housing-Development/Building-Footprints/nqwf-w8eh

equinox (March 20) is approximately equal in solar position to the fall equinox, so applying the fall equinox shadow data to the spring equinox gives me four dates worth of shadow data, one representing each season.

Next, for each date-time beginning in 2004 and ending in 2014, I calculate the shadow of each building in my buildings data with a height greater than or equal to 50 feet tall, (I implement the 50-foot restriction to limit computation time), built in that year or earlier. For each building b, I first determine the length of the shadow cast by building b at a given point in time τ using the angle of the sun above the horizon θ_τ, which I get from the Pysolar Python package,[3] and the height of building b, height_b. I calculate the length of building b's shadow as

$$\text{len}_{b\tau} = \text{height}_b / \tan(\theta_\tau).$$

I next determine the direction of shadow at a given time of day τ using the azimuth of the sun (the position of the sun in clockwise degrees from North), ϕ_τ, which I also get from the Pysolar Python package. The shadow's direction can be determined in terms of coordinates on a unit circle, $(\sin(\phi_\tau - 180), \cos(\phi_\tau - 180))$, with the positive x-axis representing east, and the positive y-axis representing north. These coordinates are multiplied by $\text{len}_{b\tau}$ to determine the relative length of the shadow in feet in the horizontal (east-west) and vertical (north-south) direction from the building.

Next, I construct the shadow polygon for each building b at time τ. For any two side-by-side coordinates in the building footprint, I consider those two coordinates to form a wall. I then project those two coordinates by the length of the shadow and in the direction of the shadow to generate two new coordinates, which are the shadow of the wall's upper corners. That is, I add $\text{len}_{b\tau} \times (\sin(\phi_\tau - 180), \cos(\phi_\tau - 180))$ to both coordinates that form the wall. Drawing a parallelogram around these four coordinates gives me the shadow of that wall. I generate a shadow for each wall of a building, and finally take the union of these shadows, which approximates the entire shadow of the building.

Figure 1.1 shows an example of the shadow projection from the New York Federal Reserve (NY Fed) building at 7am, 9:30am, 12pm, 2:30pm, and 5pm EST during the Summer (June 20), Fall (September 22), and Winter (December 21), respectively (removing 7am and 5pm for winter and fall). The base polygon of the NY Fed is in red, while the shadow polygon is in blue. All times are in Eastern Standard Time (UTC-05:00) such that 12pm is the hour closest to solar noon on all three dates. (Note that New York uses Eastern Daylight Time (UTC-04:00) on June 20 and September 22. All times in this paper follow EST (UTC-05:00).) The choice of a single timezone EST for all three dates by ignoring the timezone shift for

[3] https://pysolar.org/

Figure 1.1: Shadow Projection From the New York Federal Reserve

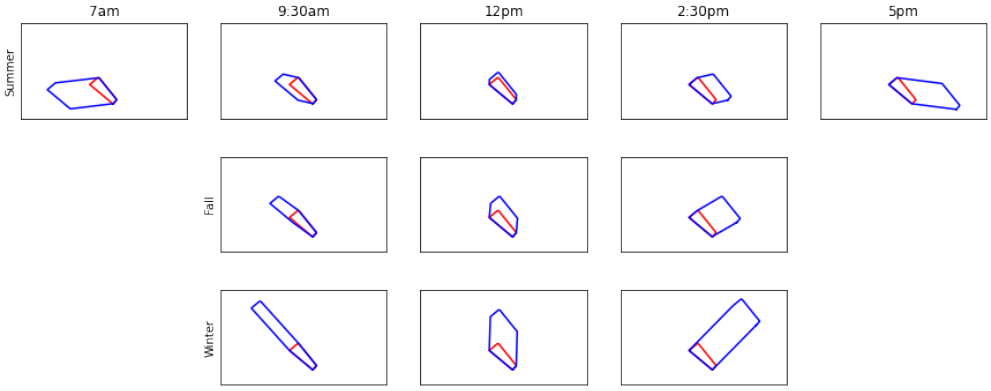

The red polygon towards the center of each plot represents the New York Federal Reserve's (NY Fed) building footprint. The surrounding blue polygon represents the NY Fed's shadow polygon. The first row labeled Summer shows the shadow on June 20. The second row labeled Fall shows the shadow on September 22. The third row labeled Winter shows the shadow on December 21. Columns represent shadows at different times of day in UTC-05:00. Shadows at 7am and 5pm are not included during Fall and Winter.

daylight savings is done for simplicity. Choice of timezone does not affect the analysis, as the measure of shadow considers the fraction of total daylight hours shadowed and is not based on specific times of day.

For each residential building r hosting a residential unit i in my transactions data, I generate a variable determining the amount of shadow on r at time τ, $\text{shadow}_{r\tau}$. $\text{shadow}_{r\tau}$ will range from 0 to 1 based on the height of the building covered by shadow, with 0 meaning that building r receives no shadow at time τ, 1 meaning that building r receives a shadow up to the height of its roof at time τ, and a number in between meaning that some fraction $\text{shadow}_{r\tau}$ of building r's height is covered in shadow.

To calculate $\text{shadow}_{r\tau}$, I first check if building b's shadow falls on building r at time τ. I simplify this step by checking only if the centroid of address r's building footprint falls within building b's shadow polygon at time τ. If it does, I then calculate the fraction of building r's height that is shadowed by comparing the height of the shadow at the coordinate of building r's centroid to the height of building r. To calculate the height of building b's shadow at the centroid of building r at time τ, $\text{height}_{br\tau}$, I next calculate the distance between the

two buildings, dist_{br}, using the Euclidean distance between the centroids of the respective buildings' footprint polygons. I then calculate the height of the shadow at the centroid of building r using the formula

$$\text{height}_{br\tau} = \frac{\text{len}_{b\tau} - \text{dist}_{br}}{\text{len}_{b\tau}} \times \text{height}_{b}.$$

To determine the amount of shadow on the residential address r, $\text{shadow}_{br\tau}$, I then calculate the height of the shadow as a fraction of building r's height, height_{r}, not to surpass 1.

$$\text{shadow}_{br\tau} = \min\left\{1, \frac{\text{height}_{br\tau}}{\text{height}_{r}}\right\}.$$

Finally, I calculate the total shadow on building r at time τ, $\text{shadow}_{r\tau}$, as a fraction of building r's height to be the maximum fraction of shadow $\text{shadow}_{br\tau}$ cast by all buildings b on building r at time τ.

$$\text{shadow}_{r\tau} = \max_{b} \text{shadow}_{br\tau}.$$

Figure 1.2 shows the red building (red polygon) casting a shadow (blue polygon) in the direction of the green building (green polygon). At 10:30am, the shadow partially covers the green building. However, for simplicity, as the centroid of the green building falls just outside of the shadow, I do not consider the green building as shadowed by the red building at 10:30am (note that other buildings may still shadow the green building). At 11am, the centroid of the green building is clearly within the shadow of the blue building, so I consider the green building to be shadowed by the red building. At 11:30am, the centroid of the green building again falls just outside of the shadow, so I do not consider the green building as shadowed by the red building. While at 10:30am and 11:30am, it appears that the green building could be considered about half shadowed by the red building, this should balance out, as if the green building's centroid were to fall just slightly inside of the blue shadow polygon, which would also lead the green building to appear half shadowed, I would calculate a full shadow on the green building. As I measure shadows at 5-minute intervals, the green building is considered shadowed by the red building for all times starting from 10:35am and ending at 11:20am.

Figure 1.3 shows how the vertical shadow of building b (red) on building r (green) is measured at time τ. The shadow is shown in blue, where θ_τ is the angle of the sun above the horizon,

Figure 1.2: Throwing Shade: Shadowing Another Building

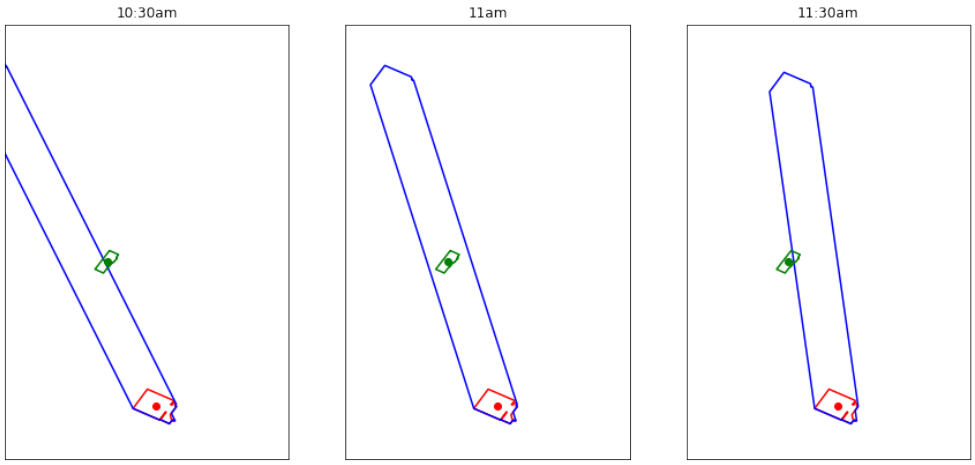

The red polygon towards the bottom of each plot represents World Trade Center One's (WTC) building footprint at 10:30am, 11am, and 11:30am (UTC-5:00) on December 21. The surrounding blue polygon represents the WTC's shadow polygon. The green building towards the center represents the footprint of a building shadowed by the WTC. Centroids are plotted at the center of each building's footprint polygon. The vertical shadow cast by the World Trade Center is measured at 11am when green building's centroid is within the WTC's shadow, but not at 10:30am or 11:30am when the green building's centroid is outside of the WTC's shadow.

height_r is the height of building r, and height_b is the height of building b. $\text{height}_{br\tau}$ is the vertical height of building b's shadow on building r at time τ

To convert these data into annual shadow data for each residential address, I calculate the average shadow received by building r across the 5 minute intervals for each of the three dates, then calculate the average shadow across the three dates for each year, double weighting the fall equinox, September 22, to account for the Spring equinox, March 19. I use this average as an approximation of the average annual shadow across a given year. That is, for residential building r in year t, I calculate

Figure 1.3: Vertical Shadow Measurement

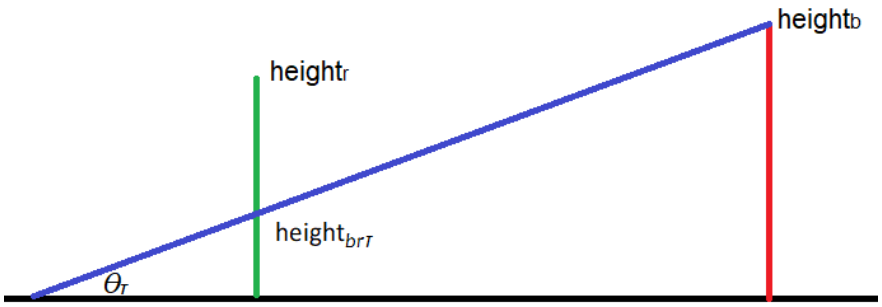

The vertical red line on the right represents the height of a building b that we are measuring the shadow of, with height height_b. The vertical green line towards the left represents a building r that is being shadowed by the red building, with height height_r. The diagonal blue line represents the maximum height of the shadow cast by building b as it falls towards the ground, where θ_τ is the angle of the sun above the horizon at time τ. $\text{height}_{br\tau}$ is the measured height of the shadow cast by building b on building r at time τ.

$$\text{shadow}_{rt} = \frac{1}{4}\sum_{S}\left[\frac{1}{N_S}\sum_{\tau \in t_S}\text{shadow}_{r\tau}\right]$$

where S is the season, representing Winter, Summer, Fall, and Spring, N_S is the number of 5 minute intervals in the respective date for season S, and t_S represents the set of 5 minute intervals in the respective date for season S in year t.

1.3 Empirical Strategy

My empirical approach compares changes in housing prices for residential units to changes in average daily shadow received by those units. I calculate the average daily shadow for unit i in year t

$$S_{it} = \text{shadow}_{r_i t},$$

where r_i is the building r containing unit i. (Calculation of shadow$_{rt}$ is detailed in the previous section.) S_{it} ranges from 0 to 1, with 0 meaning that building r_i receives no shadow at all times that I measure in year t, 1 meaning that a shadow up to the roof of building r_i is being cast on building r_i at all times that I measure in year t, and a number in between 0 and 1 meaning that some fraction S_{it} of building r_i's height is covered in shadow on average in year t. Ideally, I would know what floor or side of building r_i unit i is located on, which could allow me to calculate a more precise measure of shadow received by unit i. However, this is not allowed by the data, which does not include specific data on where in a given building an apartment or condo unit is located. As such, I approximate the shadow received by the unit using the average shadow received by its building.

I use a differenced regression model with spatial-time fixed effects, based on the spatial fixed effects model from Anderson (2020). For a given variable u_{it}, I define $\widetilde{u}_{it_j} = u_{it_j} - \bar{u}_{it_j}$, where $\bar{u}_{it_j}$ is the mean of variable u within radius k of unit i in year t_j. This regression is of the form

$$\Delta\widetilde{y}_{it_j} = \beta\Delta\widetilde{x}_{it_j} + \Delta\widetilde{\varepsilon}_{it_j},$$

with $j \in \{2, 3, \ldots, n_i\}$, where n_i is the number of observations for unit i in the sample, and Δ represents the change from t_{j-1} to t_j. That is,

$$\Delta\widetilde{x}_{it_j} = \widetilde{x}_{it_j} - \widetilde{x}_{it_{j-1}}.$$

Because transactions for each unit do not necessarily take place every year, t_{j-1} and t_j may be one or more years apart. The inclusion of spatial-time fixed effects control for average changes in housing values between periods t_{j-1} and t_j that are within radius k of unit i. This allows the regression to account for the panel imbalance by netting out time-related housing price changes that affect all units within radius k, allowing me to focus on unit-level housing price variation over time rather than neighborhood-level variation (e.g. removing the effect general housing price changes in a given neighborhood). Due to panel imbalance, the spatial-time fixed effects within the same radius may be averaged over a different set of units for any two years t_{j-1} and t_j, but for each period, the spatial time fixed effect still represents the average housing sale price within a given neighborhood in year t.

Writing out the differences in longer form, the regression can be written as

$$\widetilde{y}_{it_j} - \widetilde{y}_{it_{j-1}} = \beta\left(\widetilde{x}_{it_j} - \widetilde{x}_{it_{j-1}}\right) + \left(\widetilde{\varepsilon}_{it_j} - \widetilde{\varepsilon}_{it_{j-1}}\right).$$

Writing out the spatial-time fixed effects as well, this becomes

$$\begin{aligned}\left(y_{it_j} - \bar{y}_{it_j}\right) - \left(y_{it_{j-1}} - \bar{y}_{it_{j-1}}\right) = \beta \left(\left(x_{it_j} - \bar{x}_{it_j}\right) - \left(x_{it_{j-1}} - \bar{x}_{it_{j-1}}\right)\right) \\ + \left(\left(\varepsilon_{it_j} - \bar{\varepsilon}_{it_j}\right) - \left(\varepsilon_{it_{j-1}} - \bar{\varepsilon}_{it_{j-1}}\right)\right).\end{aligned}$$

Note that the differences and spatial-time fixed effects can be swapped.

$$\begin{aligned}\left(y_{it_j} - y_{it_{j-1}}\right) - \left(\bar{y}_{it_j} - \bar{y}_{it_{j-1}}\right) = \beta \left(\left(x_{it_j} - x_{it_{j-1}}\right) - \left(\bar{x}_{it_j} - \bar{x}_{it_{j-1}}\right)\right) \\ + \left(\left(\varepsilon_{it_j} - \varepsilon_{it_{j-1}}\right) - \left(\bar{\varepsilon}_{it_j} - \bar{\varepsilon}_{it_{j-1}}\right)\right).\end{aligned}$$

The spatial-time fixed effects can thus be thought of similarly to a standard spatial fixed effects, but in the differenced model the spatial-time fixed effects are also differenced over time.

To estimate the effect of shadows on housing prices, I regress changes in log housing prices on changes in average daily sunlight, S, using this differenced model with spatial-time fixed effects. This regression follows the equation

$$\widetilde{\Delta \log(p_{it_j})} = \beta \widetilde{\Delta S_{it_j}} + \Delta \widetilde{\varepsilon}_{it_j}. \tag{1.1}$$

In Equation (1.1), $\log(p_{it_j})$ is the log sale price of housing unit i in year t_j. S_{it_j} is the average daily shadow received by unit i, averaged over year t_j, as previously defined. ε_{it_j} is an error term.

For a given unit i in my data with n_i transactions, the number of differenced observations is $n_i - 1$ for $n_i > 1$. A unit with only one transaction does not appear as an observation for estimating the regression, but is included for the calculation of spatial-time fixed effects, as the averages are taken before differencing.

Note that due to differencing, $\widetilde{\Delta S_{it_j}}$ is only non-0 when there is a change in shadow on unit i between periods t_{j-1} and t_j, or a change in the spatial-time fixed effect of unit i resulting from a change in shadow to one or more nearby units between periods t_{j-1} and t_j. As a result of differencing, a richer set of unit-specific controls such as square footage, number of bathrooms, etc., is not needed, as differencing removes variation in pricing resulting from such fixed features of units.

Figure 1.4: Locations of new buildings over 50 feet tall constructed between 2005 and 2014

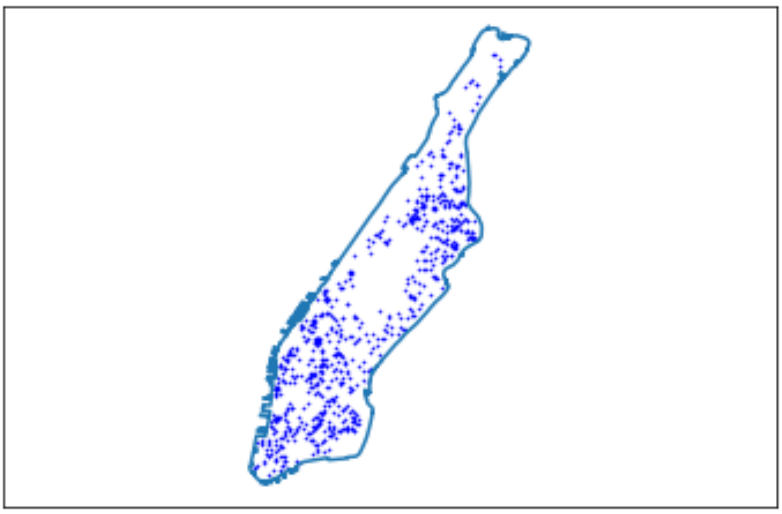

Each blue dot represents a building in Manhattan over (or equal to) 50 feet tall for which construction was completed between 2005 and 2014.

The identification of a causal effect in this model is derived from the fact that shadows generated by a new building are not spread uniformly to existing units. The additional shadow received by two units located at the same distance from a new building depends on both the direction from the new building that the existing unit is located (e.g. units directly to the south of a new building will not receive any additional shadow) and the pre-existing shadow on the new building (e.g. if the building is already fully shadowed, there will not be additional shadow generated). This allows me to identify the changes in unit prices due to shadows separately from the changes in unit prices due to other neighborhood-level externalities from the the construction of a new building.

In order to control for potential neighborhood effects resulting from the construction of a new building, I additionally run a regression which includes as controls a quadratic function of the distance to a new building over 50 feet tall within 500 meters of unit i (see Figure 1.4). This accounts for the effect of the shadow-generating building's neighborhood price effect on all housing units within 500 meters of the new building, and assumes this effect changes with distance from the new building. For the intercept term of this quadratic function, I first define a constant term $D500_{it_j}$, which is equal to the number of new buildings over 50

feet tall within 500 meters of unit i built between the start of my sample and period t_j. For example, if three new buildings over 50 feet tall within 500 meters of unit i are built between t_{j-1} and t_j, then $\Delta\text{D500}_{it_j} = 3$. The linear term in the quadratic function, $dist500_{it_j}$, is the sum of the distances to unit i of all buildings over 50 feet tall within 500 meters of unit i built between the start of my sample and period t_j. The quadratic term in the quadratic function, $dist500sq_{it_j}$, is the sum of the squares of the distances to unit i of all buildings over 50 feet tall within 500 meters of unit i built between the start of my sample and period t_j. Summing the values of the quadratic function of distance is equivalent to having a separate quadratic function of distance for each building within 500 meters of unit i with the same coefficients. This regression follows the equation

$$\Delta\widetilde{\log(p_{it_j})} = \beta\Delta\widetilde{S_{it_j}} + \gamma_1\Delta\text{D}\widetilde{500}_{it_j} + \gamma_2\Delta\text{dist}\widetilde{500}_{it_j} + \gamma_3\Delta\text{dist}\widetilde{500\text{sq}}_{it_j} + \Delta\widetilde{\varepsilon}_{it_j}. \tag{1.2}$$

Due to differencing, all terms of the quadratic function will be 0 if there is no new building built within 500 meters of unit i between t_{j-1} and t_j, and no change in the spatial-time fixed effect resulting from a new building built within 500 meters of a nearby unit.

In a third regression specification, I interact the quadratic function of distance to new buildings with a linear function of the height of the new building. There are three new controls used in this regression. height500_{it_j} is the sum of the height of all buildings over 50 feet tall within 500 meters of unit i built between the start of my sample and period t_j. $\text{dist500\_height500}_{it_j}$ is the sum of the distances to unit i times the height of that building of all buildings over 50 feet tall within 500 meters of unit i built between the start of my sample and period t_j. $\text{dist500sq\_height500}_{it_j}$ is the sum of the squares of the distances to unit i times the height of that building of all buildings over 50 feet tall within 500 meters of unit i built between the start of my sample and period t_j. These additional terms account for the possibility that the shadow-generating building's neighborhood price effect on housing units within a 500 meter radius varies with the height of the new building, independent of the additional shadow cast by that building. This regression, including the interaction with new building height terms, follows the equation

$$\begin{aligned}\Delta\widetilde{\log(p_{it_j})} = {} & \beta\Delta\widetilde{S_{it_j}} + \gamma_1\Delta\text{D}\widetilde{500}_{it_j} + \gamma_2\Delta\text{dist}\widetilde{500}_{it_j} + \gamma_3\Delta\text{dist}\widetilde{500\text{sq}}_{it_j} \\ & + \gamma_4\Delta\text{height}\widetilde{500}_{it_j} + \gamma_5\Delta\text{dist500\_}\widetilde{\text{height}}500_{it_j} + \gamma_6\Delta\text{dist500sq\_}\widetilde{\text{height}}500_{it_j} \\ & + \Delta\widetilde{\varepsilon}_{it_j}.\end{aligned} \tag{1.3}$$

Views are not directly controlled for in Equations (1.1) – (1.3). However, as long as views are not consistently more valuable in a direction that covaries with relative solar position, the

third specification should best control for the effect of views, as the height of new buildings which may block views and the distance to the new building are both controlled for separately from variation in shadow. As such, this is my preferred specification.

I additionally test whether the amount of baseline shadow in period t_{j-1}, $S_{it_{j-1}}$, modifies the effect of a given change in shadow. That is, for a given change in shadow, $S_{it_j} - S_{it_{j-1}}$, does a larger baseline shadow $S_{it_{j-1}}$ result in a larger price change? I estimate this by adding an interaction term of the change in shadow with the baseline shadow in period t_{j-1}, $S_{it_{j-1}}$, such that the interaction term is $\Delta S_{it_j} \times S_{it_{j-1}}$. Without spatial-time fixed effects, this would be a regression of the form

$$\Delta \log(p_{it_j}) = \beta \Delta S_{it_j} + \delta \left(\Delta S_{it_j} \times S_{it_{j-1}} \right) + \Delta \varepsilon_{it_j}. \tag{1.4}$$

In order to calculate the spatial-time fixed effects for $\Delta S_{it_j} \times S_{it_{j-1}}$, I first distribute the interacted baseline shadow $S_{it_{j-1}}$ across $\Delta S_{it_j} = S_{it_j} - S_{it_{j-1}}$. That is, $\Delta S_{it_j} \times S_{it_{j-1}} = S_{it_j} S_{it_{j-1}} - (S_{it_{j-1}})^2$. I then separately calculate spatial-time fixed effects for $S_{it_j} S_{it_{j-1}}$ only using units with a price observation in period t_j, and for $(S_{it_{j-1}})^2$ only using units with a price observation in period t_{j-1}. This ensures that I am taking the spatial-time fixed effects over the same set of units for $S_{it_j} S_{it_{j-1}}$ and $(S_{it_{j-1}})^2$ as for S_{it_j} and $S_{it_{j-1}}$, respectively. The regression that I thus estimate to test for the impact of baseline shadow on the effect of a change in shadow on housing price follows the equation

$$\Delta \widetilde{\log(p_{it_j})} = \beta \Delta \widetilde{S_{it_j}} + \delta \left(\widetilde{S_{it_j} S_{it_{j-1}}} - \widetilde{(S_{it_{j-1}})^2} \right) + \Delta \widetilde{\varepsilon}_{it_j}. \tag{1.5}$$

1.4 Results

Table 1.1 displays the estimates of regressions following Equations (1.1), (1.2), and (1.3) in Columns (1), (2), and (3), respectively. Standard errors are clustered using a 250 meter[2] grid. The spatial-time fixed effects use a 250 meter radius. Columns (2) and (3) both include as controls a quadratic function of distances to newly constructed buildings. Column (3) additionally interacts this quadratic function with the height of the new building.

Column (3), my preferred specification, suggests that a 10 percentage point increase in average daily shadow for a unit is associated with an approximately 3.78% decrease in unit price. To put this into context, assume a one million dollar unit receives on average 6 hours out of a 12 hour day in shadow. Suppose a new building increases the shadow on this unit

[4]Appendix Table A.1 includes estimates similar to those in Table 1.1 but using a 500 meter radius for spatial-time fixed effects and a 500 square meter spatial grid for clustering of standard errors.

by one hour to 7 out of 12 hours, equal to an 8.3 percentage point increase in average daily shadow. Assume a similar 8.3 percentage point increase in average daily shadow across all days of the year. Then the value of the housing unit would be predicted to fall by $.378 \times .083 = 3.14\%$, or approximately $31,400.

Table 1.2 displays the results from the regression given by Equation (1.5) in Column (1), with the addition of distance and distance×height controls in Columns (2) and (3), similar to those in Equations (1.2) and (1.3), respectively. Standard errors are clustered using a 250 meter^2 grid. The spatial-time fixed effects also use a 250 meter radius.

These regressions test if the amount of baseline shadow modifies the effect of a given change in shadow. That is, for a given change in shadow, does a larger baseline shadow result in a larger price change? While the interacted results are not consistently significant across specifications, they are consistently signed and of similar magnitude, suggesting there may be some (weak) effect of baseline shadow on the effect of the shadow externality. Column

Table 1.1: Effect of changes in shadow on changes in housing price

	$\Delta\widetilde{\log(p_{it_j})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it}}$	-0.451***	-0.432***	-0.378***
	(0.142)	(0.136)	(0.134)
Spatial-Time FE radius	250m	250m	250m
Cluster Grid	250m^2	250m^2	250m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	35,180	35,180	35,180
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. Column (1) estimates Equation (1.1). Column (2) estimates Equation (1.2). Column (3) estimates Equation (1.3). Spatial-time fixed effects are measured with a 250 meter radius prior to differencing. Standard errors are clustered using a 250 square meter spatial grid.[4]

[5]Appendix Table A.2 includes estimates similar to those in Table 1.2 but using a 500 meter radius for spatial-time fixed effects and a 500 square meter spatial grid for clustering of standard errors.

(3) suggests that, for a given change in shadow, a 10 percentage point higher average daily baseline shadow for a unit is associated with an approximately 1.16 percentage point larger decrease in unit price.

To put this into perspective, assume that there are two one-million dollar units, one which receives an average of 0 hours out of a 12 hour day in shadow, and the other which receives an average of 11 hours out of a 12 hour day in shadow. Suppose a new building increases the shadow on both units by one hour out of the 12 hour day. That is, the shadow on the first unit is now 1 out of 12 hours, and for the second unit, it is 12 out of 12 hours. For both units, this is equal to an 8.3 percentage point increase in daily shadow. The first unit, however, began with 0 percent of the day in shadow, and the second unit began with 91.7% of the day

Table 1.2: Effect of changes in shadow relative to baseline shadow on changes in housing price

	$\Delta\widetilde{\log(p_{it_j})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.355***	-0.343**	-0.295**
	(0.135)	(0.134)	(0.133)
$\widetilde{\left(\Delta S_{it_j}\right)} \times S_{it_{j-1}}$	-0.131*	-0.122	-0.116
	(0.079)	(0.076)	(0.079)
Spatial-Time FE radius	250m	250m	250m
Cluster Grid	250m^2	250m^2	250m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	35,180	35,180	35,180
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. $\widetilde{\left(\Delta S_{it_j}\right)} \times S_{it_{j-1}}$ represents the change in spatial-time demeaned shadow interacted with baseline shadow. Column (1) estimates Equation (1.5). Column (2) estimates Equation (1.5) with the inclusion of distance fixed effects as in Equation (1.2). Column (3) estimates Equation (1.5) with the inclusion of distance and height fixed effects as in Equation (1.3). Spatial-time fixed effects are measured with a 250 meter radius prior to differencing. Standard errors are clustered using a 250 square meter spatial grid.[5]

in shadow. Assume these percentages are similar across all days of the year. Then the value of the first unit would be predicted to fall by $.295 \times .083 = 2.44\%$, or about \$24,400, and the value of the second unit would be predicted to fall by $(.295 + .116 \times .917) \times .083 = 3.33\%$, or about \$33,300.

Robustness Check - Views

Perhaps the most important potential confounder with shadows is views. Views throughout a city, however, are difficult to quantify, as one would need to apply a numerical value to many different possible views. As described in Section 1.3, the distance and distance-times-height control variables are designed to capture the effect of views, under the assumption that there is not a consistent direction towards which views are favored that aligns with solar position. For example, this would assume that buildings to the south of a new building which do not receive as much new shadow have the same potential for blocked views of similar value when compared to buildings to the north, east, or west of a new building which would receive relatively more new shadows. This is conditional on an assumption that the quality of views in New York does not covary with relative solar position.

One way to test if the distance and distance-times-height control variables should work as intended is to see if the height of existing residential buildings co-varies with direction from new buildings, as taller residential buildings likely have better existing views. In order to perform this test, I first consider all existing residential buildings within a given radius of new buildings constructed between 2005 and 2014, where the residential building was built prior to the construction of the new building. I divide these residential buildings into four quadrants: South of the new building, North of the new building, East of the new building, and West of the new building. I perform this directional classification by considering where residential buildings within this given radius lie with respect to an X-shape centered on the new building's centroid. The X-shape is made up of two 45-degree lines running from north-west to south-east and from south-west to north-east. Residential buildings whose centroids fall to the south of the X are classified as South, to the north of the X as North, to the east of the X as East, and to the west of the X as West. For simplicity, I classify buildings whose centroids lie exactly along the lines forming the X as either South or North, as opposed to East or West. Figure 1.5 demonstrates this classification, using the New York Federal Reserve (NY Fed) as the central building. (Note that for the purpose of Figure 1.5, there is no restriction on construction year of residential buildings being prior to that of the NY Fed. The NY Fed would not be included in the sample of new buildings considered as it was not constructed between 2005 and 2014.

Table 1.3 presents summary statistics for the heights of residential buildings in the given quadrants of new buildings constructed between 2005 and 2014. This is repeated both for residential buildings within a radius of 500 meters of a new building, and within a radius equal to two times the new building's height.

Figure 1.5: 500m Directional Quadrants Centered on NY Federal Reserve

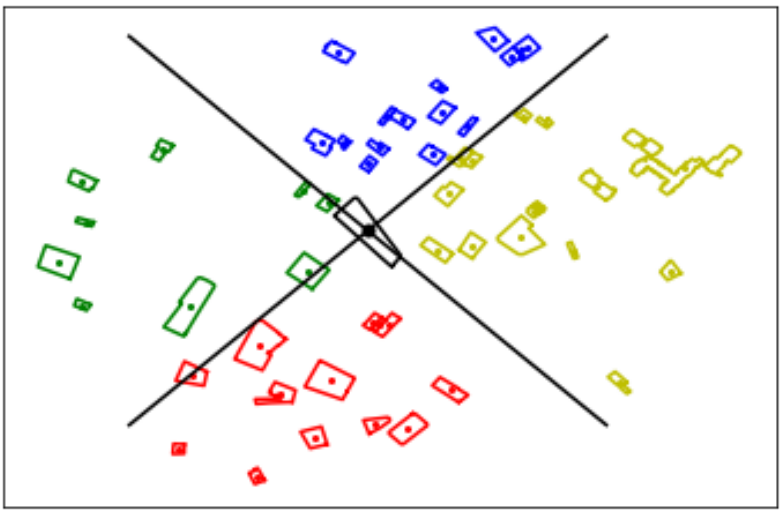

Each polygon represents a building, which contains a dot representing its centroid. The central polygon in black represents the NY Federal Reserve. The black X divides nearby residential buildings into quadrants. The red polygons represent buildings in the South quadrant. The blue polygons represent buildings in the North quadrant. The yellow polygons represent buildings in the East quadrant. The green polygons represent buildings in the West quadrant.

As I perform the directional classification of residential buildings for each new building, I allow for repeated observations of residential buildings in the sample of residential building heights: for example, if a residential building is in the South quadrant of one new building, and in the North quadrant of a different new building, I will consider its height as belonging to both the South quadrant of the first new building, and to the North quadrant of the second new building, allowing its height to be included in the calculation of mean height in both the South quadrant sample and the North quadrant sample. Similarly, if the same residential building is in both of the South quadrants of two different new buildings, the residential building's height will be repeated in the South quadrant sample.

The summary statistics in Table 1.3 show that the numbers of residential buildings within each quadrant are of a similar order of magnitude. The mean heights of the residential buildings are, at the maximum difference, 8.92 feet, when measured between the North and

Table 1.3: Height of residential buildings by direction from new buildings

	500m Radius			2×Height Radius		
	(1) Obs	(2) Mean	(3) SD	(4) Obs	(5) Mean	(6) SD
South	35616	75.69	56.81	1320	110.54	99.69
North	33026	77.17	59.46	1422	111.00	99.46
East	34645	74.80	57.88	1604	107.95	92.14
West	32841	77.50	56.96	1479	102.08	89.52

Columns (1) – (3) show summary statistics for residential buildings within 500 meters of a new building. Columns (4) – (6) show summary statistics for residential buildings within a radius of a new building equal to two times the new building's height. The columns marked Obs are the number of observations within each quadrant for the given radius. The columns marked Mean are the mean heights within each quadrant for the given radius. The columns marked SD are the standard deviation of heights within each quadrant for the given radius.

West quadrant given a radius of two times the new building's height. 8.92 feet is less than the height one might consider for a typical story (Hickey, 2014), suggesting that the direction of new shadows does not substantially co-vary with residential building height, and thus there is not a greater potential for blocked views in quadrants that are more likely to be shadowed. This allows us to more safely assume that the distance and distance-times-height control variables will control for the effect of blocked views, under the assumption that view quality is roughly equal across quadrants.

An additional test for robustness to the impact of views on the results is to determine if the results hold while removing units that are likely to have the best views in the city. While quantifying views is difficult, views of Central Park are considered to be among the most valuable views in Manhattan (Biggs, 2019). Therefore, I re-estimate the regressions in Table 1.1, but removing units falling within 500 meters of Central Park (see Figure 1.6).

Table 1.4 shows the estimates of regressions removing units within 500 meters of Central Park, which removes 5,923 observations from the sample. I keep the spatial-time fixed effects unchanged. Each column corresponds to its matching column in Table 1.1. While the coefficient estimate in each column falls somewhat in magnitude, the coefficients each remain within one standard deviation of the coefficients estimated in Table 1.1, and remain statistically significant at the 5% level.

Figure 1.6: Central Park with 500 meter buffer

The green rectangle at the center represents Central Park. The surrounding red area represents a buffer of 500 meters from Central Park.

While the building fixed effects and these robustness checks do not entirely rule out correlation between increased shadows and blocked views that may affect the results, the robustness checks provide evidence that shadows remain a significant determinant of housing prices in New York City independent of views. Note that for the purpose of the policy implications presented in the next section, that some of the estimated drop in housing prices may be the result of views rather than shadows does not affect the conclusion inasmuch as one might consider pricing blocked views as part of a price on negative externalities.

1.5 Policy Implications

In theory, such externalities as shadows from development could be handled by the market if the conditions of the Coase theorem were satisfied. Two assumptions of the Coase theorem

[6]Appendix Table A.3 includes estimates similar to those in Table 1.4 but using a 500 meter radius for spatial-time fixed effects and a 500 square meter spatial grid for clustering of standard errors. Appendix Tables A.4 and A.5 include estimates similar to those in Tables 1.2 and A.2, respectively, but removing observations within 500 meters of Central Park.

Table 1.4: Effect of changes in shadow on changes in housing price, excluding units within 500 meters of Central Park

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.386**	-0.375**	-0.330**
	(0.154)	(0.148)	(0.148)
Spatial-Time FE radius	250m	250m	250m
Cluster Grid	250m^2	250m^2	250m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	29,257	29,257	29,257
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. Column (1) estimates Equation (1.1). Column (2) estimates Equation (1.2). Column (3) estimates Equation (1.3). Observations do not include units within 500 meters of Central Park. Spatial-time fixed effects are measured with a 250 meter radius prior to differencing and prior to removing observations near Central Park. Standard errors are clustered using a 250 square meter spatial grid.[6]

that fail in this instance are the assumption of clearly defined property rights, and that there must be few affected parties. Regarding property rights: do incumbent buildings have the right to not be shadowed, or do new developers have the right to build and create shadows as they see fit? In the former case, the owners of a new development could pay incumbent nearby residential unit owners for the right to shadow their units. In the latter case, the owners of shadowed units could pay the owners of a new development to not throw shade in their direction. Regarding the number of affected parties: there are both many potential unit owners who could be shadowed by a new building, and many potential new buildings that could throw shade at a given residential unit.

Currently, NYC has a system of height restrictions based on a floor area ratio (FAR), commonly called "air rights." The FAR is the maximum ratio of floor area to lot area allowed by city zoning laws. For example, a building on a lot with a FAR of 10 could be up to 10 stories tall if taking up the whole lot, 20 stories tall if taking up half of the lot, 40 stories tall if taking up a quarter of the lot, etc. From the law's perspective, two adjoining lots can merge their air rights, allowing the owner of one lot to sell unused air rights to the neighboring lot.

For example, if two equally sized adjoining lots each have an allowable FAR of 10, but the building on one of the lots only has a FAR of 5, the owner of that lot could sell the unused FAR to the neighboring lot, allowing the neighboring lot to build up to a FAR of 15.

If the intent of height regulation is to maximize social welfare, a policy directly targeted at the externalities of height could improve welfare relative a policy that targets height itself. For example, if shadows were the only externality from height, an efficient policy would directly price this externality by taxing the creation of shadows.

An optimal pricing policy which maximizes welfare would set a tax on construction equal to the net externality of new development. Besides shadows, there may be other be other externalities of new construction. This might include positive externalities of development, such as providing additional jobs and amenities to residents (Couture, 2016; Li, 2022). I refer to negative externalities minus positive externalities as the net negative externality. At the optimal height, the marginal net negative externality of building taller would equal the marginal profit to developers of building taller.

In order for there to be an optimal height of new construction which maximizes social welfare, I first assume that the marginal net negative externality is increasing. This assumes that either the marginal positive externality of new construction is decreasing with height, or the marginal negative externality of of new construction is increasing with height, or both. The former can be rationalized in that many of the positive spillovers from density and new construction, such as new public infrastructure and amenities, can be assumed to be a concave function of the size of a new development. For example, a larger development may encourage more and better shops at the ground level or make it more likely that a new bus or metro line will come through the neighborhood, but only so many shops and public transit lines can be packed into a given area even as buildings grow ever taller. The latter case of increasing marginal negative externalities from height can be seen in this research: while an equally sized shadow created by the 10th floor or the 100th floor of a new building may have the same average effect, the taller a building is relative to those around it, the more additional shadow the new building is likely to create for each new story, as the building is more likely to rise above the height of existing shadows. Together, the concavity of positive externalities and convexity of negative externalities from height rationalizes an increasing marginal net negative externality from height. Figure 1.7 shows this marginal net negative externality.

I next assume that marginal profits to developers are decreasing in height, or downward-U shaped, such that marginal profits are eventually decreasing in height, even if initially increasing. To justify this assumption, I first assume that marginal construction costs are increasing with height (Glaeser, Gyourko, and Saks, 2005; Chau et al., 2007; Ahlfeldt and McMillen, 2018). With respect to marginal revenue, research has shown increasing marginal revenues with height (Wong, 2004; Chau et al., 2007). However, it must be that

Figure 1.7: Graph of marginal net negative externality

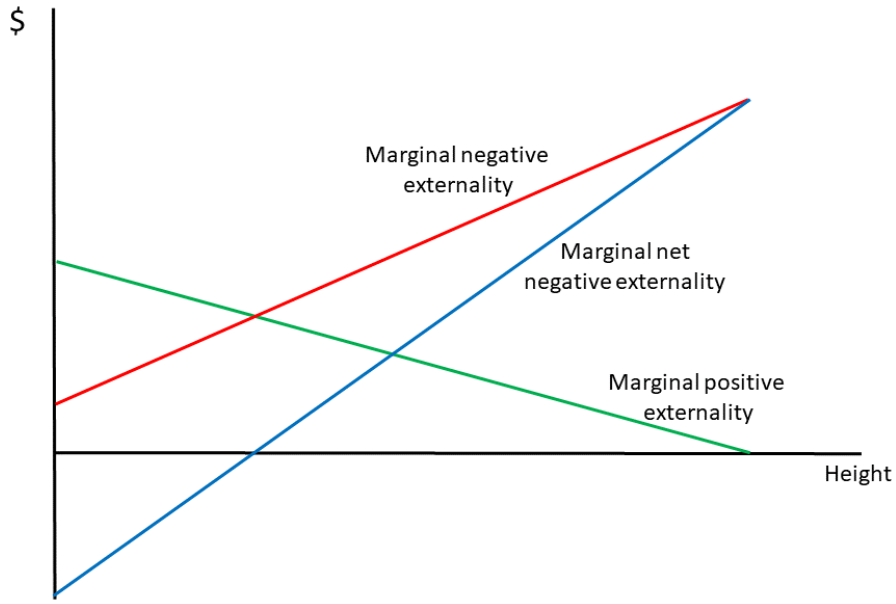

The red line represents the marginal negative externality. The green line represents the marginal positive externality. The blue line represents the marginal net negative externality. Marginal net negative externality is the difference between marginal negative externality and marginal positive externality.

marginal construction costs begin below marginal revenues from height, but eventually surpass marginal revenues; otherwise, an unregulated city would see either zero construction, or infinite building heights. Therefore, it is necessary that marginal construction costs increase more quickly than construction costs (implying downward-sloping marginal profits), or that either marginal construction costs are convex or that marginal revenues are concave in height (or both) - implying downward-U shaped marginal profits. Therefore, marginal profits of building taller must either be decreasing, or downward-U shaped. Figure 1.8 demonstrates downward sloping marginal profits with linear marginal construction cost and marginal revenue functions.[7]

[7]Appendix Table A.1 demonstrates a downward-U shaped marginal profit curve with both a convex marginal construction cost function and a concave marginal revenue function.

Figure 1.8: Graph of marginal profits

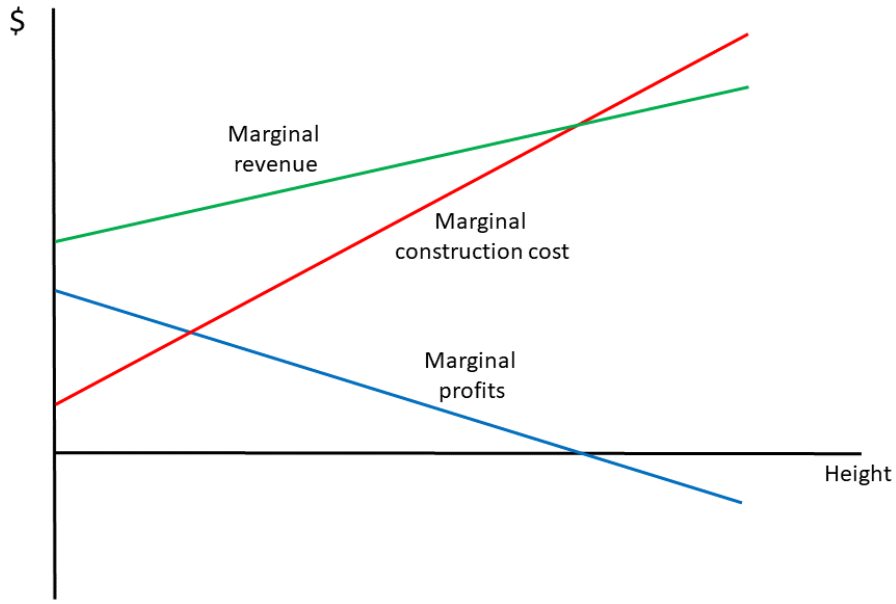

The red line represents the linear marginal construction cost. The green line represents the marginal revenue. The blue line represents the marginal profit to the developer. Marginal profits is the difference between marginal revenue and marginal construction cost.

Combining the assumptions of an increasing marginal net negative externality and decreasing or downward-U shaped marginal profits, an optimal height exists where the increasing marginal net negative externality from height crosses the eventually-decreasing marginal profits of developers. Figure 1.9 shows the optimal height at the intersection of the marginal net negative externality and marginal profits.[8]

In a city in which incumbent residents have all of the power in deciding construction policy, and in which policies are designed efficiently by incumbents to maximize their own welfare, the FAR would allow height up to the point at which the marginal net negative externality of height is zero. Assuming further that incumbents live in a city in order to benefit from some of the positive externalities associated with density, I assume that the marginal net

[8] Appendix Table A.2 shows the optimal height at the intersection of the marginal net negative externality and marginal profits with a downward-U shaped marginal profits curve.

Figure 1.9: Finding optimal height

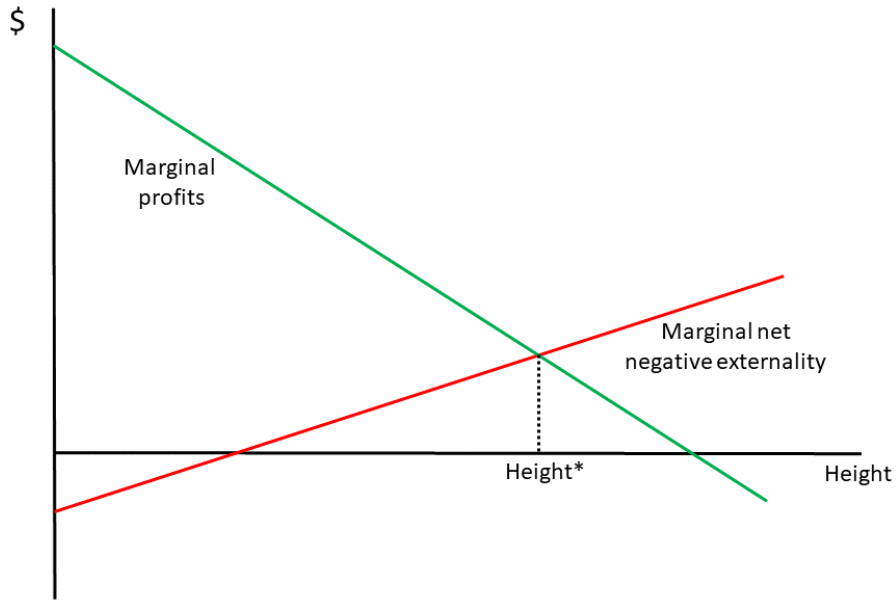

The red line represents the marginal net negative externality. The green line represents the marginal profit to the developer. The optimal height, labelled Height*, can be found where the marginal net negative externality equals the marginal profits.

negative externality is initially negative (that is, for lower stories, the marginal positive spillovers from development, such as the amenity effect, outweigh negative spillovers), such that incumbents would want some positive level of nearby development. In an unregulated city in which developers have all of the power, developers would build beyond the socially optimal height to maximize their own welfare, reaching the point where marginal profits from development are zero. Assuming current height regulations are designed by policymakers considering both the positions of rational incumbents and developers, and assuming that developers would build beyond current height regulations if they could, the FAR policy would be set such that building heights are somewhere in between the two extremes of the height preferred by incumbents where marginal net negative externalities are zero, and the optimal height preferred by developers where marginal profits are zero. I therefore assume that under current policy, both the marginal net negative externality and the marginal profits to developers are greater than or equal to zero.

Somewhere in between the two extremes of the best policy for incumbents and the best policy for developers lies the optimal regulation. At this optimal regulation, marginal profits to developers equal the marginal net negative externality, maximizing total surplus. However, without a full accounting of all positive and negative externalities of new construction, it would be practically impossible for a policymaker to pinpoint the specific height regulation which would achieve this optimum. However, it is possible for policymakers to use incomplete information on externalities to design "better" policies that push regulation closer to this optimum. A tax equal to the marginal net negative externality for shadows that are in excess of those that would be allowed under current FAR policy would be weakly social welfare improving.

I refer to the shadows allowed under current FAR regulations as "baseline" shadow. The proposed tax on shadows created in excess of baseline shadow is weakly social welfare improving in that if current FAR regulation is to the "left" of the optimal regulation (meaning that current height restrictions are more stringent than at the social optimum), developers could build beyond the height allowed under the FAR regulation in exchange for paying a tax for the additional shadows they create. Developers would then equalize their marginal pre-tax profits from height to the tax, increasing social surplus. Figure 1.10 demonstrates the effect of the proposed shadow tax when FAR is to the left of the optimal regulation.

If current FAR regulation is to the "right" of the optimal regulation (meaning current height restrictions are less stringent than at the optimum), this policy would have no effect: since the tax would already be in excess of the pre-tax marginal profit to developers, developers would not build past what is already allowed under current FAR regulations. Figure 1.11 shows the effect of the proposed shadow tax when FAR is to the right of the optimal regulation.

The reason for not proposing to tax baseline shadow already allowed under current FAR policy is to ensure that the policy proposal does not potentially push the policy away from the optimum level of development. Taxing all shadow, including baseline shadow, could potentially improve social welfare in the scenario that current FAR regulation is to the right of the optimal regulation, meaning that that current regulations allow for more development than socially optimal, as the tax in this case would tend to reduce development. However, such a tax on baseline shadow, absent subsidies for positive externalities from development, would have the potential to reduce social welfare. Notably, a tax on shadow does not require a full accounting of positive externalities of height. It is possible that with such a tax on baseline shadow, the tax, being potentially larger than the marginal net negative externality, could reduce new development even when current regulations are already to the left of the optimal regulation. In this scenario, the tax would tend to push policy away from the social optimum rather than towards it. By only taxing shadows in excess of baseline shadows, this tax proposal avoids reducing development away from the social optimum.

In practice, a tax on shadows would effectively act as an indirect transfer of welfare from

Figure 1.10: FAR to the left of the optimal regulation

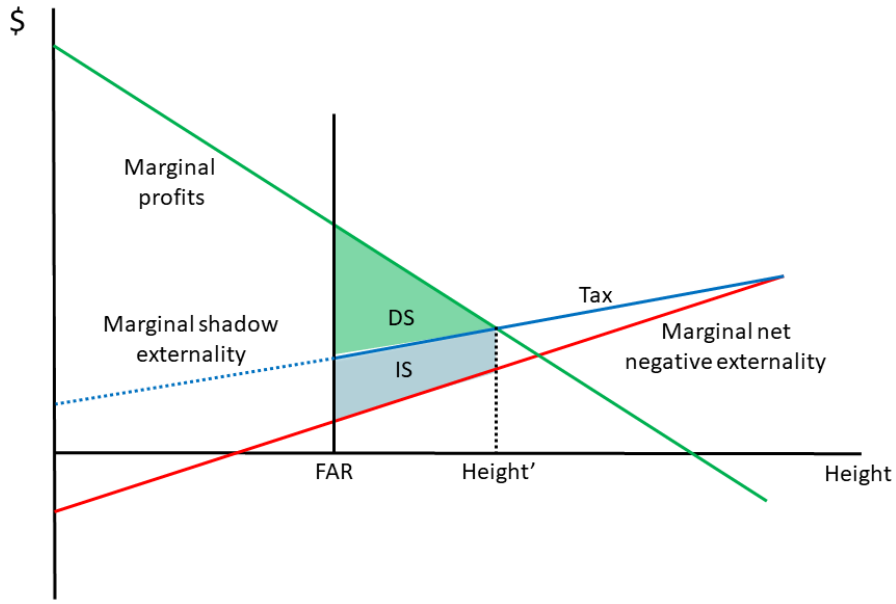

The red line represents the marginal net negative externality. The green line represents the marginal profit to the developer. The blue dotted line represents the marginal shadow externality. The solid blue line is the tax on shadows, which applies only to shadows in excess of FAR. In this simplified example, the FAR policy acts as a quantity restriction on height. The tax results in building heights of Height', closer to the social optimum than FAR. Assuming government surplus is transferred to incumbents, incumbent residents get an additional incumbent surplus IS relative to the FAR policy, represented by the shaded blue area. Developers get an additional developer surplus DS relative to the FAR policy, represented by the shaded green area.

developers to incumbent residents via the public budget. The underlying idea is that if developers build taller than current height restrictions allow, this will impose a larger net negative externality on incumbent residents. Assuming public taxation and spending are designed on behalf of a city's residents to improve their welfare, incumbent residents would indirectly benefit from increases in public revenue. This welfare gain could come through reducing other taxes on incumbent residents, or by increasing spending on public goods. Therefore, the proposed tax frees developers to build taller than current height restrictions allow, but with the tax offsetting the welfare loss to incumbent residents.

Figure 1.11: FAR to the right of the optimal regulation

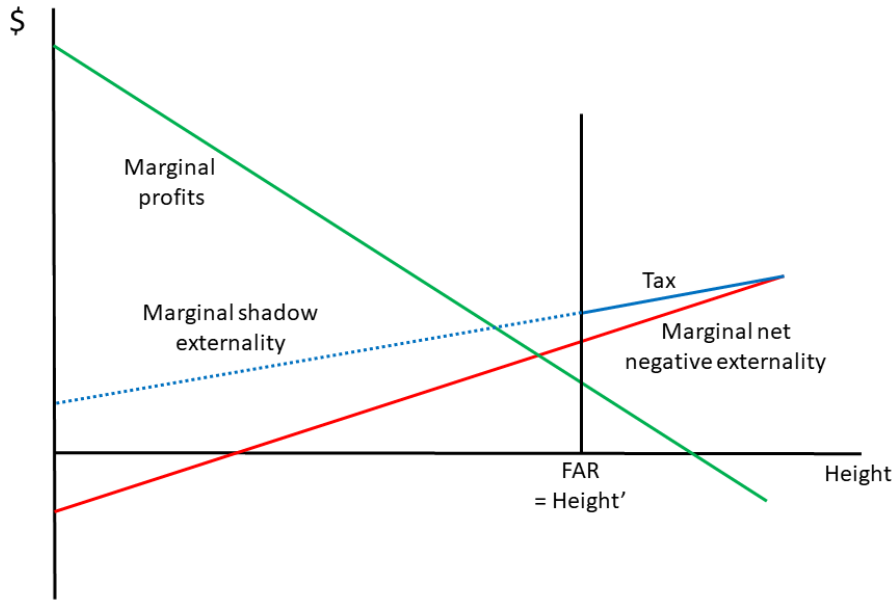

The red line represents the marginal net negative externality. The green line represents the marginal profit to the developer. The blue dotted line represents the marginal shadow externality. The solid blue line is the tax on shadows, which applies only to shadows in excess of FAR. In this simplified example, the FAR policy acts as a quantity restriction on height. Developers would build up to Height', which equals the height allowed under FAR. Since FAR is already to the right of the social optimum, the tax has no effect on building heights.

The remaining question is: how large should this tax be? The tax could in theory be directly based on the expected externality on existing units. That is, one could calculate the increase in shadows on each existing unit, determine the value of those units, and make the tax directly proportional to that external cost. Some issues arise here: one is that each existing unit in potentially many different buildings would need to be valued at the time that any new development project is considered. Another is that as the external cost is proportional to the value of the unit being shadowed, this could disadvantage lower-income incumbent residents, who would on average have units that are less expensive to the developer to cast a shadow on, all else equal.

A more straightforward calculation of the tax would be to make the size of the tax proportional to the value of the new building, rather than directly proportional to the external cost. That is, one would calculate the total new shadow made by the new building, find the value of the new building, and make the tax on this new shadow proportional to the value of the new building. Taxing the new building in proportion to its own value would simplify accounting without needing to value all shadowed units, while also preventing developers from receiving effective tax breaks by building in lower income areas with less expensive existing residential units.

To calculate the shadow created by the new building, one could consider the potential sunlight blocked by the new building. This would be equal to the proportion of unshadowed area on the new building itself. This is equivalent to calculating $1 - \text{shadow}_{rt}$, where shadow_{rt} is defined as as in Section 2, r is the new building, and t is the year of building approval or completion.

To avoid taxing buildings that satisfy the allowable FAR, and to ensure that only new shadows created due to exceeding the allowable FAR are taxed, we would also need to know how much of the shadow created by the new building is due to the building exceeding current FAR regulations. FAR regulations allow buildings of many shapes, which could lead to varying amounts of shadow from buildings satisfying FAR. As such, without an exact method to calculate which specific shadows of the new building are created due to the building exceeding FAR, this amount could be approximated as proportional to the ratio by which the FAR regulation has been exceeded. With FAR representing the FAR of the new building, and FAR' representing the allowable FAR under current regulations, the ratio by which the new building would have exceeded the FAR regulation is

$$\frac{\text{FAR} - \text{FAR}'}{\text{FAR}}.$$

It remains to multiply this ratio by the amount of shadow created by the new building to determine the amount of shadow exceeding FAR. This follows the formula

$$(1 - \text{shadow}_{rt}) \times \frac{\text{FAR} - \text{FAR}'}{\text{FAR}}.$$

This value can then be multiplied by the negative of one of our point estimates from Table 1.1, and again by 100, to determine the applicable shadow tax on the building's value as a percentage of the building's total value.

Case Study - Extell Building

This subsection will consider the shadow cast by a 69-story residential tower currently under development at 50 West 66th Street in Manhattan, which I refer to as the "Extell Building," after its developer, the Extell Development Company. The Extell Building's plans, with a height of 775 feet, created much controversy due to the inclusion of 198 feet of vertical mechanical space in its design, allowing the building to be made taller without counting against its FAR (Kim, 2020). New York City has since amended its laws to limit the loophole to the FAR policy caused by mechanical "voids."[9] A New York State Senate bill that is currently in committee, Senate Bill S3820A, would more strictly define stories and limit mechanical voids for the purpose of FAR regulations: up to 20 feet for a ground floor, up to 12 feet for other floors, and up to 5% of the building's total height for mechanical spaces not counted towards the FAR (New York State Senate, 2019).

In this case study, I will consider the tax as it might be applied to the Extell Building under its current design, assuming a 5% limit on untaxed mechanical space, and further consider what might be the developer's optimal height under the proposed tax. The Extell Building has received $967 million in construction financing, and aims to generate $188 million in sales for the first 22 apartments, with 127 units in total(TRD Staff, 2022). Extrapolating from here, under the assumption that $967 million represents the total construction cost, and that total revenues for the developer are $1085 million (assuming all apartments sell for the same price as the first 22), I estimate developer profits of $118 million.

I next approximate the amount of space in excess of the building's FAR. As the developer was able to buy air rights, the exact FAR allowed exceeds the building's zoning. As such, I will use the ratio of excess height as a stand-in for the exact calculation of the ratio of excess FAR. As the mechanical space is 198 feet and the overall height is 775 feet tall, the building's height without the mechanical space is 577 feet. Assuming the developer had only been allowed up to 5% mechanical space before being taxed, this would lead to an untaxed building height of 607 feet, with 30 feet of mechanical space. I therefore estimate the FAR ratio to be $(775 - 607)/775 = .22$, suggesting that 22% of the 775-foot building's height exceeds its allowable height under a FAR policy that incorporates a 5% limit on mechanical space.

I next approximate the shadow cast onto the Extell Building, shadow_{rt}, where r is the Extell Building, and setting $t = 2014$, since this is when my current shadow data ends. I find that $\text{shadow}_{rt} = .05$. This implies that $(1 - .05) \times .22 = 20.7\%$ of the building should fall under the shadow tax. With a shadow tax rounded slightly based on Table 1.1, Column (3) of 3.75%, the tax would amount to 7.8% of the building's total value, or $84.1 million based on $1085 million in expected sales. Note that this amounts to over two-thirds of the expected total profits from the building: the developer would still go ahead with construction, but

[9] https://www1.nyc.gov/site/planning/plans/voids/voids.page

in exchange, the city would receive $84.1 million in tax revenue and the developer would receive $33.9 million in profits.

What would be the optimal height of the Extell Building, assuming the developer was freed from height limits under the new tax regime? This exercise requires a few additional assumptions. First, I need to determine the value and construction cost of an additional story. For simplicity, I assume a marginal revenue and marginal cost per story equal to average revenue and average cost per story. However, I additionally assume that there are no sales on the mechanical space floors or the lobby. As such, I divide the height of the building without the mechanical space by the height of the building, $(775 - 198)/775 = .74$ to determine the fraction of existing stories that are not mechanical space. Based on this formula, with 69 stories of height, approximately 51 stories are not mechanical space. Removing the lobby, this gives 50 floors that generate sales revenue for the developer. I therefore calculate a revenue per story of $1085/50 = \$21.7$ million. I approximate a cost per story of $967/69 = \$14.0$ million. This leaves an estimated $7.7 million in marginal profit per upper story.

Based on the height of the building and the number of stories, I approximate the height of each story to be $775/69 = 11.2$ feet. With an additional 11.2 feet of height, shadow_{rt} would remain about the same. The excess height ratio, however, would increase to about .23 (accounting for an $11.2 \times 5\%$ increase in allowable mechanical space), leading to the tax rate increasing to 8.2%, and the total additional tax therefore being $.082 \times (1085 + 21.7) - .078 \times 1085 = \$7.1$ million in tax for adding the 70th story, leading to a marginal profit of $0.6 million for the 70th story. The developer would therefore build the 70th story.

I continue this process to determine the optimal height of the building. I find that the difference in tax between a 73-story Extell Building and a 72-story building is $7.8 million, surpassing the estimated net profit per story. Therefore, the optimal height of the Extell Building would be 72 stories, or about 808.6 feet. With an additional three stories, the government would receive $106 million in tax revenue, and the developer would receive $34.9 million in profit.

While the tax on top of the FAR regulation does encourage the Extell Building's developer to build taller than they would have even with the mechanical void, there are two important distinctions that make the newly proposed policy social welfare improving. First, the negative external cost of the mechanical space is transferred back to the city via the tax, which could benefit incumbent city residents through increased public spending or via tax reductions elsewhere. Second, paying the tax frees the developer to convert the mechanical space into additional useful space, thus allowing for a potential increase in the housing supply without needing to create additional shadow.

1.6 Conclusion

This paper measures the externality of urban shadows in New York City by estimating the impact of shadows created by new highrise construction on nearby housing prices. I find that a 10 percentage point increase in average daily shadow is associated with a 3.78% decrease in housing price, and provide robustness checks to demonstrate that this result is not likely to be driven by views. I present a policy that would potentially benefit developers, while offsetting the external costs to incumbent residents on average, pushing us closer to the socially optimal level of development, in the form of a tax on the externality created in excess of that allowed under existing regulation. I then provide a case study for a building that had exploited a previous loophole in current regulations by creating a mechanical "void," and show that while the building would potentially be taller under the proposed tax, the city would collect a substantial proportion of the profit from the building, which it could use to benefit incumbent residents that otherwise might receive the shadow externality, thus pushing development closer to a "socially optimal" level.

Future research in this area could be expanded upon by more precise shadow calculations. This paper made simplifying assumptions in calculating shadows for computational reasons, such as limiting the the number of days for which shadows are calculated, relying on 2D footprint data rather than 3D building models, and limiting the buildings for which shadows are calculated to those over 50-feet tall. A data limitation is that we do not know specifically which story of a given building the residential units in our sample are on, requiring us to estimate the shadow on a given unit using the average shadow received by the building containing it. Such future research could benefit from improved shadow computation that allows for the use of 3D shadow models for both the new buildings creating shadows, and the residential buildings receiving the shadows. Future research could also benefit from data that specifies the specific story or side of a building that residential units are on. In addition, the policy design in this paper uses a tax based on the shadow externality found in this paper. A Pigouvian tax that fully accounts for both negative and positive externalities of height could improve welfare beyond that of the tax proposed in this paper. As such, more research on both the positive and negative externalities of new construction would help policy makers to design better policy regulations for building height.

Chapter 2

Estimating the Relationship Between Inter-Climate Migration and Air Conditioning Adoption[1]

2.1 Introduction

As Earth's climate changes, people may respond to these changes through adaptation. Such adaptations may be in response to changes in productivity, such as a drop in crop yield, or for personal health and comfort, such as to reduce heat stress. Many such adaptations exist, including migration towards milder climates, or investments in temperature-regulating devices such as air conditioning (AC).

In this paper, we investigate the question of how people adapt to changes in climate resulting from migration by investment in air conditioning. In particular, we explore whether the rate of air conditioning adoption is higher in states made up of more migrants experiencing a relatively larger increase in climate-related temperature changes as a result of their migration. To answer this question, we look at data from the United States between 1960 and 1990 on air conditioning adoption by state, and migration between states from census data. We additionally use data on cooling degree days (CDDs) as a measure of how hot a given climate is. Linking data on migration and CDDs, we can determine the average increase (or decrease) in cooling degree days experienced by migrants to a given state. We can then determine if states with migrants experiencing larger average increases in CDDs have relatively higher air conditioning adoption rates.

We find that air conditioning adoption rates are higher in states with migrants who experience larger average temperature increases, controlling for the overall inter-state migrant

[1]The material in this chapter was co-authored with Léopold Biardeau.

Figure 2.1: Evolution of residential AC adoption rates in the contiguous US between 1960 and 2000

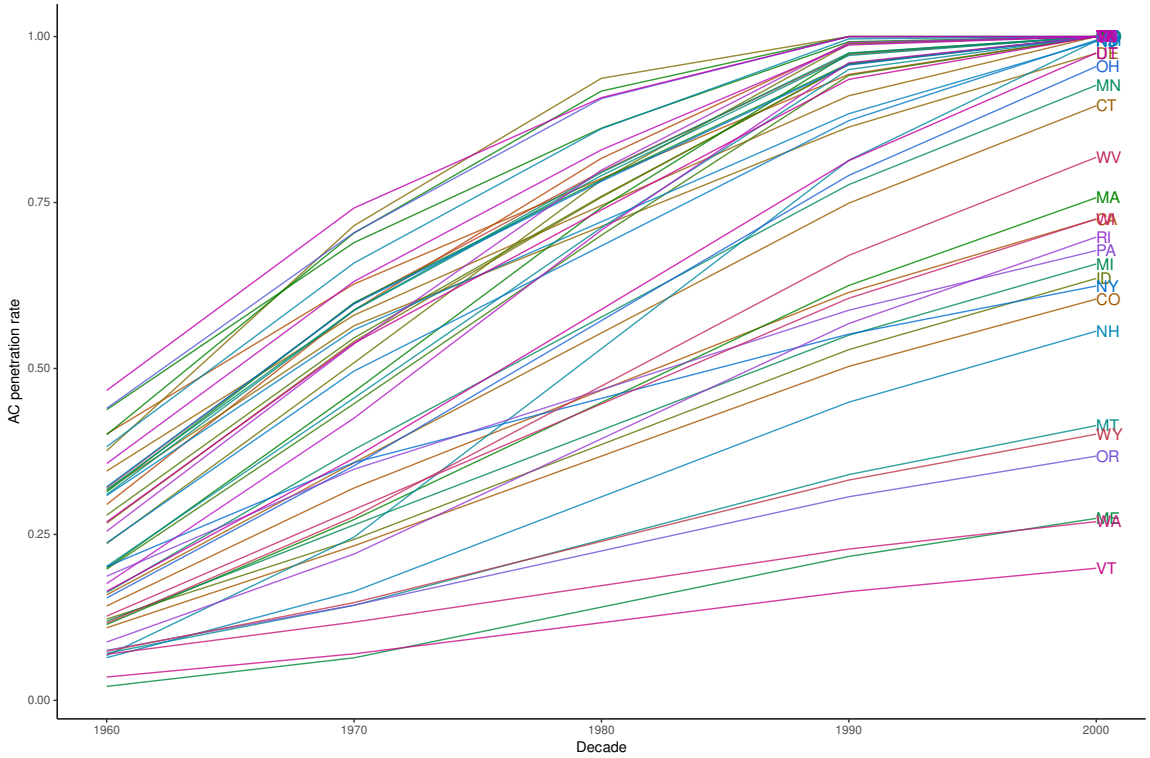

This figure displays the evolution of residential AC adoption between 1960 and 2000 at a decadal level for all contiguous US states (excluding Alaska and Hawaii).

share of the population. This result suggests that migrants who experience relatively larger increases in temperature are more likely to own air conditioners, although some of this effect could be driven by spillovers to other migrant groups or to the local population. Our results imply that relative increases in temperature matter for determining adaptive responses, not just absolute temperature.

It is plausible that reverse causality or omitted variables may be driving our result. One consideration that would imply reverse causality is that states with higher rates of air conditioning adoption are better able to draw migrants who would experience a relatively larger temperature increase due to moving between climates. To provide evidence that our result is not driven by reverse causality or omitted variables bias, we use an instrumental variables approach, instrumenting for migration. Our coefficient estimate is relatively unchanged and remains significant. Although reverse causality does not seem to be driving our result, note

that we do not rule out that air conditioning may attract migrants from relatively cooler climates to relatively warmer climates. This question is an interesting research topic in its own right, but outside of the scope of our analysis.

This paper builds on literatures in air conditioning adoption and migration. With respect to air conditioning, Biddle (2008) investigates some of the potential drivers of air conditioning adoption in the US from the 1950s until 1980, including cooling degree days, although he does not consider migration as a driver of air conditioning adoption. Barreca et al. (2016) show that air conditioning adoption in the US has modified the relationship between extreme temperatures and mortality, reducing heat-related deaths. Auffhammer (2014) shows that air conditioning adoption in China is higher the year following a hot summer.

While this paper does not directly address the causes of migration, we find the literature on migration in response to environmental factors to be relevant to our study. Deschênes and Moretti (2009) find that migration from colder states to warmer states in the US resulted in a decrease in mortality due to a reduction in exposure to cold weather. Chen, Oliva, and Zhang (2022) find that in China, increased air pollution in a county results in outmigration to other counties. Bohra-Mishra, Oppenheimer, and Hsiang (2014) find that increases in temperatures in Indonesia above 25 degrees Celsius result in permanent outmigration within Indonesia. Beine and Parsons (2015), however, do not find evidence of international migration resulting directly from long-run climate variation.

Previous papers combining questions on migration and air conditioning generally consider if air conditioning has resulted in migration to warmer climates. However, the literature finds mixed evidence of this. Rappaport (2007) finds that while people in the US have been moving away from cold winters, and that air conditioning was important in the growth of some states with hot climates, people have also moved towards cool summers. Biddle (2012) does not find strong evidence that air conditioners led to increased migration to warmer climates in the US. He does, however, find evidence that migration to a hot climate is less attractive to those from a cool climate than it is to those from another hot climate. This latter finding provides some justification for our finding that larger positive temperature differentials between origin and destination state would increase air conditioning adoption more for relatively warm states with more migrants from relatively cooler states.

Our paper fits into this broader literature in two ways. First, we may want to know if adaptive responses to climate change are based not only on the static climate at a given point in time, but on the relative changes in climate. This paper adds to the climate literature by suggesting that relative changes are an important determinant of adaptive responses to climate change. That is, even if two regions end up at the same temperature in the next century, if one started from a lower baseline temperature in this century, we may expect to see larger adaptive responses in that region, such as increased rates of air conditioning ownership.

Second, this paper fits into the migration and climate literature by investigating if there is an interaction between migration and adaptive responses to climate differentials. Migration can itself be an adaptation to climate change, but adaptive technologies could in turn facilitate migration to less favorable climates. Therefore, regions with a higher proportion of migrants from different climates may have differing rates of adoption of adaptive technologies when compared to regions with fewer migrants, or compared to regions with a high proportion of migrants but from more similar climates. Our paper helps us to understand this relationship as it pertains to temperature and air conditioning.

This paper proceeds as follows. In Section 2, we describe our data and their sources. In Section 3, we describe our estimation strategy and our instrumental variables approach. In Section 4, we present our results and provide a test for reverse causality, a robustness check, and data simulations. Section 5 concludes.

2.2 Data

Migration History data for individual US states was obtained from America's Great Migrations project, curated by Dr. James Gregory from the University of Washington.[2] Covering the period 1850-2017, it allows for the disaggregation of migrant stocks at the race level for all state pairs. The following categories are present: Asian, Black, Latinx, Native, Other, White. For the purpose of this analysis, we retain the observations concerning all races at the aggregate level, over the period 1960–1990.

The prevalence of residential air conditioning for each state comes from Barreca et al. (2016). Using US Census of Population ownership estimates, the authors initially performed a linear interpolation to obtain data at the state-year level between 1960 and 2004. For the purpose of this analysis, we aggregate the data back at the state-decade level between 1960 and 1990 in order to match the resolution and span of the migration data.

Temperature data came from Berkeley Earth's daily gridded temperatures experimental dataset.[3] We proceed to calculate daily Cooling Degree Days (CDDs), using a threshold of 18.3°C (65°F),[4] for each cell and sum them up at the year level, before calculating the annual average for each decade at the cell level. Last, the estimates are aggregated at the state level by taking the average of all the cells falling within the boundaries of each state. The resulting variable captures the average annual CDDs at the state level for each decade over the period 1960–1990.

Income and population data were obtained at the state-year level from the Bureau of Eco-

[2] https://depts.washington.edu/moving1/migrationhistory-states.shtml

[3] https://berkeleyearth.org/data/#Gridded%20Data

[4] A common threshold used in the AC potential literature, as discussed by Biardeau et al. (2020).

nomic Analysis' Regional Data.[5] The income variable, which corresponds to the aggregate income level in the state, is measured in millions of current USD. Both variables were subsequently averaged at the state-decade level.

Annual Cost-Of-Living (COL) estimates were obtained thanks to Dr. Richard Fording from Berry, Fording, and Hanson (2000). The authors updated their original estimates in 2009 to cover the period 1960–2007. The benchmark of 100 corresponds to the median cost of living in 2007, calculated using the average value of the index for the two states New Mexico and Wyoming. The decadal average was then calculated for each state for the period 1960–1990.

Last, lat-lon coordinates of each state's centroid, contiguities and bilateral distances between pairs of state centroids were calculated directly from the Census Cartographic Boundary shapefiles.

2.3 Empirical Model

In order to measure if a higher proportion of migration from relatively cooler states to relatively warmer states influences the rate of air conditioning adoption in the destination state, we regress the air conditioning adoption rate, AC_{dt}, on the average percent increase in CDDs experienced by migrants to the destination state, weighted by the migrant share of the destination state's population. We model this percent increase in CDDs as

$$\frac{\sum_o \left(\frac{\text{CDD}_{dt}-\text{CDD}_{ot}}{(\text{CDD}_{dt}+\text{CDD}_{ot})/2}\right) \text{Migrants}_{odt}}{\sum_o \text{Migrants}_{odt}} \times \frac{\sum_o \text{Migrants}_{odt}}{\text{Population}_{dt}}. \tag{2.1}$$

Here, o refers to origin state, d refers to destination state, and t refers to time. (Note that we do not include observations with $o = d$ such that a destination state is its own origin state. We therefore use $\sum_o$ throughout this paper as equivalent to $\sum_{o \neq d}$, dropping $o \neq d$ for notational simplicity.) The term $\frac{\text{CDD}_{dt}-\text{CDD}_{ot}}{(\text{CDD}_{dt}+\text{CDD}_{ot})/2}$ inside the sum in the numerator on the left-hand side represents the relative percent increase in CDDs experienced by a migrant from state o to state d, in which the percentage is measured as the change in CDDs, $\text{CDD}_{dt} - \text{CDD}_{ot}$, divided by the midpoint between the CDDs of the two states, given by $(\text{CDD}_{dt} + \text{CDD}_{ot})/2$. Using the midpoint percentage change allows for the percent increase in CDDs when moving from a given origin state to destination state to be equivalent to the percent decrease in CDDs when moving in the opposite direction, i.e. reversing the origin and destination states. We then weight this term by the number of migrants in the destination state who came from a given origin state, Migrants_{odt}, sum across origin states for a given destination state, then divide by the total number of migrants from all origin

[5] https://apps.bea.gov/iTable/iTable.cfm?reqid=70&step=1&acrdn=2

states $\sum_o \text{Migrants}_{odt}$ to get an average percent change in CDDs experienced by all migrants to destination state d, as shown on the left of the multiplication sign $\times$ in Equation (2.1).

As an example of why we use a midpoint percentage change, as opposed to a standard percentage change measured that is measured relative to origin state CDDs, consider a migrant from New York to Florida. As the CDDs in New York are small compared to those in Florida, using a traditional percentage change, there would be a relatively large magnitude percent increase in CDDs experienced by this migrant. Next, consider a migrant from Florida to New York, which is also measured in our sample. Using a traditional percentage change, this migrant, moving from Florida to New York, would experience a relatively small magnitude percent decrease in CDDs, as the CDDs in Florida are large relative to those in New York. Additionally, as states cannot have negative CDDs, this would restrict the CDD change term to the interval $[-1, \infty)$, as one could not have more than a 100% decrease in CDDs. Using midpoint percentage changes, the percent change in CDDs moving from New York to Florida, or from Florida to New York, would be the additive inverse of the other. This would be better suited to ensuring consistently measured magnitudes of the percent change in CDDs between states. As such, we prefer to use midpoint percent changes in our model.

The right-hand-side of Equation (2.1) is the fraction of migrants from the origin states in our sample relative to the destination state's population, giving us the migrant share of the destination state's population. Equation (2.1) simplifies to

$$\sum_o \left(\frac{\text{CDD}_{dt} - \text{CDD}_{ot}}{(\text{CDD}_{dt} + \text{CDD}_{ot})/2} \right) \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}. \tag{2.2}$$

The variable in Equations 2.1 and 2.2 is the independent variable of interest in our regressions, which we refer to as the migration-induced relative change in CDDs.

In addition to this term, our regressions include a variable equal to the migrant share of the destination state's population, $\sum_o \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$ to control for the overall effect of migration on air conditioning adoption, independent of relative changes in climate. Note that this term is equal to the right-hand side of Equation (2.1). Our regression also includes controls for the log of income per capita in the destination state, $\log(\text{Income\_pc}_{dt})$, the consumer price index in the destination state, CPI_{dt}, and the log of the CDDs in the destination state, $\log(\text{CDD}_{dt})$. We also include destination state fixed effects, ν_d, and time fixed effects, τ_t. Our main regression equation is thus

$$\begin{aligned}\text{AC}_{dt} =&\beta_1 \sum_o \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}} + \beta_2 \sum_o \left(\frac{\text{CDD}_{dt} - \text{CDD}_{ot}}{(\text{CDD}_{dt} + \text{CDD}_{ot})/2} \right) \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}} \\ &+ \beta_3 \log(\text{Income\_pc}_{dt}) + \beta_4 \text{CPI}_{dt} + \beta_5 \log(\text{CDD}_{dt}) + \nu_d + \tau_t + \varepsilon_{dt}.\end{aligned} \tag{2.3}$$

where AC_{dt} is the air conditioning adoption rate in state d at time t and ε_{dt} is an error term. β_2 will be our main coefficient of interest, as this tells us the relationship between the migration-induced relative change in CDDs and air conditioning adoption in the destination state.

Instrumental Variables

We additionally use an instrumental variables approach. One possible concern in our main regression using ordinary least squares is that of reverse causality: that a higher rate of existing air conditioning adoption may attract people from relatively cooler states to migrate to relatively warmer states. The possibility of such an effect makes it plausible that an increase in the rate of AC adoption causes an increase in the migration-induced relative change in CDDs.

In addition, it is plausible that our regression could suffer from omitted variables bias. For example, it is possible that lower construction costs in a state results in both an increase in migration to that state due to a lower cost of housing, and in higher rates of AC adoption due to a newer housing stock that is more likely to include central air conditioning. Another possible omitted variable is unmeasured intra-state migration, which might crowd out inter-state migration while simultaneously increasing AC adoption rates. While AC adoption resulting from intra-state migration would be important to measure in understanding the effect of migration on AC adoption, we lack the data to incorporate the effect of intra-state migration into our estimates.

In order to control for the possible effects of reverse causality and omitted variables bias, we make use of an instrumental variables approach to estimate Equation (2.3). While we need to instrument for the entire independent variable of interest, which is the migration-induced relative change in CDDs (Equation (2.2)), the target of our instruments is the relative share of migrants in the destination state from each origin state, $\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. We therefore consider instruments which may cause changes in migration between two states, without directly causing changes in air conditioning adoption in the destination state. CDDs themselves are plausibly exogenous, so we include the percent difference in CDDs between origin and destination states in defining our instruments. We additionally instrument for the control variable representing the migrant share in the destination state population, $\sum_o \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$, as this control variable is made up of the same migration terms instrumented for in our independent variable of interest.

We first describe our instruments for the migrant share of the population control variable. The construction of instruments for the migration-induced relative change in CDDs will be similar in construction, but multiplying the instruments for migrant shares by the relative percent increase in CDDs experienced by a migrant before summing. The instrumental variables for these terms will be made up of instruments for migrant shares from a given origin state, weighted and summed across origin states, for each given destination state.

The instruments for the migrant share of the population for each origin-destination pair, $\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$, include a combination of origin-state-specific variables, including the rate of air conditioning adoption in the origin state, AC_{ot}, the log of income per capita in the origin state, $\log(\text{Income\_pc}_{ot})$, the log of the consumer price index in the origin state, CPI_{ot}, and the log of the CDDs in the origin state, $\log(\text{CDD}_{ot})$. The instruments also include fixed exogenous parameters that might affect migration between an origin and destination state, including the log of distance between the origin and destination state, $\log(\text{Distance}_{od})$, and an indicator for whether or not the origin and destination state share a border, Border_{od}

For each destination state, in order to construct an instrumental variable for the sum of migrant shares from each origin state, we weight each of these instruments by the relative population of the given origin state relative to the total of possible origin state populations in period t, $\frac{\text{Population}_{ot}}{\sum_o \text{Population}_{ot}}$, and then sum across origin states. For example, we construct the instrumental variable representing air conditioning adoption in origin states as

$$\sum_o \text{AC}_{ot} \frac{\text{Population}_{ot}}{\sum_o \text{Population}_{ot}}. \tag{2.4}$$

We similarly construct the instrumental variables representing log income per capita in origin states, the CPI in the origin states, the log of CDDs in the origin states, the log of distance between the origin states and destination state, and indicators for bordering origin states (note that this latter term is no longer an indicator after weighting by population and summing across origin states).

The construction of instrumental variables for our independent variable of interest given in Equations 2.1 and 2.2, the migration-induced relative change in CDDs, follows similarly. We use the same set of instruments as we do for the migrant share of the population for each origin-destination pair. To construct the instrumental variables for a given destination state, for each origin state, we multiply each of the instruments for migrant shares by the relative percent increase in CDDs experienced by a migrant from the origin state to the destination state, in addition to weighting by the relative population of the given origin state, before summing across origin states. That is, we construct the instrumental variable representing air conditioning adoption in origin states as

$$\sum_o \left(\frac{\text{CDD}_{dt} - \text{CDD}_{ot}}{(\text{CDD}_{dt} + \text{CDD}_{ot})/2} \right) \text{AC}_{ot} \frac{\text{Population}_{ot}}{\sum_o \text{Population}_{ot}}. \tag{2.5}$$

As before, we similarly construct the instrumental variables representing log income per capita in origin states, the CPI in the origin states, the log of CDDs in the origin states, the log of distance between the origin states and destination state, and indicators for bordering origin states.

2.4 Estimation Results

Table 2.1 displays the results of an ordinary least squares regression based on Equation (2.3), and an instrumental variables regression using the instruments described in Section 2.3. We cluster our errors by destination state d. In both regressions, the coefficient on the migration-induced relative change in CDDs is positive and significant at confidence levels below 1%. The simplest interpretation of the coefficient on the instrumental variables estimate, although extrapolating outside of the range of the data, would be that in a state made up entirely of migrants who originate from a state with approximately 500 annual CDDs (close to the average in our sample of 475.8) and move to a state with 100 more annual CDDs, the rate of AC adoption in the destination state would on average be $1.13 \times \left(\frac{600-500}{(600+500)/2} \right) = .21$ higher due to the climate shift, which is about a 21 percentage point increase in AC adoption.

The first stage of the instrumental variables regression for the control variable representing the migrant share may suffer from weak instruments: the F-statistic is 8.62. However, the F-statistic on the migration-induced relative change in CDDs is substantially higher, at 39.37. The estimated coefficients in the instrumental variables regression remain fairly close to the OLS results, suggesting that our result is less likely to be driven by reverse causality or omitted variables bias.

If interpreting these results as the effect of a migration-induced relative change in CDDs for an average resident of a given origin state, then one might be concerned about selection – those who can tolerate higher temperatures would be more likely to migrate to relatively warmer states. This would tend to bias our estimate downward, as, unconditional on moving, those with higher heat tolerance would be less likely to invest in air conditioning compared to others, all else equal. Similarly, if those with higher heat tolerance are more likely to migrate to states with lower existing rates of AC adoption, this would also tend to bias our estimate downward, as migrants would sort towards states where existing rates of AC adoption better match their preferences. This implies that when considering the effect of a migration-induced relative change in CDDs on an average origin state resident, the true

Table 2.1: Effect of migrant share and migration-induced relative change in CDDs on AC adoption

	Dependent variable:	
	AC Adoption	
	(1)	(2)
	OLS	IV
$\sum$ Migrant share	0.454^{*}	0.323
	(0.262)	(0.436)
$\sum$ %ΔCDD $\times$ Migrant share	0.840^{***}	1.126^{***}
	(0.232)	(0.276)
log(Income per capita)	0.239^{*}	0.223
	(0.134)	(0.160)
CPI	-0.001	0.001
	(0.003)	(0.003)
log(CDD)	-0.099	-0.152
	(0.125)	(0.137)
Observations	196	196
R^2	0.961	0.959
Adjusted R^2	0.945	0.943
Residual Std. Error (df = 139)	0.068	0.069
Note:	*p<0.1; **p<0.05; ***p<0.01	

Column (1) estimates Equation (2.3) using OLS. Column (2) estimates Equation (2.3) using two-stage least squares, instrumenting for $\sum$ Migrant share and $\sum$ %ΔCDD $\times$ Migrant share using the instrumental variables described in Section 2.3. $\sum$ Migrant share is the migrant share of the population, $\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. $\sum$ %ΔCDD $\times$ Migrant share is the migration-induced relative change in CDDs, $\sum_o \left(\frac{\text{CDD}_{dt}-\text{CDD}_{ot}}{(\text{CDD}_{dt}+\text{CDD}_{ot})/2}\right) \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. Standard errors are clustered by destination state.

coefficient on our independent variable of interest may be larger than the estimates we provide.

Investigating Reverse Causality With a Granger Causality Test

While this paper analyzes the adoption of air conditioners by states with a higher proportion of migrants from relatively cooler states, it is possible that migrants are more likely to move to

locations that already have higher rates of air conditioning adoption. One example of why existing rates of air conditioning adoption could encourage migration is that the expense associated with adopting air conditioning, either by installing central air conditioning or purchasing a new window unit, causes people to prefer locations with a higher stock of pre-installed central air conditioning or used air conditioning units.

In order to test this hypothesis, we run a Granger causality test to see if a lag in air conditioning adoption predicts an increase in the migration-induced relative change in CDDs. A positive and significant outcome would indicate the possible presence of reverse causality. While our instrumental variables approach would help to reduce the impact of such reverse causality on our result regardless, it would nonetheless be interesting to see if such reverse causality exists, and help justify our use of an instrumental variables design.

We follow Holtz-Eakin, Newey, and Rosen (1988) to estimate a Granger test in a panel setting. In particular, rewriting the migration-induced relative change in CDDs from Equations 2.1 and 2.2 as C_{dt} for notational simplicity, we run a regression to estimate the equation

$$C_{dt} - C_{dt-1} = \delta_1(C_{dt-1} - C_{dt-2}) + \delta_2(\mathrm{AC}_{dt-1} - \mathrm{AC}_{dt-2}) + \tau_t + \varepsilon_{dt}. \tag{2.6}$$

Here, a significant and positive value of δ_2 would suggest the presence of reverse causality. We present the coefficient estimates of this regression in Table 2.2.

We find in Table 2.2 that δ_2 is positive and significant. This suggests that air conditioning adoption rates may cause future migration from relatively cooler climates, and so reverse causality may be present. This justifies our use of an instrumental variables approach to our main regression Equation (2.3) in Table 2.1.

Robustness check - Falsification with HDDs

We perform a robustness check to see if the variation in air conditioning adoption is being driven by the relative change in CDDs experienced by migrants, and not potentially the result of another variable that is closely correlated with the difference in CDDs between the origin and destination states of migrants. This acts as a falsification test, and lends credibility to our analysis as we cannot easily instrument for CDDs to rule out omitted variables bias in the relationship between relative CDD change and AC adoption. We begin by using a regression to estimate an equation similar to Equation (2.3), but instead of using the relative change in CDDs, we use an independent variable representing the migration-induced relative change in heating degree days (HDDs). That is, we define our independent variable in this regression as

Table 2.2: Granger causality test for effect of AC adoption on the migration-induced relative change in CDDs

	Dependent variable:
	$C_{dt} - C_{dt-1}$
$C_{dt-1} - C_{dt-2}$	-0.464^{***}
	(0.082)
$\text{AC}_{dt-1} - \text{AC}_{dt-2}$	0.455^{***}
	(0.058)
Observations	98
R^2	0.465
Adjusted R^2	0.447
Residual Std. Error	0.039 (df = 94)
Note:	*p<0.1; **p<0.05; ***p<0.01

This table estimates Equation (2.6) using OLS. $C_{dt} - C_{dt-1}$ is the change in the migration-induced relative change in CDDs in state d between period $t-1$ and period t (similarly for $C_{dt-1} - C_{dt-2}$ as the change between period $t-2$ and $t-1$). $\text{AC}_{dt-1} - \text{AC}_{dt-2}$ is the change in air conditioning rate in state d between period $t-2$ and $t-1$.

$$\sum_{o} \left(\frac{\text{HDD}_{dt} - \text{HDD}_{ot}}{(\text{HDD}_{dt} + \text{HDD}_{ot})/2} \right) \frac{\text{Migrants}_{odt}}{\text{Population}_{dt}} \tag{2.7}$$

CDDs and HDDs are closely correlated, but not identical. When the average temperature changes above 65°F, only CDDs are affected. Conversely, when the average temperature changes below 65°F, only HDDs are affected. Warmer climates, on average, will have more CDDs and fewer HDDs. HDDs thus have a strong inverse correlation with CDDs, with a correlation of -0.928 in our sample.

We should therefore expect the migration-induced relative change in HDDs to be a significant negative predictor of AC adoption in the destination state. However, when including both the migration-induced relative change in CDDs and the migration-induced relative change in HDDs in the same regression, the migration-induced relative change in HDDs should not have a strong influence on the AC adoption rate, as we expect that variation in the migration-induced relative change in CDDs explains any variation in the AC adoption rate that could be explained by the migration-induced relative change in HDDs, plus some additional variation that would be more relevant to AC adoption than that of HDDs.

Table 2.3: Effect of migrant share and migration-induced relative change in CDDs and HDDs on AC adoption

	Dependent variable:		
	AC Adoption		
	(1)	(2)	(3)
$\sum$ Migrant share	0.454*	0.075	0.476**
	(0.262)	(0.300)	(0.231)
$\sum$ %ΔCDD × Migrant share	0.840***		0.906**
	(0.232)		(0.419)
$\sum$ %ΔHDD × Migrant share		−1.244*	0.152
		(0.658)	(0.881)
log(Income per capita)	0.239*	0.242	0.249
	(0.134)	(0.174)	(0.149)
CPI	−0.001	−0.004	−0.001
	(0.003)	(0.004)	(0.003)
log(CDD)	−0.099	0.026	−0.111
	(0.125)	(0.110)	(0.114)
Observations	196	196	196
R^2	0.961	0.958	0.961
Adjusted R^2	0.945	0.942	0.944
Residual Std. Error	0.068 (df = 139)	0.070 (df = 139)	0.068 (df = 138)
Note:			*p<0.1; **p<0.05; ***p<0.01

Column (1) estimates Equation (2.3) using OLS. Column (2) estimates Equation (2.3), but replacing the migration-induced relative change in CDDs with the migration-induced relative change in HDDs. Column (3) estimates Equation (2.3), but including both the migration-induced relative change in CDDs and the migration-induced relative change in HDDs. $\sum$ Migrant share is $\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. $\sum$ %ΔCDD × Migrant share is the migration-induced relative change in CDDs, $\sum_o \left(\frac{\text{CDD}_{dt}-\text{CDD}_{ot}}{(\text{CDD}_{dt}+\text{CDD}_{ot})/2}\right)\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. $\sum$ %ΔHDD × Migrant share is the migration-induced relative change in HDDs, $\sum_o \left(\frac{\text{HDD}_{dt}-\text{HDD}_{ot}}{(\text{HDD}_{dt}+\text{HDD}_{ot})/2}\right)\frac{\text{Migrants}_{odt}}{\text{Population}_{dt}}$. Standard errors are clustered by destination state.

In Table 2.3, we show the OLS results of regressions based on variations of Equation (2.3). Column (1) is the main result based on Equation (2.3), with the migration-induced relative

change in CDDs as the main independent variable. This is the same as Column (1) of Table 2.1. Column (2) takes Equation (2.3) but replaces the migration-induced relative change in CDDs with the migration-induced relative change in HDDs. Column (3) takes Equation (2.3) but includes separate terms for both the migration-induced relative change in CDDs and the migration-induced relative change in HDDs. We cluster errors by destination state in all three columns.

As can be seen in Column (2), the migration-induced relative change in HDDs inversely predicts the rate of AC adoption, but with a much larger standard error relative to the coefficient and therefore less significance as compared to the estimate for the migration-induced relative change in CDDs in Column (1). Column (3) shows that including both the migration-induced relative change in CDDs and the migration-induced relative change in HDDs in our regression leaves the coefficient estimate on the migration-induced relative change in CDDs roughly the same as in Column (1), albeit with a somewhat larger standard error, but still significant at a 5% level. On the other hand, the coefficient on the migration-induced relative change in HDDs changes in sign and shrinks in magnitude relative to Column (2), while its standard error increases, removing its significance as a predictor of AC adoption.

The regressions in Table 2.3 thus provide evidence that it is unlikely that the effect of the relative change in CDDs experienced by migrants on AC adoption rates is driven by another variable that closely correlates with the relative CDD change experienced by migrants. We do not additionally provide a version of these regressions using the instrumental variables approach used in Column (2) of Table 2.1, as the migration-induced relative change in HDDs would require the same set of instruments as the migration-induced relative change in CDDs.

Simulation

We further consider the presence of potential model-generated biases that could manifest themselves when dealing with panel data. Such biases could generate spurious correlations that would lead us to attribute a causal effect to a variable when there is in fact none.

To mitigate the risk that such biases are driving the results described in the previous sections, we perform 10,000 randomization tests in which we randomly reallocate three treatment variables used in the analysis (CDDs in the origin state, CDDs in the destination state, and migrant stocks in each destination state) in three different ways, before re-estimating Equation (2.3) using the OLS and IV regressions that we used in Table 2.1, a procedure described by Hsiang and Jina (2014). The coefficients on the migrant share of the population ($\hat{\beta_1}$) and the migration-induced relative change in CDDs ($\hat{\beta_2}$) obtained through the OLS and IV specifications are then plotted and compared to the ones presented earlier in the analysis to show how likely it was that one would have gotten the coefficients that were presented in Table 2.1 had the data been randomly distributed rather than how it was actually observed. As in Hsiang and Jina (2014), the randomization is performed in the three following ways :

1. **Entire sample:** The treatment variable to be reassigned is randomly assigned across time (decade) and space (origin/destination state, depending on the treatment variable).
2. **Between panel variable:** In this case, we reassign the respective variable to another state, while preserving the decade that the observation was associated with. This particular randomization procedure preserves the time structure within the data, thus allowing us to check whether temporal trends might lead to spurious correlations.
3. **Within panel variable:** For this third and last procedure, we shuffle each state's time-series, while ensuring that the observations remain assigned to their respective state. This particular procedure allows us to test whether cross-sectional patterns might generate spurious correlations.

The distributions of the $\widehat{\beta_1^{OLS}}$, $\widehat{\beta_2^{OLS}}$, $\widehat{\beta_1^{IV}}$, and $\widehat{\beta_2^{IV}}$ coefficients obtained after shuffling the underlying variables CDD_{ot}, CDD_{dt} and Migrants_{odt} 10,000 times each, for each of the three randomization procedures detailed above, are shown in Figures 2.2 and 2.3.

These simulations show that the coefficients associated with the actual data would have been very unlikely to be found had the data been distributed any other way. Indeed, with the exception of the within randomization procedure[6], our estimates fall in the tails of the distributions of the coefficients obtained after randomizing the data. Two-sided t-tests systematically reject the hypothesis that the mean of the distribution could fall within an interval that would contain the estimates obtained when using the actual data.

Taken together, this analysis implies that the results obtained in Table 2.1 rely on the particular data generating process that yielded the data we gathered, as other data generating processes would have given very different results on average to those that we presented.

2.5 Conclusion

In this paper, we show that a higher proportion of migrants from relatively cooler climates increases a state's air conditioning adoption rate compared to states with a smaller proportion of migrants from relatively cooler states. This effect is positive and significant, and separate from a pure migration effect, which we control for. This suggests that larger relative changes in climate may lead to larger adaptation to climate change. Climate change may therefore result in larger adaptive responses to future temperature increases in a given location than could be predicted by comparing to the adaptive response in a location with a similar temperature today.

[6]Which is based on the permutations of four decades for each state, yielding 24 possible permutations in total.

Figure 2.2: Distributions of coefficient estimates for migrant share of the population from 10,000 randomizations

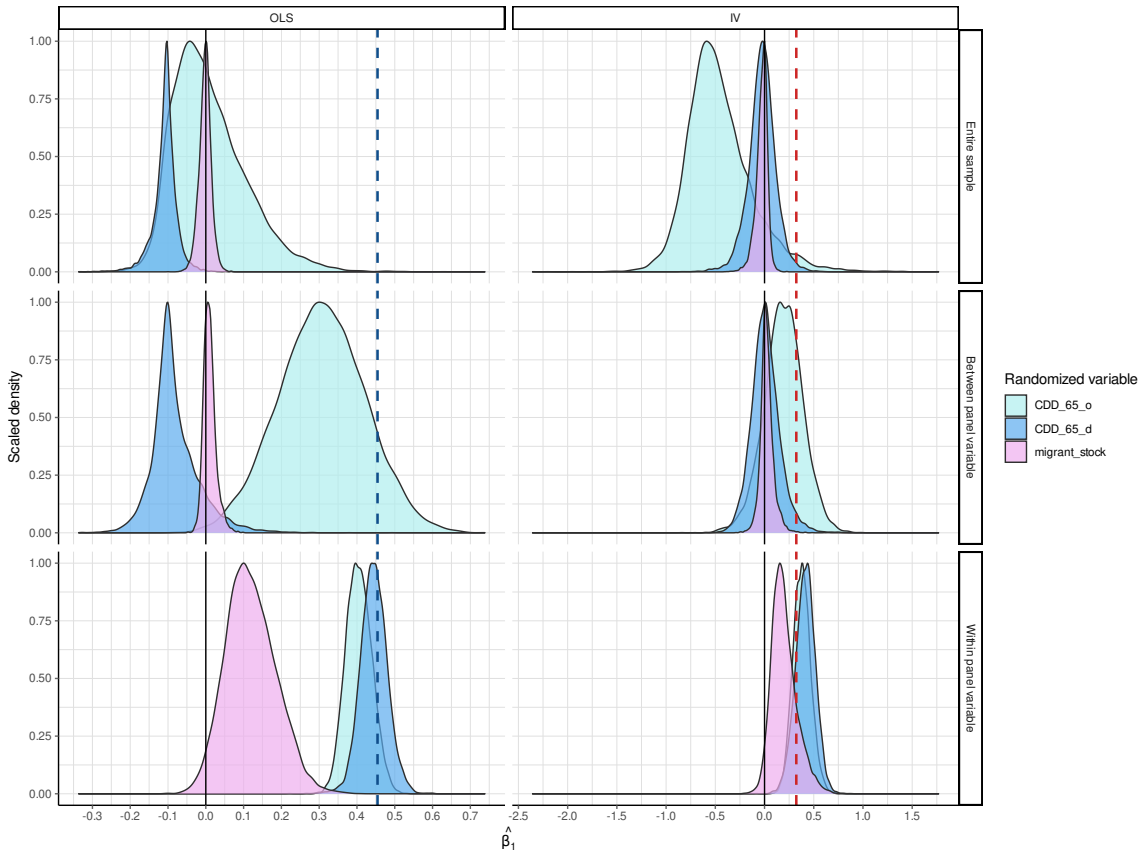

Distributions of the coefficients associated with the migrant share of the population obtained after the variables CDD_{ot} (CDD_65_o), CDD_{dt} (CDD_65_d) and Migrants_{odt} (migrant_stock) were randomized in three different ways (entire sample, between panel variable, and within panel variable) 10,000 times each. After each randomization, we re-estimated Equation (2.3) on the resampled data using the OLS and IV regressions that we used in Table 2.1. The solid black vertical line marks a value of zero for the β coefficient. The blue vertical dotted line corresponds to the value of $\widehat{\beta_1^{OLS}}$ estimated by OLS using the true observations ($\widehat{\beta_1^{OLS}} = 0.454$), while the red vertical dotted line corresponds to the value of $\widehat{\beta_1^{IV}}$ estimated by IV using the true observations ($\widehat{\beta_1^{IV}} = 0.323$).

Figure 2.3: Distributions of coefficient estimates for migration-induced relative change in CDDs from 10,000 randomizations

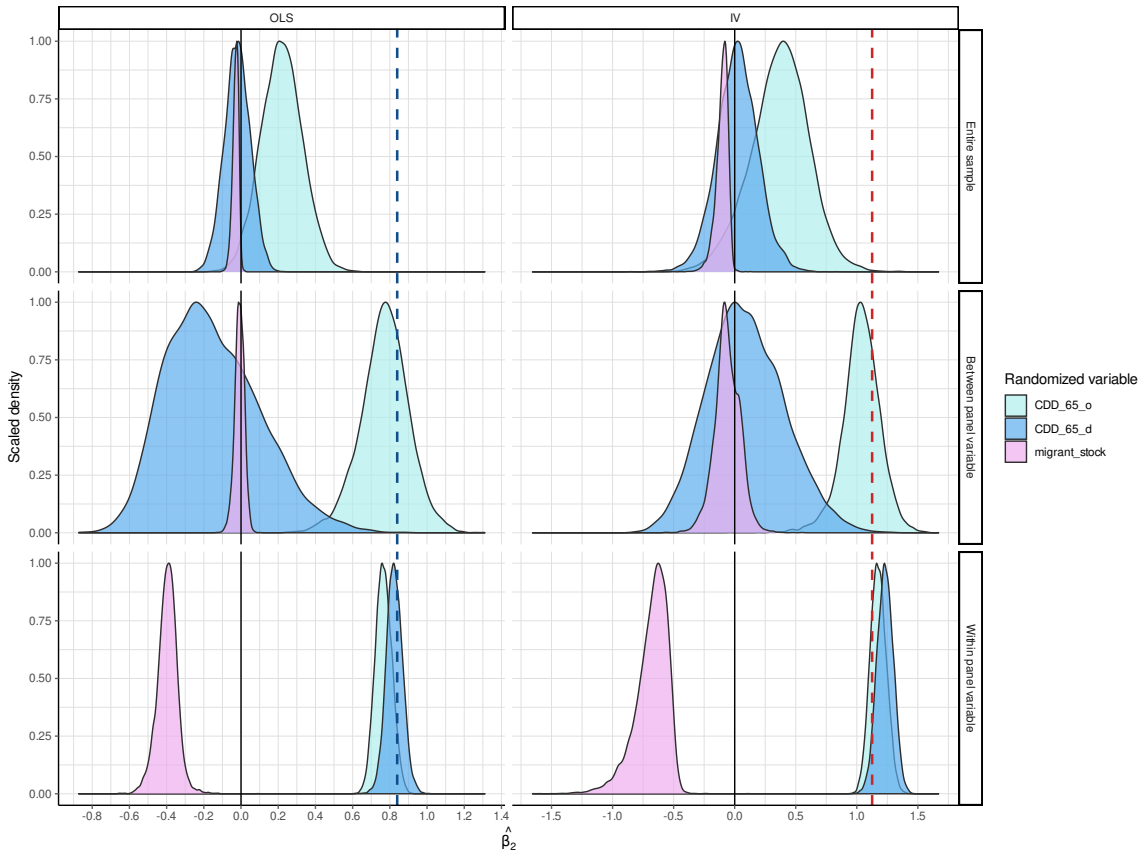

Distributions of the coefficients associated with the migration-induced relative change in CDDs are obtained after the variables CDD_{ot} (CDD_65_o), CDD_{dt} (CDD_65_d) and Migrants_{odt} (migrant_stock) were randomized in three different ways (entire sample, between panel variable, and within panel variable) 10,000 times each. After each randomization, we re-estimated Equation (2.3) on the resampled data using the OLS and IV regressions that we used in Table 2.1. The solid black vertical line marks a value of zero for the β coefficient. The blue vertical dotted line corresponds to the value of $\widehat{\beta_2^{OLS}}$ estimated by OLS using the true observations ($\widehat{\beta_2^{OLS}} = 0.840$), while the red vertical dotted line corresponds to the value of $\widehat{\beta_2^{IV}}$ estimated by IV using the true observations ($\widehat{\beta_2^{IV}} = 1.126$).

Our results also demonstrate how migrants may shift the adoption rate of adaptive technologies based on differences between their home climate and destination climate. As it is often assumed that air conditioning made warmer climates more attractive to those from cooler climates, it could be expected that migrants from cooler climates would invest more heavily in air conditioning than those from already warmer climates. This paper provides evidence that higher rates of migration from relatively cooler climates increases adoption of adaptive technologies in the destination.

Future research on this topic could benefit from the use of micro-level data on air conditioning adoption by individual migrants or groups of migrants from specific origin locations. Our paper relies on macro-level census data by state, so we are unable to make inferences regarding the specific increase in statewide air conditioning adoption that can be attributed to adoption by out-of-state migrants. Micro-level data could allow us to distinguish if individual migrant groups are more likely to adopt air conditioning in the destination state. Micro-level data could also allow us to investigate if some of the observed increase in air conditioning adoption in states with a higher proportion of migrants from relatively cooler states may be driven by spillover effects, such that locals or migrants groups from relatively warmer climates adopt more air conditioning in response to the higher rates of adoption by migrants from relatively cooler states.

Chapter 3

Evaluating the Effectiveness of Very Large Marine Protected Areas at Deterring Fishing Effort[1]

3.1 Introduction

In October 2010, the tenth meeting of the Conference of the Parties adopted in Nagoya (Japan) the Strategic Plan for Biodiversity. The plan included the Aichi Biodiversity Targets, a set of goals named after the prefecture that hosted the meeting. Aichi Target 11 of the United Nations' Convention on Biological Diversity (CBD) stipulated that 10% of the ocean should be protected by some form of Marine Protected Area (MPA) by 2020 (Convention on Biological Diversity (CBD), 2010). While the 10% target was not achieved by 2020, the global coverage of MPAs did more than triple, from 2.5% coverage in 2010 to 7.7% coverage by 2020, totaling an additional 19.1 million km^2 (UN Environment Programme WCMC and IUCN, 2020).

MPAs are marine areas where certain activities are limited, or prohibited, in an attempt to protect part, or all, of the natural resources it contains. They are classified in various categories by the International Union for Conservation of Nature (IUCN) (International Union for Conservation of Nature (IUCN), 1988). Marine reserves (IUCN category Ia), also known as "no-take" areas, refer to MPAs in which fishing, hunting, or collecting are entirely prohibited in order to protect sensitive habitats or threatened species.

There is scientific evidence that MPAs *can* be effective at increasing biomass, species richness, and population size within the boundaries of the reserves (Lester et al., 2009; Zupan et al., 2018). In particular, marine reserves appear to be the most effective type of MPA when it

[1]The material in this chapter was co-authored with Léopold Biardeau and David Zilberman.

comes to achieving conservation targets (Costello and Ballantine, 2015; Sala and Giakoumi, 2017; Claudet et al., 2008; Edgar et al., 2014).

However, this body of knowledge has focused on smaller MPAs. This is not surprising, given the fact that 87% of the 13,000+ existing MPAs are no larger than 100km^2, while 95.7% are smaller than 1,000 km^2 (MCI, 2019). Recent years however have seen the emergence of Very Large Marine Protected Areas (VLMPAs), MPAs spanning more than 100,000 km^2. Establishing large MPAs has become a new trend in an attempt to meet the CBD conservation targets (Boonzaier and Pauly, 2016; Dulvy, 2013; Magris and Pressey, 2018; Toonen et al., 2013). As such, 31 of the 35 VLMPAs created by 2019 have been created since 2009 (MCI, 2019), as shown in Panel **B** of Figure 3.1.

There exists scientific support for large MPAs. In theory, they may better at protecting diverse ecosystem than their smaller counterparts, since their size makes them more likely to encompass species' ranges (O'leary et al., 2018). Even so, VLMPAs are not yet well-understood, an observation shared by researchers (Gruby et al., 2016; Ban et al., 2017) and policymakers alike. As such, Dr. Jane Lubchenco, the head of the US National Oceanic and Atmospheric Administration declared in 2011: "*We don't have the resources that we need to actually monitor, enforce and understand these areas*" (Cressey, 2011).

A lack of enforcement could lead to a failure in meeting the initial objectives of the creation of MPAs (Jone, Qiu, and De Santo, 2011; Dureuil et al., 2018; Zupan et al., 2018). Indeed, the effectiveness of MPAs in providing ecological benefits has been shown to be eroded by illegal fishing effort (Arias et al., 2015; Pollnac et al., 2010; Bergseth, Russ, and Cinner, 2015; Bergseth et al., 2018).

In the meantime, recent years have marked the emergence of new data which have offered new means to assess fishing effort across the globe. In particular, the Global Fishing Watch (GFW), a non-profit organization, has made available to the public a database which tracks vessels using their Automatic Identification System (AIS) data. This database allows users to track individual fishing vessels at an unprecedented spatio-temporal resolution[2] and manages to detect whether these vessels are currently fishing, by looking at their speed and trajectory (Kroodsma et al., 2018) (see Panel **A** of Figure 3.1 to see the evolution of the number of fishing hours and vessels tracked by GFW). A growing number of studies have used this database to study fishing effort in a variety of contexts (Guiet et al., 2019; Queiroz et al., 2019; Dureuil et al., 2018; Englander, 2019; Bradley et al., 2019; McDermott et al., 2019).

In this paper, we use data on fishing effort gathered by GFW to study whether the eight "no-take" VLMPAs implemented between 2012 and 2018 (Figure 3.2 and Table 3.1) have been successful at deterring fishing effort. We also investigate the characteristics of the vessels

[2]The data are available at a $0.1° \times 0.1°$ resolution, on a daily basis from January 2012 to December 2020.

Figure 3.1: Evolution of tracked vessels and fishing hours and expansion of VLMPAs over time

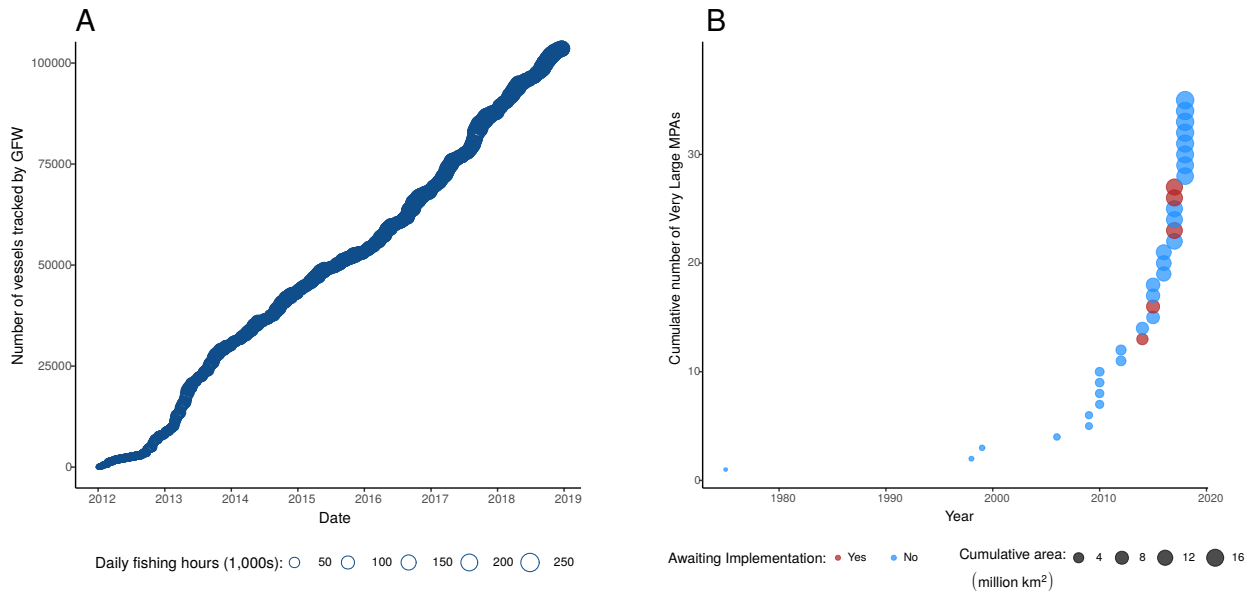

Figure (A) presents the evolution of the number of vessels and fishing hours tracked by GFW over the period from 2012 through 2018. The horizontal axis represents the date. The vertical axis represents the number of vessels tracked. The size of the (overlapping) dots represents the number of daily global fishing hours on each date. Figure (B) presents the expansion of all types of Very Large Marine Protected Areas over time between 1975 and 2019. VLMPAs now cover 17,002,904 km^2, which includes 7,802,962 km^2 of marine reserves, i.e. areas where no fishing is permitted (also known as "no-take" areas).

associated with illegal fishing effort in these protected areas.

Our analysis shows that the introduction of VLMPAs has reduced overall fishing effort in these important ecological areas, although the impact varies by region. Since each VLMPA is supporting its own unique marine life, it is important to enforce policies that would reduce fishing throughout all of these reserves.

This paper proceeds as follows. Section 2 provides an overview of our data. Section 3 analyzes the average reduction in fishing effort across VLMPAs. Section 4 investigates the displacement of fishing effort within each VLMPA. Section 5 analyzes the profile of infringing vessels. Section 6 concludes.

Figure 3.2: Eight Very Large Marine Protected Areas (VLMPAs) in sample

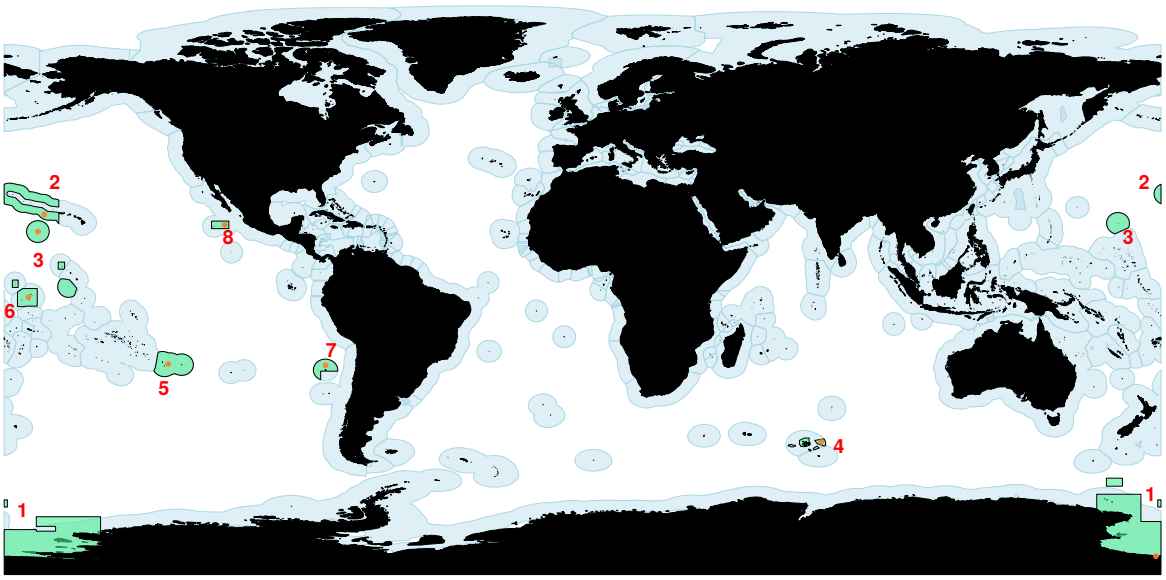

VLMPAs are shaded in green. Exclusive Economic Zones (EEZ) are shaded in blue. The deepest location inside each of the eight VLMPAs is shown in orange. **1** corresponds to the Ross Sea Protected Area, **2** to the Papahānaumokuākea Marine National Monument, **3** to the Pacific Remote Islands Marine National Monument, **4** to the Terres Australes Françaises, **5** to the Pitcairn Islands Marine Reserve, **6** to the Phoenix Islands Protected Area, **7** to the Nazca-Desventuradas Marine Park, and **8** to the Revillagigedo National Park.

3.2 Data

Data on daily fishing effort are provided by the Global Fishing Watch (GFW), a non-profit organization that offers near real-time tracking of global commercial fishing activity. The database gathers Automatic Identification System (AIS) data gathered from ships at sea. It allows to track individual vessels by using their individual Maritime Mobile Service Identity (MMSI) and is available at a daily, $0.1° \times 0.1°$ resolution from January 2012 to December 2018. Fishing effort detection is based on a machine learning algorithm that tracks the behaviors associated with specific fishing activities (long-lining, purse seining, trawling, etc.),

a methodology described in Kroodsma et al. (2018).

MPA shapefiles are obtained from UNEP-WCMC and IUCN (2019). The land polygons drawn in Figure 3.2 are downloaded from *Natural Earth*. We obtain EEZ boundaries from Flanders Marine Institute (n.d.).

Seven out of the eight areas under consideration in this paper are managed by one of only six sovereign states, including one Small Island Developing State (SIDS) (see Table 3.1). The world's largest VLMPA, the Ross Sea Protected Area, is managed by the Commission for the Conservation of Antarctic Marine Living Resources (CCAMLR), an international commission composed of 26 Members. Table 3.1 shows that it has taken on average 525 days (17.5 months) for the no-take areas to come into effect after the original announcement that an area would eventually become protected was made.

Figure 3.3 shows the daily sum of fishing hours per 1,000 km^2 in each of the 8 VLMPAs over the period 2012 to 2018. It distinguishes hours fished by vessels tracked by GFW prior to the announcement of the creation of each of the areas from hours fished by vessels first tracked after the announcement. One can see that fishing intensity varies widely between the 8 areas. Prior to the date when the the announcement of the creation of the Phoenix Islands Protected Area became public, fishing effort averaged 0.263 hours per 1,000 km^2, as opposed to 0.013 hours per 1,000 km^2 for the extended portion of the Papahānaumokuākea Marine National Monument. For comparison, the Global Fishing Watch has measured on average 83,991 daily fishing hours over the period 2012–2018 across all oceans. Given that oceans span approximately 360 million km^2, this would yield 0.223 hours (14 minutes) of fishing effort per 1,000 km^2, a conservative estimate of fishing effort intensity (Weimerskirch et al., 2020), given that many active fishing vessels are yet to be tracked by AIS. At a glance, it would appear that some areas experience overall declines in the intensity of fishing effort after the fishing ban comes into effect (Phoenix Islands Protected Area, Revillagigedo National Park and Ross Sea Protected Area) while others do not (Nazca-Desventuradas Marine Park, Pacific Remote Islands Marine National Monument, Papahānaumokuākea Marine National Monument,Pitcairn Islands Marine Reserve, Terres Australes Françaises). In spite of the apparent heterogeneity in the intensity of fishing effort over time between these eight VLMPAs, it would appear that on average, fishing effort dropped after the enforcement of the ban relative to the period preceding the original announcement, especially when one restricts their attention to the vessels that were tracked prior to the announcement date (blue dots in Figure 3.3). This is highlighted further in Figure 3.4 by averaging the daily sums fishing hours for all eight VLMPAs.

The type of fishing effort practiced in each of these eight areas also varies, as shown in Table 3.2. The table nonetheless shows that the most prevalent fishing methods in the VLMPAs are set longlines and drifting longlines, regardless of the period under consideration (pre-announcement, post-announcement & pre-implementation and post-implementation).

Table 3.1: Characteristics of the eight VLMPAs under consideration.

VLMPA name	Managing Country/Entity	Announcement Date	Implementation Date	Size (millions km^2) of no-take area
Ross Sea Protected Area	CCAMLR (International)	October 27, 2016	December 1, 2017	1.12
Papahānaumokuākea Marine National Monument	United States	January 29, 2015	August 26, 2016	1.15
Pacific Remote Islands Marine National Monument	United States	May 20, 2014	September 25, 2014	1.06
Pitcairn Islands Marine Reserve	United Kingdom	March 18, 2015	September 12, 2016	0.83
Terres Australes Françaises	France	June 14-15, 2016	March 31, 2017	0.12
Phoenix Islands Protected Area	Kiribati	May 9, 2014	January 1, 2015	0.41
Nazca-Desventuradas Marine Park	Chile	February 23, 2013	August 24, 2016	0.30
Revillagigedo National Park	Mexico	July 17, 2016	November 24, 2017	0.15

The Announcement date corresponds to the earliest evidence we could find that an area would eventually be designated to become a marine protected area (see Appendix B.2 for the timelines leading to the creation of each VLMPA). The Implementation date corresponds to the day when the fishing ban becomes effective. The Size column indicates the size of the "no-take" area withing each VLMPA. The VLMPAs can also be comprised of areas where some type of fishing effort is allowed.

Figure 3.3: Evolution of daily hours of fishing effort per 1,000 km^2 in each VLMPA

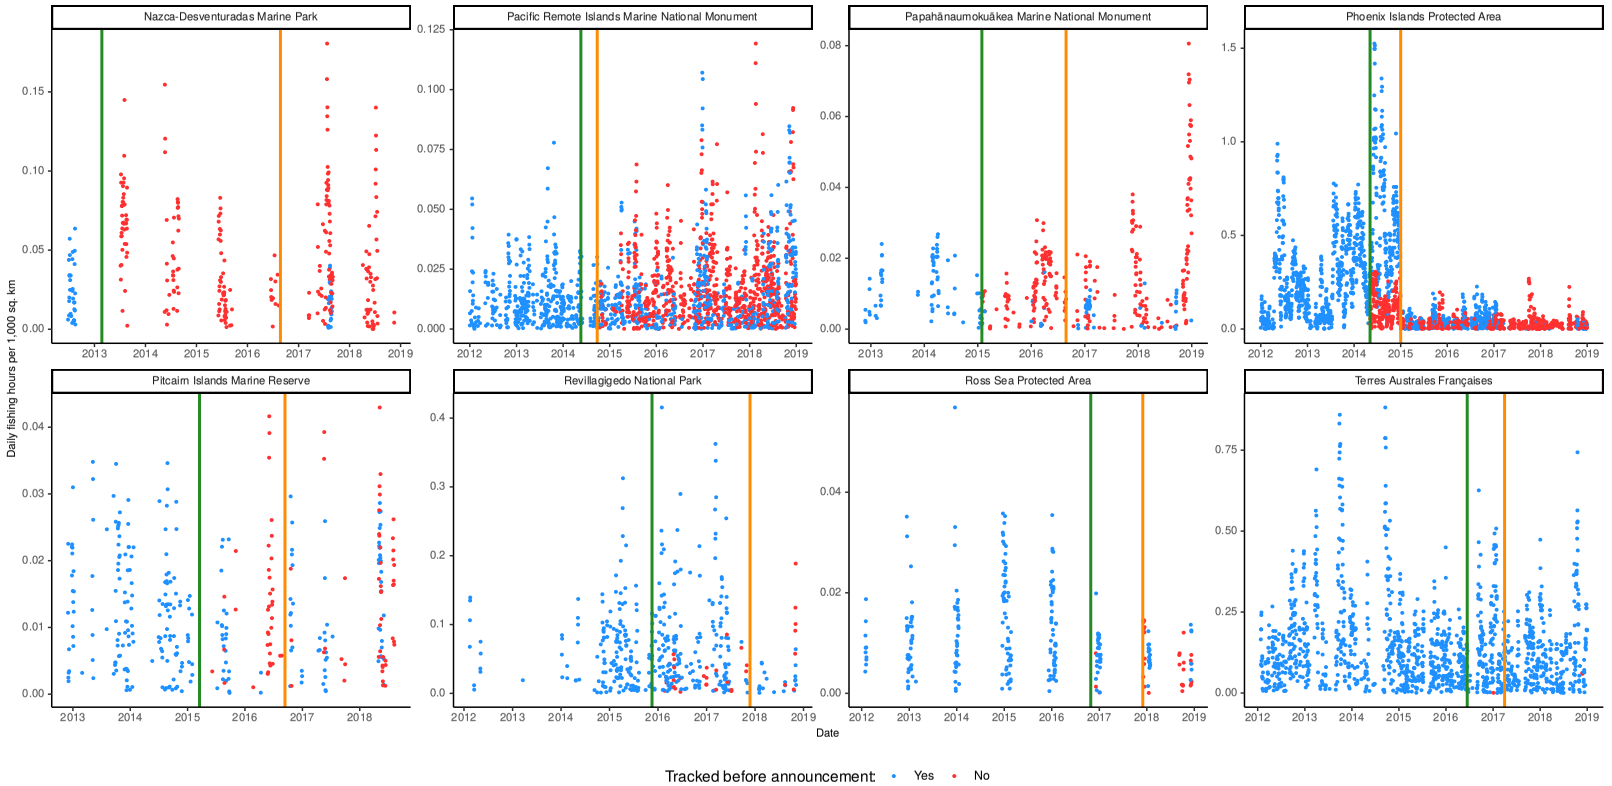

Blue dots measure the sum of fishing hours spent by vessels emitting prior to the announcement of the creation of an VLMPA. The date announcing a possible implementation implementation of a fishing ban is shown in green. Red dots show the sum of fishing hours spent by vessels emitting for the first time after the announcement. The solid orange vertical line corresponds to the date when the fishing ban becomes effective in each VLMPA. The length of the vessels tracked prior to the announcement averages 54.95 meters while that of those tracked after it averages 48.58 meters.

Figure 3.4: Evolution of average daily hours of fishing effort per 1,000 km^2 in all VLMPAs

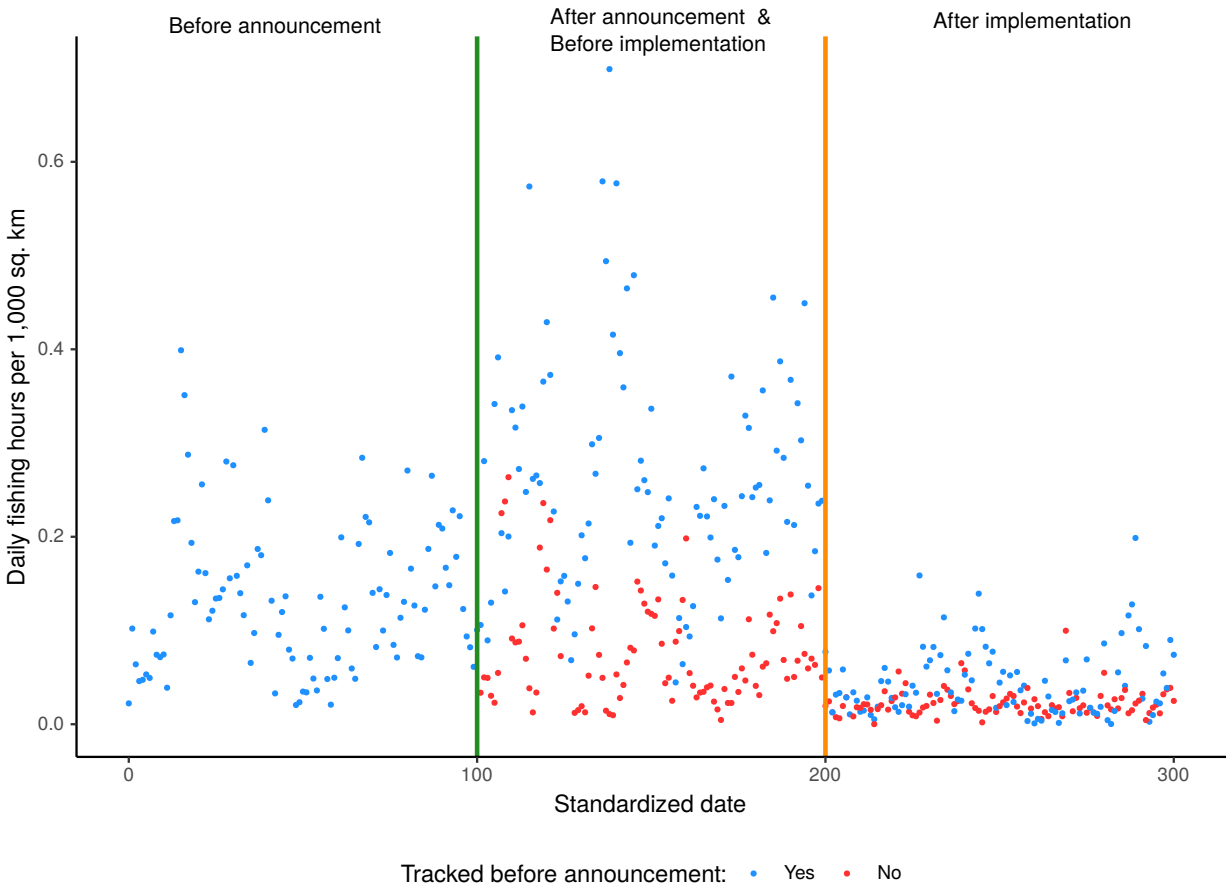

Blue dots measure the sum of fishing hours spent by vessels emitting prior to the announcement of the creation of an VLMPA. Red dots show the sum of fishing hours spent by vessels emitting for the first time after the announcement. The dates have been standardized so that each period – before announcement, between announcement & implementation, and after implementation – would last for 100 days, meaning that the period ranging from January 1, 2012 to December 31, 2018 corresponds to 300 days under this standardized date format. The announcement of the creation of an VLMPA occurs on the 100$^{\text{th}}$ day (solid green line). The implementation occurs on the 200$^{\text{th}}$ day (solid orange line).

Trawling, a particularly destructive fishing method where a net is dragged across the ocean sea floor, is observed at some point between 2012 and 2018 in three out of the eight VLMPAs, and even accounts for 1% of all of the illegal fishing effort monitored in the Terres Australes Françaises VLMPA.

3.3 Estimating the Average Reduction in Fishing Effort[3]

We begin by investigating the average deterrence effect across VLMPAs by estimating the reduction in the average daily hours fished per 1,000 km^2 in the VLMPAs in our sample after the ban comes into effect, compared to the average daily hours fished per 1,000 km^2 prior to the announcement of the ban. We additionally estimate the change in average daily fishing hours during the time period between the announcement of the VLMPAs and their implementation (announcement period). We also consider the possibility of spillovers of fishing activity from VLMPAs into their surrounding waters after implementation of the ban by estimating a separate announcement and implementation effect for "donut" shaped areas of varying sizes surrounding the VLMPAs, ranging from the waters within 10 km of a VLMPA, up to the waters within 50 km of a VLMPA.

As shown in panel **A** of Figure 3.1, the cumulative number of vessels tracked by the GFW database has skyrocketed over the years 2012 – 2018, as countries started mandating their fleets to acquire AIS tracking devices[4] Incorrectly assuming that a vessel only started fishing upon being tracked by AIS could lead to biases in our analysis. For instance, we cannot rule out that vessels appearing after the implementation of an MPA were not fishing prior to being monitored. Should such vessels fish in the area yet to be announced respect the ban and get tracked by the GFW after the announcement, our results would underestimate the extent of the deterrence effect. Therefore, our regression only considers fishing effort during the period after 2014, and only includes fishing effort by ships that were tracked prior to January 1, 2014. This 2014 cutoff strikes a balance between allowing us to include the pre-announcement periods for most of the considered VLMPAs in the analysis (all but one, Nazca-Desventuradas Marine Park, which was announced prior to 2014, although it was not implemented until 2016[5]), while still leaving us with fishing data for approximately 28.5% of all ships that GFW began tracking between 2012 and 2019.

We use a difference-in-difference design with multiple treatment effects to simultaneously estimate the effect of the implementation of the VLMPA on fishing effort inside the VLMPA, the effect of implementation of the VLMPA on the donut immediately surrounding the VLMPA

[3]This section was produced using updated data from GFW. Other sections have not yet been edited to include the updated data.

[4]This can also be inferred from Figure B.2.

[5]See Table 3.1.

Table 3.2: Breakdown of fishing effort in each of the 8 VLMPAs

MPA name	Prevalence Rank	Before announcement		Between announcement and implementation		After implementation	
		Fishing type	Percentage total	Fishing type	Percentage total	Fishing Type	Percentage total
Ross Sea Protected Area	1.	Set longlines	97.2	Set longlines	93.2	Set longlines	73
	2.	Drifting longlines	2.3	Drifting longlines	6.8	Drifting longlines	27
	3.	Tuna purse seines	0.4	/	/	/	/
Papahānaumokuākea Marine National Monument	1.	Drifting longlines	69.6	Drifting longlines	91.3	Drifting longlines	98.3
	2.	Pole and line	27.7	Tuna purse seines	8.6	Trolling	1
	3.	Unidentified	2.7	Other fishing	0.2	Research	0.6
Pacific Remote Islands Marine National Monument	1.	Drifting longlines	91.2	Drifting longlines	75.1	Drifting longlines	92.3
	2.	Other fishing	4.9	Tuna purse seines	15.2	Tuna purse seines	5.9
	3.	Pole and line	2.4	Pole and line	9.7	Pole and line	0.9
Pitcairn Islands Marine Reserve	1.	Drifting longlines	100	Drifting longlines	99.9	Drifting longlines	100
	2.	/	/	Trawling	0.1	/	/
	3.	/	/	/	/	/	/
Terres Australes Françaises	1.	Set longlines	72.6	Set longlines	99.9	Set longlines	96.6
	2.	Trawling	26.6	Trawling	0.1	passenger	2.1
	3.	Passenger	0.7	/	/	Trawling	1
Phoenix Islands Protected Area	1.	Drifting longlines	97.2	Drifting longlines	86.1	Tuna purse seines	62
	2.	Tuna purse seines	2.7	Tuna purse seines	13.9	Drifting longlines	38
	3.	Pole and line	0.1	/	/	/	/
Nazca-Desventuradas Marine Park	1.	Drifting longlines	99.2	Drifting longlines	99.6	Drifting longlines	99.5
	2.	Trawling	0.8	Squid jigger	0.3	Squid jigger	0.5
	3.	/	/	Trawling	0.1	/	/
Revillagigedo National Park	1.	Tuna purse seines	59.3	Tuna purse seines	82.3	Tuna purse seines	99.9
	2.	Pole and line	40.7	Pole and line	17.7	Pole and line	0.1
	3.	/	/	/	/	/	/

The three most prevalent fishing methods in each VLMPA during each of the three periods characterizing the life cycle of an MPA: before announcement, after announcement & before implementation and after implementation.Estimation of the type of fishing effort is provided by GFW.

(spillover effect from implementation), the effect of the announcement of the VLMPA on fishing effort both inside the VLMPA, and the effect of the announcement in the donut immediately surrounding the VLMPA. We estimate the following linear regression model:

$$\log(y_{it}) = \alpha_i + \alpha_y + \alpha_{U_i m} + \beta_V^A V_{it}^A + \beta_V^I V_{it}^I + \beta_D^A D_{it}^A + \beta_D^I D_{it}^I + \mathbf{X}_{it}\gamma + \eta_{it}. \tag{3.1}$$

In Equation (3.1), $\log(y_{it})$ represents the natural log of average daily fishing hours per 1,000 km^2 occurring in area i during period t. α_i is a fixed effect for area i. α_y represents annual fixed effects for year y. $\alpha_{U_i m}$ represents joint fixed effects for each VLMPA-donut pair U_i and each month m (across years), accounting for seasonal variation in a given VLMPA and its surrounding donut. V_{it}^A is an indicator for VLMPAs during the announcement period, which is 1 if i is a VLMPA and t is during the announcement period for area i. V_{it}^I is an indicator for VLMPAs after implementation, which is 1 if i is a VLMPA and t after implementation of VLMPA i (this value takes on a fraction during the month of implementation, described in the next paragraph). D_{it}^A is an indicator for the donuts around VLMPAs during the announcement period, which is 1 if i is a VLMPA's donut and t is during the announcement period for the VLMPA associated with i (this value also takes on a fraction during the month of announcement, and the month of implementation, described in the next paragraph). V_{it}^I is an indicator for the donuts around VLMPAs after implementation, which is 1 if i is a VLMPA's donut and t is after implementation of the VLMPA associated with i. $\mathbf{X}_i t$ is a vector of control variables, for which we consider the log of average wind speed and the log of average net primary productivity (NPP). η_{it} is an error term.

Before running the regression, we aggregate the data to monthly averages. That is, period t represents a year-month. Since we know the specific dates of each VLMPA's announcement and implementation, we allow the indicators V_{it}^A, V_{it}^I, D_{it}^A, and D_{it}^I to take on values between 0 and 1 for the month of announcement or implementation based on the fraction of days of the month that were within the respective announcement period or post-implementation period. For example, if a VLMPA was implemented on November 10th of a given year (a 30 day month), then V_{it}^A and D_{jt}^A would be $9/30 = .3$, and V_{it}^I and D_{jt}^I would be $21/30 = .7$ for that VLMPA i and its respective donut j during that month (assuming the announcement did not also take place during that same month).

β_V^A, β_V^I, β_D^A, and β_D^I each represent one of four treatment effects. β_V^A is the announcement period effect for VLMPAs. β_V^I is the implementation effect for VLMPAs. β_D^A is the spillover effect to the waters surrounding VLMPAs during the announcement period. β_D^I is the spillover effect to the waters surrounding VLMPAs after implementation. We consider β_V^I to be the most important treatment effect that we estimate, as this measures the reduction in fishing effort within the VLMPA itself resulting from implementation of the VLMPA.

To control for changes in global fishing trends, we use as a control group fishing effort in the rest of the world, beyond the VLMPAs and their surrounding donuts. In order to minimize the possible impact of spillovers from the VLMPAs into the "rest of world" control, we exclude all fishing within a buffer of 100 km from each VLMPA from the rest of the world data. The largest donuts that we consider for spillover effects stretch out to 50 km from the VLMPAs, such that all donuts fall within the 100 km buffer from the "rest of world" area. Since we only have one control observation and eight each of VLMPAs and their respective donuts, we weight the control observation by eight, such that control, VLMPA, and donut observations are balanced. Note that the rest of world control is given its own monthly fixed effects $\alpha_{U_i m}$.[6]

There are some 0's in the data for average daily fishing hours by month in some areas and months. In order to take logs of y_{it}, we add a small value to y_{it}. We choose one-one millionth, or 0.000001, as this is a small value two orders of magnitude below the smallest non-zero value of y_{it} observed in the data.

Table 3.3 presents estimates for the regression in Equation (3.1), with standard errors clustered by VLMPA-donut pairs. Each column represents a regression using a different size for the donut surrounding the VLMPA: 10 km, 30 km, and 50 km. The main coefficient representing the treatment effect of implementing the VLMPA, listed as VLMPA Implementation, is large and significant across specifications. This table provides statistical evidence that the creation of the VLMPAs has reduced fishing effort on average.

Notably, however, we are unable to detect a significant positive announcement effect or spillover effect in our sample. The sign on VLMPA Announcement is negative across specifications, with marginal significance in some of the specifications, suggesting a possible reduction in average fishing effort already during the announcement period. The sign is also negative for Donut Implementation, suggesting that spillovers may lead to a reduction of fishing effort nearby the VLMPA. However, we cannot firmly draw any conclusions regarding spillover effects resulting from the implementation VLMPAs due to the large standard errors on these coefficients, and so we do not rule out the possible existence of positive spillover effects from the creation of VLMPAs.

In the next section, we investigate the changes in fishing effort within the individual VLMPAs to better understand how fishing effort has been displaced within each individual VLMPA.

[6]The rest of the world data leaves out three VLMPAs that were implemented in 2018, and their respective 100 km buffers: Rapa Nui Rahui, Mar de Juan Fernández, and Coral Sea Australian Marine Park.

Table 3.3: Effect of VLMPA announcement and implementation on average daily fishing

	log(Fishing Hours)					
	(1)	(2)	(3)	(4)	(5)	(6)
	10 km	30 km	50 km	10 km	30 km	50 km
VLMPA Announcement	−1.478	−1.449	−1.503*	−1.477*	−1.427*	−1.507*
	(0.801)	(0.789)	(0.771)	(0.775)	(0.760)	(0.752)
VLMPA Implementation	−2.353*	−2.320**	−2.367*	−2.325**	−2.240**	−2.363**
	(1.028)	(0.999)	(1.029)	(0.968)	(0.917)	(0.962)
Donut Announcement	−0.611	0.021	0.117	−0.618	0.036	0.107
	(1.206)	(1.057)	(0.904)	(1.189)	(1.049)	(0.897)
Donut Implementation	−1.063	−1.058	−1.134	−1.039	−0.978	−1.132
	(0.858)	(1.011)	(0.926)	(0.834)	(1.006)	(0.942)
log(Wind Speed)	N	N	N	Y	Y	Y
log(NPP)	N	N	N	Y	Y	Y
Observations	1,020	1,020	1,020	1,020	1,020	1,020
R^2	0.772	0.779	0.778	0.772	0.779	0.778
Adjusted R^2	0.740	0.748	0.747	0.740	0.748	0.747

Note: *p<0.1; **p<0.05; ***p<0.01

Columns (1) and (4), (2) and (5), and (3) and (6) represent the inclusion of donuts extending 10 km, 30 km, and 50 km from the VLMPA, respectively. Columns (1) – (3) do not include any control variables, while Columns (4) – (6) include controls for the log of average wind speed and the log of average net primary productivity (NPP). VLMPA Announcement is an indicator for being in a VLMPA during the announcement period, V_{it}^A. VLMPA Implementation is an indicator for being in a VLMPA during after implementation, V_{it}^I. Donut Implementation is an indicator for being in the donut of a VLMPA during the announcement period, D_{it}^A. Donut Implementation is an indicator for being in the donut of a VLMPA after implementation, D_{it}^I. Standard errors are clustered by VLMPA-donut pairs.

3.4 Investigating Displacement of Fishing Effort

In order to analyze the deterrence effect of the VLMPAs in our sample, we investigate the change between annualized fishing effort inside the area prior to the announcement and after the implementation of the fishing ban. For this endeavor, we rely on the Incursion Ratio (IR), which we define as the ratio of the distance between the location of fishing effort and the closest point on the VLMPA border, over the distance to the border for the deepest location inside the VLMPA, i.e.

$$IR = \frac{\text{distance of fishing effort to border}}{\text{distance of deepest point to border}} \tag{3.2}$$

For example, the most isolated point within the Revillagigedo National Park is located 130.7 km away from the nearest boundary point. On January 31, 2018, a vessel was identified to be illegally fishing 23.8 km away from the nearest point on the border. It would consequently correspond to $IR = 0.18$. We represent the deepest point inside each VLMPA in Figure 3.2 (note that this point may not be unique). The evolution of the distribution of fishing effort inside each of the eight VLMPAs is shown in Figure 3.5, along with the distribution shift for all areas combined (last panel). We display the counts of fishing hours per incursion ratio to convey the change in magnitude along with the level of displacement. The horizontal axis in Figure 3.5 displays the Incursion Ratio. The two vertical green lines show the limits of the range of values IR could take. Whenever $IR = 0$, this would imply that a vessel if fishing right on the border of the VLMPA. Whenever $IR = 1$, the vessel would be fishing at the deepest possible location inside of the VLMPA.

At first glance, it would appear that there is also heterogeneity between the eight VLMPAs when it comes to the displacement of fishing effort during the period following the implementation of the fishing ban relative to the period preceding the announcement that the area would eventually become a marine reserve. The entire distribution shifted towards the border for Phoenix Islands Protected Area and for Revillagigedo National Park. The spread of the distributions seems to have remained virtually unchanged for the remaining areas. When it comes to the overall magnitude, it would appear that the overall count of fishing hours dropped in seven out of eight areas. Pacific Remote Islands Marine National Monument appears to be the only VLMPA witnessing a slight increase in the count of post-implementation annualized fishing hours relative to the pre-announcement period. As was the case for Figure 3.3, Figure 3.5 would indicate that, at the aggregate level, fishing effort decreased after the ban, and would also indicate a general displacement of fishing effort towards the border of the VLMPA.

In Figure 3.5 and for the rest of our analysis, we exclude the period of time between the announcement and the implementation of the VLMPAs. Provided that fishing effort during

Figure 3.5: Displacement of fishing effort

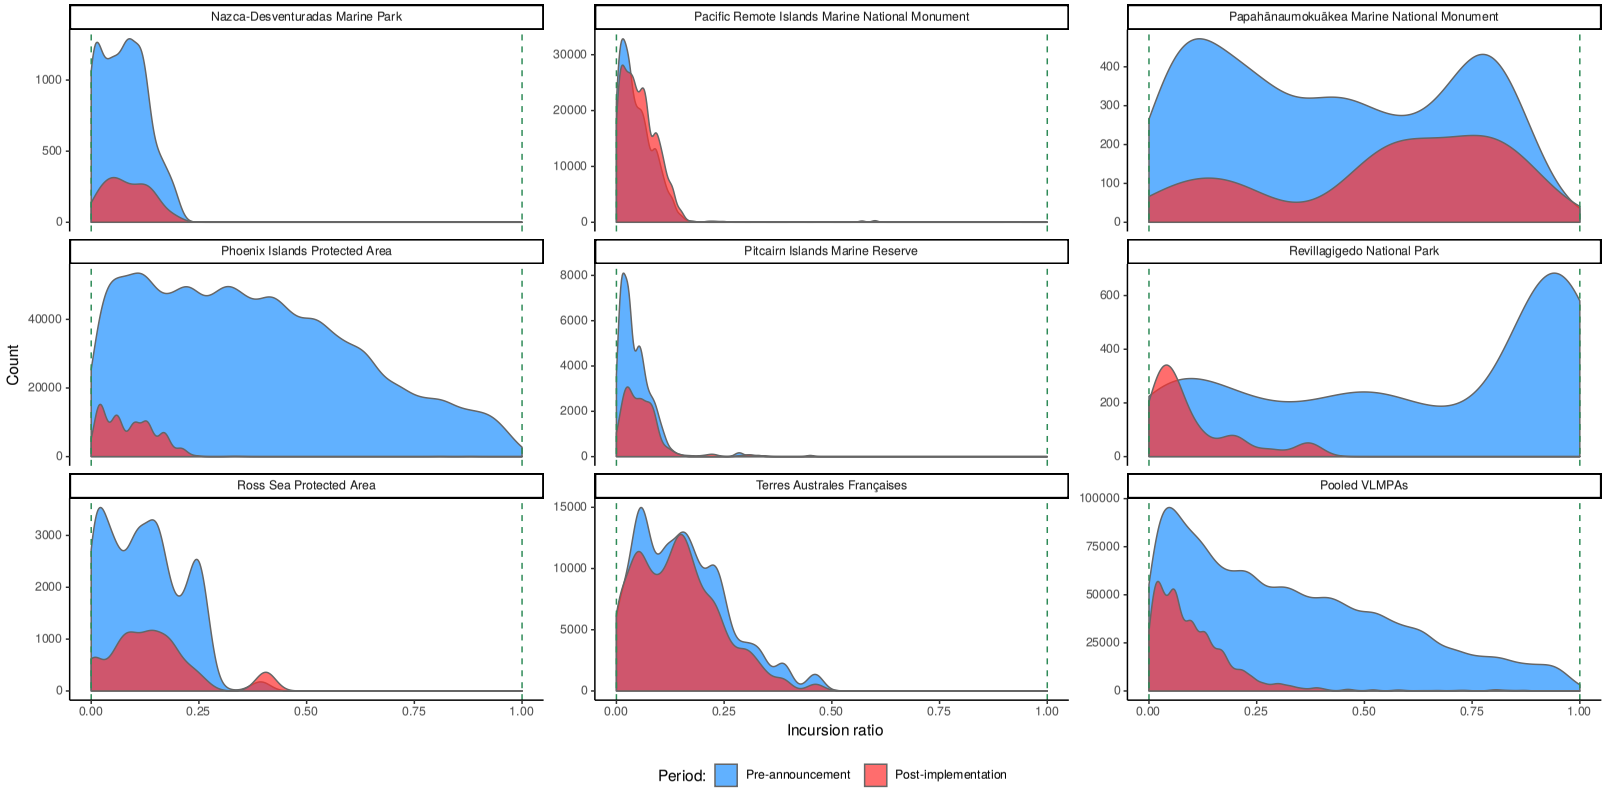

Comparison of the annualized count of fishing effort inside the VLMPAs prior to the announcement of their creation and after their implementation. Counts of fishing hours are restricted to the vessels tracked by the GFW database prior to the announcement date that a restriction on fishing effort may eventually take place. A version of the density plot that includes all the vessels is shown in Figure B.1.

the announcement period does not cause stocks to fall below the minimum viable stock threshold, a permanent ban that successfully deters fishing effort will eventually allow the resource to recover inside the marine reserve. By excluding this period, we therefore see how fishing effort has decreased relative to the time when vessels fished in an area not scheduled to be protected. We additionally focus solely on the evolution of fishing effort inside the areas for vessels tracked by GFW prior to the announcement of the VLMPAs.[7].

In order to measure the displacement of fishing effort within each VLMPA in our sample, we estimate the following fourth-order polynomial equation:

$$d_i = \sum_{k=0}^{4} \beta_k \cdot IR_i^k + \varepsilon_i, \tag{3.3}$$

where d_i corresponds to the difference in the number of annualized fishing hours after the implementation of the fishing ban and the number of annualized fishing hours preceding the announcement that a marine reserve would eventually be created. IR_i corresponds to the incursion ratio, rounded to the second decimal, and ε_i corresponds to heteroskedasticity-robust disturbances. The fitted regression curves and their associated 95% confidence intervals are shown in Figure 3.6.

After estimating Equation (3.3), we calculate the area between the curve and the horizontal axis. A positive (negative) number would imply that fishing effort has increased (decreased) overall after the implementation of the ban, relative to the period preceding the announcement of the possible creation of a marine reserve. The area calculations are displayed in Table 3.4.[8]

When pooling all the eight VLMPAs together using this method, we would conclude that "no-take" VLMPAs have overall been successful at deterring fishing effort since the area between the fitted fourth-order polynomial and 0 (represented by the horizontal red-dashed line) would be negative (-86.9) and significant at a 5% significance level with an associated 95% confidence interval of $[-120.1; -53.7]$. Should we only value the aggregate effectiveness of VLMPAs to deter fishing effort, then this is good news. Unfortunately, a closer look at each of the VLMPAs projects a more nuanced picture. As shown in Table 3.4, seven out of the eight VLMPAs are associated with an overall reduction in fishing effort after the ban. The Pacific

[7]Restricting the number of vessels practicing banned fishing to those tracked prior to the announcement of an eventual ban has a heterogeneous impact across countries. China, Taiwan, the United States, Japan and South Korea are the most affected by the restriction while most of the other countries have their fleet virtually unaffected by the restriction

[8]In Table B.3, we display the area calculations corresponding to third-order polynomial and fifth-order polynomial curves, along with their corresponding heteroskedasticity-robust 95% confidence intervals.

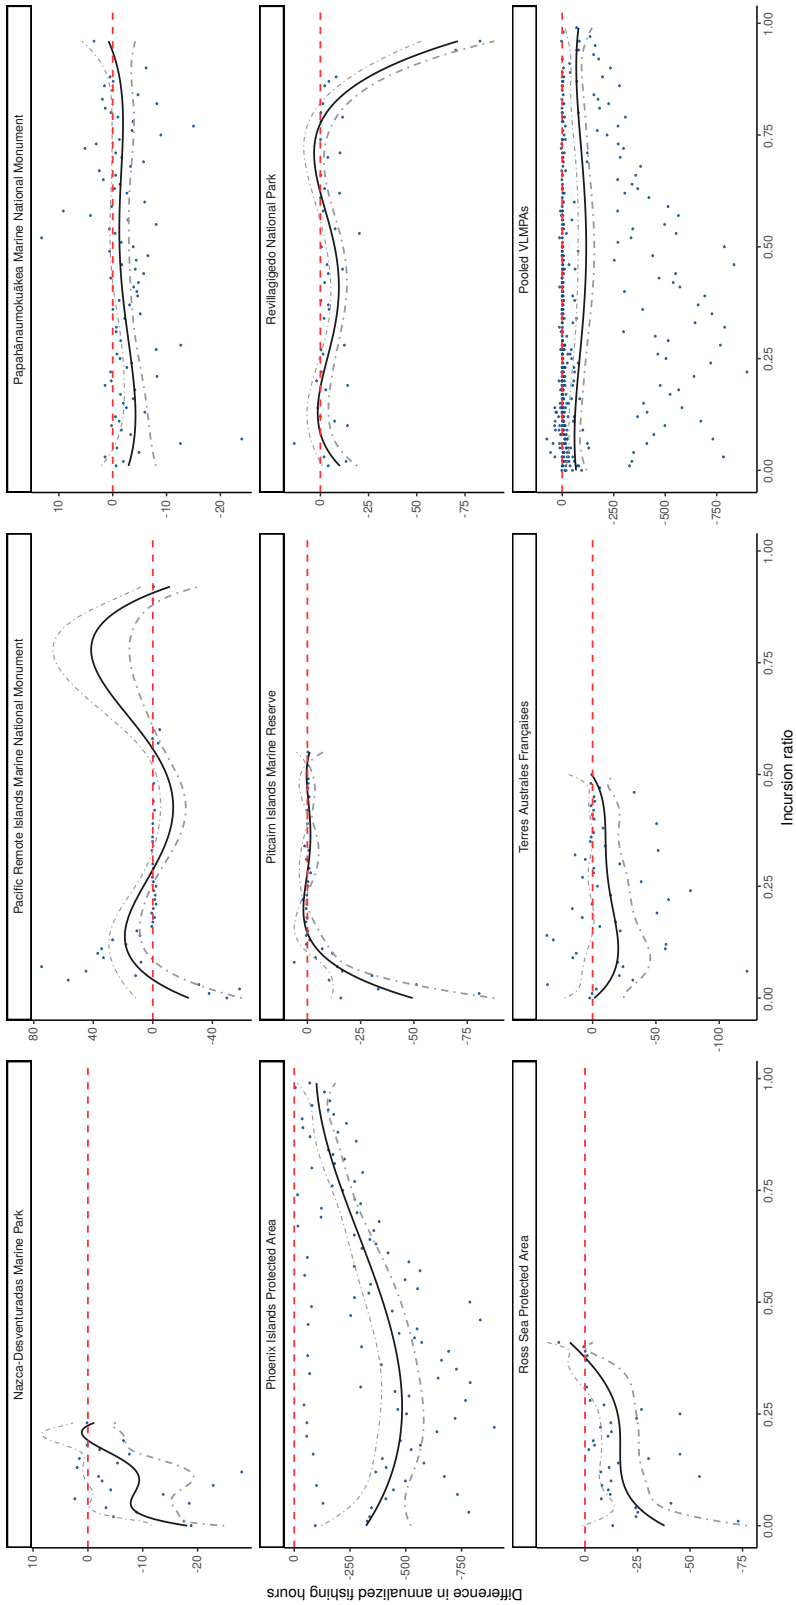

Figure 3.6: Quartic polynomial curves of best fit and their associated 95% confidence intervals

The scatter points represent the difference between the annualized fishing hours at a given incursion ratio after the implementation of the ban and prior to the announcement of the creation of the VLMPA. Negative values signify that there were fewer fishing hours at a given depth inside the VLMPA after the ban was put in place than during the period preceding the announcement of a potential marine reserve.

Remote Islands Marine National Monument, an VLMPA managed by the United States, is associated with an overall increase in annualized fishing effort during the period following the creation following the implementation of the ban, relative to the period preceding the announcement that a ban would eventually be implemented. Four of the VLMPAs – the Ross Sea Protected Area, Papahānaumokuākea Marine National Monument, Revillagigedo National Park and Phoenix Islands Protected Area – are associated with reductions in fishing effort that are significant at the 5% level. The largest reduction in fishing effort occurs in Phoenix Islands Protected Area, a VLMPA managed by the Republic of Kiribati. Kiribati is a country listed among the Small Island Developing States (SIDS).

Table 3.4: Area under the 4^{th}-order polynomial curve estimating the difference in annualized fishing hours

VLMPA name	Lower bound	Estimate	Upper bound
Ross Sea Protected Area	-10.6	-6.1	-1.6
Papahānaumokuākea Marine National Monument	-4.1	-2.1	-0.1
Pacific Remote Islands Marine National Monument	-3.4	9	21.3
Pitcairn Islands Marine Reserve	-5.4	-2.5	0.4
Terres Australes Françaises	-14.3	-6.2	1.8
Phoenix Islands Protected Area	-407.9	-332.7	-257.4
Nazca-Desventuradas Marine Park	-3.1	-1.5	0.2
Revillagigedo National Park	-12.6	-7.1	-1.6
Pooled VLMPAs	-120.1	-86.9	-53.7

A positive (negative) number means that annualized banned fishing effort exceeds (falls below) annualized fishing effort during the period preceding the announcement of a possible VLMPA creation.

Figure 3.5 shows that with the exception of the Papahānaumokuākea Marine National monument, a VLMPA managed by the United States, the VLMPAs where fishing took place deep within the to-be-protected area prior to the announcement of the VLMPA witnessed a significant shift of the distribution of fishing effort by the vessels tracked prior to the announcement towards the border of the respective VLMPA (Phoenix Islands Protected Area, Revillagigedo National Park). We fail to observe similar shifts in the spatial distribution of fishing effort for VLMPAs where the bulk of pre-announcement fishing effort took place at an incursion ratio smaller than halfway to the deepest point inside of the VLMPA.

3.5 Analyzing the Profile of Infringing Vessels

As we have at our disposal a plethora of information concerning the vessels responsible for fishing effort inside these areas, we are able to investigate the characteristics of ships engaging in illegal fishing effort. Indeed, thanks to the MMSI uniquely identifying each fishing vessel, GFW is able to track each vessel's length, tonnage, engine size, class, and its flag. Table 3.5 ranks the countries of origin of all the vessels associated with banned fishing by the number of infractions committed by vessels carrying their flag.

Table 3.5 shows that the five countries most responsible for banned fishing are associated with 80.7% of all the infractions observed in the eight VLMPAs, while their trespassing fleet represents 72.3% of all infringing vessels.

We also note that the average size of the trespassing vessels stands at 58.4 meters long, which corresponds to about 2.3 times the size of the average vessel (25.7 meters long) tracked by GFW over the 2012–2018 period. This might be explained by the remoteness of the VLMPAs under analysis, which would require larger vessels to make the journey and allow for a catch large enough to make the fishing excursion profitable.

3.6 Discussion

We leverage a database tracking fishing effort at an unprecedented resolution to evaluate the effectiveness of "no-take" very large marine protected areas (areas spanning more than 100,000 km^2) to deter fishing effort, measured in hours spent fishing inside the protected areas. We find evidence that, on average, VLMPAs have been successful at deterring fishing effort. This finding is encouraging. However, a more granular approach reveals some heterogeneity in the reduction in fishing effort across the eight VLMPAs under consideration, which might be problematic if each area was designed to preserve a distinct ecosystem. We find find that the area associated with the largest reduction in fishing effort (Phoenix Islands Protected Area) is managed by the Republic of Kiribati, a Small Island Developing States, while the least successful one (Pacific Remote Islands Marine National Monument) is managed by the United States.

The wealth of information on the vessels tracked by GFW allows us to identify the characteristics of the infringing vessels. We consequently analyze the profile of fishing vessels practicing banned fishing in the VLMPAs and find that they can be traced back to just 24 countries, and that the five countries most responsible for banned fishing are associated with more than 80% of all the infractions observed in the eight VLMPAs. We also find that trespassing vessels are on average more than twice as large as the average vessel in our sample.

Table 3.5: Ranking of the countries associated with prohibited fishing inside of the established VLMPAs

Rank	Country of Origin	All VLMPAs		
		Infractions	Vessels	Mean Length
1.	South Korea	1,527	112	58
2.	Taiwan	1,264	122	38
3.	China	872	114	50
4.	United States	613	84	42
5.	France	559	8	58
6.	Japan	376	52	50
7.	Kiribati	254	12	68
8.	Vanuatu	104	23	51
9.	Spain	89	5	64
10.	Papua New Guinea	70	17	75
11.	Micronesia (Federated States)	38	12	73
12.	Marshall Islands	33	7	71
13.	Honduras	31	1	54
14.	Mexico	31	15	72
15.	Solomon Islands	26	3	68
16.	Slovakia	24	2	29
17.	New Zealand	23	4	50
18.	Maldives	13	2	55
19.	Russia	6	1	56
20.	Australia	5	1	68
21.	Ukraine	4	1	55
22.	Cook Islands	3	1	57
23.	El Salvador	3	2	83
24.	Kosovo (UNK)	1	1	72
	Unidentified	21	6	44
	Total	5,990	608	58.4

We define infractions at the VLMPA-daily-vessel level. If a vessel were to fish illegally two days in a row in a given VLMPA, each daily occurrence would be counted as a distinct infraction. If a vessel were to fish in two VLMPAs during the same day, these would also get counted as two separate infractions.

The results in this paper provide insights into the effectiveness of VLMPAs implemented during the last decade. Increasing our understanding of VLMPAs and their ability to deter fishing effort will be important to policy makers moving forward. In 2021, Chile, Costa Rica, France, the United Kingdom, and the United States announced the development of a partnership to advance the use of MPAs as a tool to fight climate change (National Oceanic and Atmospheric Administration, 2021). At the UN Ocean Conference in 2022, a target of protecting 30% of the global ocean by MPAs by 2030 received the voluntary commitments of over 150 countries, with the Protecting Our Planet Challenge offering to invest at least $1 billion USD into creating, expanding, and managing MPAs (UN News Africa Renewal, 2022). Learning from the implementation of VLMPAs and their past successes and failures will aid policymakers in better designing and enforcing marine protection policies in the future.

Bibliography

[1] Ahlfeldt, Gabriel M. and Daniel P. McMillen. "Tall Buildings and Land Values: Height and Construction Cost Elasticities in Chicago, 1870–2010". In: *The Review of Economics and Statistics* 100 (5 2018), pp. 753 –768.

[2] An, Mihyang, Stephen M. Colarelli, Kimberly O'Brien, and Melanie E. Boyajian. "Why We Need More Nature at Work: Effects of Natural Elements and Sunlight on Employee Mental Health and Work Attitudes". In: *PLoS ONE* 11.5 (2016), pp. 1 –17.

[3] Anderson, Michael L. "As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality". In: *Journal of the European Economic Association* 18 (4 2020), pp. 1886 –1927.

[4] Arias, Adrian, Joshua E Cinner, Rhondda E Jones, and Robert L Pressey. "Levels and drivers of fishers' compliance with marine protected areas". In: *Ecology and Society* 20.4 (2015), pp. 1–14.

[5] Auffhammer, Maximilian. "Cooling China: The Weather Dependence of Air Conditioner Adoption". In: *Frontiers of Economics in China* 9.1 (2014), pp. 70 –84.

[6] Ban, Natalie C, Tammy E Davies, Stacy E Aguilera, Cassandra Brooks, Michael Cox, Graham Epstein, Louisa S Evans, Sara M Maxwell, and Mateja Nenadovic. "Social and ecological effectiveness of large marine protected areas". In: *Global Environmental Change* 43 (2017), pp. 82–91.

[7] Barr, Jason M. *Skyscrapers and Cities: A Q&A Interview with Edward Glaeser (Part I)*. Ed. by buildingtheskyline.org. `https://buildingtheskyline.org/glaeser-interview-1/`. Accessed: 2022-07-29. 2018.

[8] Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro. "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century". In: *Journal of Political Economy* 124.1 (2016), pp. 105 –159.

[9] BBC News. *Budget 2015: Pitcairn Islands get huge marine reserve*. `https://www.bbc.com/news/science-environment-31943633`. [Online; accessed 2019-04-09]. Mar. 2015.

[10] Beaulieu, Catherine and Marie Dumont. "Light exposure in the natural environment: Relevance to mood and sleep disorders". In: *Sleep Medicine* 8 (2007), pp. 557 –565.

[11] Beine, Michel and Christopher Parsons. "Climatic Factors as Determinants of International Migration". In: *The Scandinavian Journal of Economics* 117.2 (2015), pp. 723 –767.

[12] Bergseth, Brock J, Georgina G Gurney, Michele L Barnes, Adrian Arias, and Joshua E Cinner. "Addressing poaching in marine protected areas through voluntary surveillance and enforcement". In: *Nature Sustainability* 1.8 (2018), p. 421.

[13] Bergseth, Brock J, Garry R Russ, and Joshua E Cinner. "Measuring and monitoring compliance in no-take marine reserves". In: *Fish and Fisheries* 16.2 (2015), pp. 240–258.

[14] Berry, William D, Richard C Fording, and Russell L Hanson. "An annual cost of living index for the American states, 1960-1995". In: *Journal of Politics* 62.2 (2000), pp. 550–567.

[15] Biardeau, Léopold T, Lucas W Davis, Paul Gertler, and Catherine Wolfram. "Heat exposure and global air conditioning". In: *Nature Sustainability* 3.1 (2020), pp. 25–28.

[16] Biddle, Jeff. "Air Conditioning, Migration, and Climate Related Wage and Rent Differentials". In: *Research in Economic History* 28 (2012), pp. 1 –41.

[17] Biddle, Jeff. "Explaining the spread of residential air conditioning, 1955–1980". In: *Explorations in Economic History* 45 (2008), pp. 402 –423.

[18] Biggs, Caroline. *Rooms With a View (and How Much You'll Pay for Them)*. Ed. by New York Times. `https://www.nytimes.com/2019/05/10/realestate/rooms-with-a-view-and-how-much-youll-pay-for-them.html`. Accessed: 2022-12-10. 2019.

[19] Bliss, Laura. *How the Battle for Sunlight Shaped New York City*. Ed. by bloomberg.com. `https://www.bloomberg.com/news/articles/2016-12-18/new-york-city-zoning-and-the-fight-for-sunlight`. Accessed: 2022-07-29. 2016.

[20] Bohra-Mishra, Pratikshya, Michael Oppenheimer, and Solomon M. Hsiang. "Nonlinear Permanent Migration Response to Climatic Variations but Minimal Response to Disasters". In: *PNAS* 111.27 (2014), pp. 9780 –9785.

[21] Boonzaier, Lisa and Daniel Pauly. "Marine protection targets: an updated assessment of global progress". In: *Oryx* 50.1 (2016), pp. 27–35.

[22] Bradley, Darcy, Juan Mayorga, Douglas J McCauley, Reniel B Cabral, Patric Douglas, and Steven D Gaines. "Leveraging satellite technology to create true shark sanctuaries". In: *Conservation Letters* 12.2 (2019), e12610.

[23] Chau, Kwong-Wing, S. K. Wong, Y. Yau, and A. K. C. Yeung. "Determining Optimal Building Height". In: *Urban Studies* 44.3 (2007), pp. 591 –607.

[24] Chen, Shuai, Paulina Oliva, and Peng Zhang. "The effect of air pollution on migration: Evidence from China". In: *Journal of Development Economics* 156 (2022).

[25] Claudet, Joachim, Craig W Osenberg, Lisandro Benedetti-Cecchi, Paolo Domenici, José-Antonio García-Charton, Ángel Pérez-Ruzafa, Fabio Badalamenti, Just Bayle-Sempere, Alberto Brito, Fabio Bulleri, et al. "Marine reserves: size and age do matter". In: *Ecology letters* 11.5 (2008), pp. 481–489.

[26] Convention on Biological Diversity (CBD). *Strategic Plan for Biodiversity 2011-2020 and the Aichi biodiversity targets.* `https://www.cbd.int/decision/cop/?id=12268`. 2010.

[27] Costello, Mark J and Bill Ballantine. "Biodiversity conservation should focus on no-take Marine Reserves: 94% of Marine Protected Areas allow fishing". In: *Trends in ecology & evolution* 30.9 (2015), pp. 507–509.

[28] Couture, Victor. "Valuing the Consumption Benefits of Urban Density". In: (2016). Working Paper.

[29] Cressey, Daniel. *Ocean conservation: Uncertain sanctuary.* `https://www.nature.com/news/ocean-conservation-uncertain-sanctuary-1.9568`. Dec. 2011.

[30] Deschênes, Olivier and Enrico Moretti. "Extreme Weather Events, Mortality, and Migration". In: *The Review of Economics and Statistics* 91.4 (2009), pp. 659 –681.

[31] Donovan, Geoffrey H. and David T. Butry. "The value of shade: Estimating the effect of urban trees on summertime electricity use". In: *Energy and Buildings* 41 (2009), pp. 662 –668.

[32] Dougherty, Conor. *The Great American Single-Family Home Problem.* Ed. by New York Times. `https://www.nytimes.com/2017/12/01/business/economy/single-family-home.html`. Accessed: 2022-07-29. 2017.

[33] Dulvy, Nicholas K. "Super-sized MPAs and the marginalization of species conservation". In: *Aquatic Conservation: Marine and Freshwater Ecosystems* 23.3 (2013), pp. 357–362.

[34] Dureuil, Manuel, Kristina Boerder, Kirsti A Burnett, Rainer Froese, and Boris Worm. "Elevated trawling inside protected areas undermines conservation outcomes in a global fishing hot spot". In: *Science* 362.6421 (2018), pp. 1403–1407.

[35] Edgar, Graham J, Rick D Stuart-Smith, Trevor J Willis, Stuart Kininmonth, Susan C Baker, Stuart Banks, Neville S Barrett, Mikel A Becerro, Anthony TF Bernard, Just Berkhout, et al. "Global conservation outcomes depend on marine protected areas with five key features". In: *Nature* 506.7487 (2014), p. 216.

[36] Englander, Gabriel. "Property rights and the protection of global marine resources". In: *Nature Sustainability* 2.10 (2019), pp. 981–987.

[37] Flanders Marine Institute. *Maritime Boundaries and Exclusive Economic Zones (200 NM), version 10.* `http://www.marineregions.org/`.

[38] Fleming, David, Arthur Grimes, Laurent Lebreton, David Maré, and Peter Nunns. "Valuing Sunshine". In: *Regional Science and Urban Economics* 68 (2018), pp. 268 –276.

[39] Glaeser, Edward L., Joseph Gyourko, and Raven Saks. "Why is Manhattan So Expensive? Regulation and the Rise in Housing Prices". In: *Journal of Law and Economics* 48 (2005), pp. 331 –369.

[40] Glaeser, Edward L., Scott Duke Kominers, Michael Luca, and Nikhil Naik. "Big Data and Big Cities: The Promises and Limitations of Improved Measures for Urban Life". In: *Econoic Inquiry* 56 (1 2018), pp. 114 –137.

[41] Gruby, Rebecca L, Noella J Gray, Lisa M Campbell, and Leslie Acton. "Toward a social science research agenda for large marine protected areas". In: *Conservation Letters* 9.3 (2016), pp. 153–163.

[42] Guiet, Jérôme, Eric Galbraith, David Kroodsma, and Boris Worm. "Seasonal variability in global industrial fishing effort". In: *PloS one* 14.5 (2019), e0216819.

[43] Hickey, Walt. *The Tall Apple: Number of Stories in NYC Skyscrapers Might Floor You*. Ed. by FiveThirtyEight. `https://fivethirtyeight.com/features/the-tall-apple-number-of-stories-in-nyc-skyscrapers-might-floor-you/`. Accessed: 2022-12-11. 2014.

[44] Holick, Michael F. "Sunlight and vitamin D for bone health and prevention of autoimmune diseases, cancers, and cardiovascular disease". In: *The American Journal of Clinical Nutrition* 80.6 (2004), 1678S –1688S.

[45] Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. "Estimating Vector Autoregressions with Panel Data". In: *Econometrica* 56.6 (1988), pp. 1371 –1395.

[46] Hsiang, Solomon M and Amir S Jina. *The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones*. Tech. rep. National Bureau of Economic Research, 2014.

[47] International Union for Conservation of Nature (IUCN). *17th session of the general assembly of IUCN AND 17th technical meeting*. `https://portals.iucn.org/library/efiles/documents/GA-17th-011.pdf`. 1988.

[48] Jone, PJS, Wanfei Qiu, and EM De Santo. *Governing marine protected areas: getting the balance right*. United Nations Environment Programme (UNEP), 2011.

[49] Kim, Elizabeth. *Judge Voids City Permit For Controversial Upper West Side Luxury Mega Tower*. Ed. by gothamist.com. `https://gothamist.com/news/judge-voids-city-permit-controversial-upper-west-side-luxury-mega-tower`. Accessed: 2022-07-29. 2020.

[50] Kroodsma, David A, Juan Mayorga, Timothy Hochberg, Nathan A Miller, Kristina Boerder, Francesco Ferretti, Alex Wilson, Bjorn Bergman, Timothy D White, Barbara A Block, et al. "Tracking the global footprint of fisheries". In: *Science* 359.6378 (2018), pp. 904–908.

[51] Lambert, G. W., C. Reid, D. M. Kaye, G. L. Jennings, and M. D. Esler. "Effect of sunlight and season on serotonin turnover in the brain". In: *Lancet* 360 (2002), pp. 1840 –1842.

[52] Lester, Sarah E, Benjamin S Halpern, Kirsten Grorud-Colvert, Jane Lubchenco, Benjamin I Ruttenberg, Steven D Gaines, Satie Airamé, and Robert R Warner. "Biological effects within no-take marine reserves: a global synthesis". In: *Marine Ecology Progress Series* 384 (2009), pp. 33–46.

[53] Li, Xiaodi. "Do New Housing Units in Your Backyard Raise Your Rents?" In: *Journal of Economic Geography* 22 (6 2022), pp. 1309 –1352.

[54] Magris, Rafael A and Robert L Pressey. "Marine protected areas: Just for show?" In: *Science* 360.6390 (2018), pp. 723–724.

[55] McDermott, Grant R, Kyle C Meng, Gavin G McDonald, and Christopher J Costello. "The blue paradox: Preemptive overfishing in marine reserves". In: *Proceedings of the National Academy of Sciences* 116.12 (2019), pp. 5319–5325.

[56] MCI. *MPAtlas*. [Online; accessed 2019-04-08]. 2019. URL: `www.mpatlas.org`.

[57] Miranda, Fabio, Harish Doraiswamy, Marcos Lage, Luc Wilson, Mondrian Hsieh, and Cláudio T. Silva. "Valuing Sunshine". In: *IEEE Transactions on Visualization and Computer Graphics* 25.3 (2019), pp. 1559 –1574.

[58] National Geographic Society & Oceana. *Islas desventuradas – Biodiversidad marina y propuesta de conservación*. Report. Feb. 2013.

[59] National Oceanic and Atmospheric Administration. *New global partnership to elevate marine protected areas as tool in climate response*. `https://www.noaa.gov/news-release/new-global-partnership-to-elevate-marine-protected-areas-as-tool-in-climate-response`. 2021.

[60] New York State Senate. *Senate Bill S3820A*. `https://www.nysenate.gov/legislation/bills/2019/s3820`. Accessed: 2022-07-29. 2019.

[61] O'leary, Bethan C, Natalie C Ban, Miriam Fernandez, Alan M Friedlander, Pablo García-Borboroglu, Yimnang Golbuu, Paolo Guidetti, Jean M Harris, Julie P Hawkins, Tim Langlois, et al. "Addressing criticisms of large-scale marine protected areas". In: *Bioscience* 68.5 (2018), pp. 359–370.

[62] Pollnac, Richard, Patrick Christie, Joshua E Cinner, Tracey Dalton, Tim M Daw, Graham E Forrester, Nicholas AJ Graham, and Timothy R McClanahan. "Marine reserves as linked social–ecological systems". In: *Proceedings of the National Academy of Sciences* 107.43 (2010), pp. 18262–18265.

[63] Queiroz, Nuno, Nicolas E Humphries, Ana Couto, Marisa Vedor, Ivo Da Costa, Ana MM Sequeira, Gonzalo Mucientes, António M Santos, Francisco J Abascal, Debra L Abercrombie, et al. "Global spatial risk assessment of sharks under the footprint of fisheries". In: *Nature* 572.7770 (2019), pp. 461–466.

[64] Rappaport, Jordan. "Moving to Nice Weather". In: *Regional Science and Urban Economics* 37 (2007), pp. 375 –398.

[65] Sala, Enric and Sylvaine Giakoumi. "No-take marine reserves are the most effective protected areas in the ocean". In: *ICES Journal of Marine Science* 75.3 (2017), pp. 1166–1168.

[66] Science News. *Mega-Marine Reserve Going Off-Limits to Fishers*. `https://www.sciencemag.org/news/2014/05/mega-marine-reserve-going-limits-fishers`. [Online; accessed 2019-04-02]. May 2014.

[67] TAAF. *Bilan d'activités 2016 de la réserve naturelle des Terres australes françaises*. 2017.

[68] Toonen, Robert J, T'Aulani Wilhelm, Sara M Maxwell, Daniel Wagner, Brian W Bowen, Charles RC Sheppard, Sue M Taei, Tukabu Teroroko, Russell Moffitt, Carlos F Gaymer, et al. "One size does not fit all: the emerging frontier in large-scale marine conservation". In: *Marine pollution bulletin* 77.1-2 (2013), pp. 7–10.

[69] TRD Staff. *Extell's UWS tower nabs $967M construction loan*. Ed. by therealdeal.com. `https://therealdeal.com/2022/02/10/extells-uws-tower-nabs-967m-construction-loan/`. Accessed: 2022-07-29. 2022.

[70] UN Environment Programme WCMC and IUCN. *Protected Planet Report 2020*. `https://livereport.protectedplanet.net/`. 2020.

[71] UN News Africa Renewal. *UN Ocean Conference ends with call for greater ambition and global commitment to address dire state of the Ocean*. `https://www.un.org/africarenewal/magazine/july-2022/un-ocean-conference-ends-call-greater-ambition-and-global-commitment-address-dire`. 2022.

[72] UNEP-WCMC and IUCN. *Marine Protected Planet*. [Online; accessed 2019-04-08]. 2019. URL: `www.protectedplanet.net`.

[73] UNESCO. *World Heritage List*. `https://whc.unesco.org/en/list/1510`. [Online; accessed 2019-04-10]. July 2016.

[74] Weimerskirch, Henri, Julien Collet, Alexandre Corbeau, Adrien Pajot, Floran Hoarau, Cédric Marteau, Dominique Filippi, and Samantha C Patrick. "Ocean sentinel albatrosses locate illegal vessels and provide the first estimate of the extent of nondeclared fishing". In: *Proceedings of the National Academy of Sciences* 117.6 (2020), pp. 3006–3014.

[75] Wong, K. M. Grace. "Vertical cities as a solution for land scarcity: the tallest public housing development in Singapore". In: *Urban Design International* 9 (2004), pp. 17–30.

[76] Zupan, Mirta, Eliza Fragkopoulou, Joachim Claudet, Karim Erzini, Bárbara Horta e Costa, and Emanuel J Gonçalves. "Marine partially protected areas: drivers of ecological effectiveness". In: *Frontiers in Ecology and the Environment* 16.7 (2018), pp. 381–387.

Appendix A

Chapter 1 Appendix

A.1 Additional tables and figures

Additional tables

Table A.1: Effect of changes in shadow on changes in housing price with 500 meter fixed effects and cluster grid

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.331**	-0.334**	-0.292**
	(0.148)	(0.147)	(0.143)
Spatial-Time FE radius	500m	500m	500m
Cluster Grid	500m^2	500m^2	500m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	35,180	35,180	35,180
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. Column (1) estimates Equation (1.1). Column (2) estimates Equation (1.2). Column (3) estimates Equation (1.3). Spatial-Time Fixed Effects are measured with a 500 meter radius prior to differencing. Standard errors are clustered using a 500 square meter spatial grid.

Table A.2: Effect of changes in shadow relative to baseline shadow on changes in housing price with 500 meter fixed effects and cluster grid

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.259*	-0.273**	-0.237*
	(0.136)	(0.139)	(0.132)
$\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$	-0.111*	-0.094	-0.087
	(0.065)	(0.064)	(0.074)
Spatial-Time FE radius	500m	500m	500m
Cluster Grid	500m^2	500m^2	500m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	35,180	35,180	35,180
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. $\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$ represents the change in spatial-time demeaned shadow interacted with baseline shadow. Column (1) estimates Equation (1.5). Column (2) estimates Equation (1.5) with the inclusion of distance fixed effects as in Equation (1.2). Column (3) estimates Equation (1.5) with the inclusion of distance and height fixed effects as in Equation (1.3). Spatial-Time Fixed Effects are measured with a 500 meter radius prior to differencing. Standard errors are clustered using a 500 square meter spatial grid.

Table A.3: Effect of changes in shadow on changes in housing price with 500 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.303**	-0.301**	-0.249
	(0.153)	(0.150)	(0.152)
Spatial-Time FE radius	500m	500m	500m
Cluster Grid	500m^2	500m^2	500m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	29,257	29,257	29,257
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. Column (1) estimates Equation (1.1). Column (2) estimates Equation (1.2). Column (3) estimates Equation (1.3). Observations do not include units within 500 meters of Central Park. Spatial-Time Fixed Effects are measured with a 500 meter radius prior to differencing and prior to removing observations near Central Park. Standard errors are clustered using a 500 square meter spatial grid.

Table A.4: Effect of changes in shadow relative to baseline shadow on changes in housing price with 250 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.258*	-0.254*	-0.214
	(0.145)	(0.144)	(0.146)
$\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$	-0.169*	-0.160*	-0.155
	(0.095)	(0.093)	(0.095)
Spatial-Time FE radius	250m	250m	250m
Cluster Grid	250m^2	250m^2	250m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	29,257	29,257	29,257

Note: *p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. $\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$ represents the change in spatial-time demeaned shadow interacted with baseline shadow. Column (1) estimates Equation (1.5). Column (2) estimates Equation (1.5) with the inclusion of distance fixed effects as in Equation (1.2). Column (3) estimates Equation (1.5) with the inclusion of distance and height fixed effects as in Equation (1.3). Observations do not include units within 500 meters of Central Park. Spatial-Time Fixed Effects are measured with a 250 meter radius prior to differencing and prior to removing observations near Central Park. Standard errors are clustered using a 250 square meter spatial grid.

Table A.5: Effect of changes in shadow relative to baseline shadow on changes in housing price with 500 meter fixed effects and cluster grid, excluding units within 500 meters of Central Park

	$\Delta\widetilde{\log(p_{it})}$		
	(1)	(2)	(3)
$\Delta\widetilde{S_{it_j}}$	-0.246*	-0.258*	-0.207
	(0.146)	(0.147)	(0.144)
$\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$	-0.085	-0.065	-0.064
	(0.076)	(0.076)	(0.087)
Spatial-Time FE radius	500m	500m	500m
Cluster Grid	500m^2	500m^2	500m^2
Distance terms	No	Yes	Yes
Height terms	No	No	Yes
Observations	29,257	29,257	29,257
Note:			*p<0.1; **p<0.05; ***p<0.01

$\Delta\widetilde{\log(p_{it_j})}$ represents the change in spatial-time demeaned log housing price. $\Delta\widetilde{S_{it}}$ represents the change in spatial-time demeaned shadow. $\widetilde{(\Delta S_{it_j})} \times S_{it_{j-1}}$ represents the change in spatial-time demeaned shadow interacted with baseline shadow. Column (1) estimates Equation (1.5). Column (2) estimates Equation (1.5) with the inclusion of distance fixed effects as in Equation (1.2). Column (3) estimates Equation (1.5) with the inclusion of distance and height fixed effects as in Equation (1.3). Observations do not include units within 500 meters of Central Park. Spatial-Time Fixed Effects are measured with a 500 meter radius prior to differencing and prior to removing observations near Central Park. Standard errors are clustered using a 500 square meter spatial grid.

Additional figures

Figure A.1: Graph of downward-U shaped marginal profits

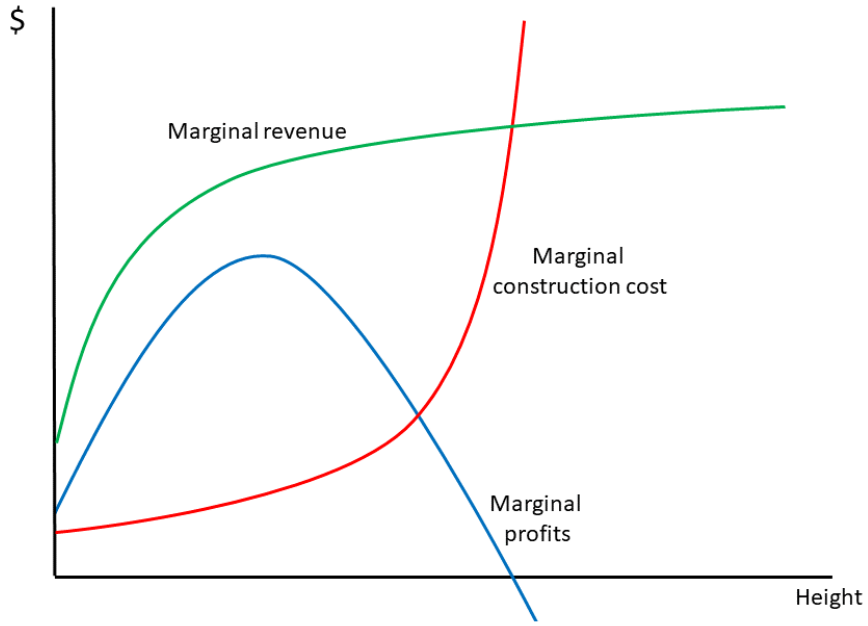

The red curve represents a convex marginal construction cost. The green curve represents a concave marginal revenue. The blue curve represents the marginal profit to the developer. Marginal profits is the difference between marginal revenue and marginal construction cost.

Figure A.2: Finding optimal height with downward-U shaped marginal profits

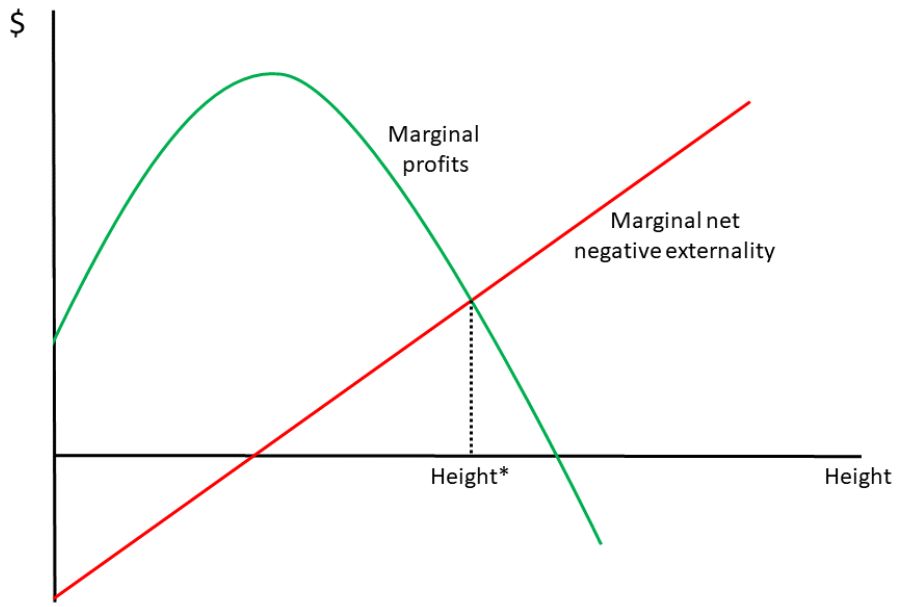

The red line represents the marginal net negative externality. The green curve represents the marginal profit to the developer. The optimal height, labelled Height*, can be found where the marginal net negative externality equals the marginal profits.

Appendix B

Chapter 3 Appendix

B.1 Additional tables and figures

Additional tables

Table B.1: Effect of VLMPA announcement and implementation on average daily fishing, all donut sizes, no controls

	log(Fishing Hours)				
	(1)	(2)	(3)	(4)	(5)
	10 km	20 km	30 km	40 km	50 km
VLMPA Announcement	−1.478	−1.419	−1.449	−1.439*	−1.503*
	(0.801)	(0.797)	(0.789)	(0.770)	(0.771)
VLMPA Implementation	−2.353*	−2.192*	−2.320**	−2.296*	−2.367*
	(1.028)	(1.031)	(0.999)	(1.022)	(1.029)
Donut Announcement	−0.611	−0.001	0.021	0.265	0.117
	(1.206)	(1.145)	(1.057)	(1.058)	(0.904)
Donut Implementation	−1.063	−0.796	−1.058	−0.882	−1.134
	(0.858)	(1.091)	(1.011)	(1.008)	(0.926)
log(Wind Speed)	N	N	N	N	N
log(NPP)	N	N	N	N	N
Observations	1,020	1,020	1,020	1,020	1,020
R^2	0.772	0.772	0.779	0.781	0.778
Adjusted R^2	0.740	0.741	0.748	0.751	0.747

Note: *p<0.1; **p<0.05; ***p<0.01

Columns (1) – (5) represent the inclusion of donuts extending 10 km, 20 km, 30 km, 40 km, and 50 km from the VLMPA, respectively. Control variables are not included in any of the columns. VLMPA Announcement is an indicator for being in a VLMPA during the announcement period, V_{it}^A. VLMPA Implementation is an indicator for being in a VLMPA during after implementation, V_{it}^I. Donut Implementation is an indicator for being in the donut of a VLMPA during the announcement period, D_{it}^A. Donut Implementation is an indicator for being in the donut of a VLMPA after implementation, D_{it}^I. Standard errors are clustered by VLMPA-donut pairs.

Table B.2: Effect of VLMPA announcement and implementation on average daily fishing, all donut sizes, with controls

	log(Fishing Hours)				
	(1)	(2)	(3)	(4)	(5)
	10 km	20 km	30 km	40 km	50 km
VLMPA Announcement	−1.477*	−1.405	−1.427*	−1.432*	−1.507*
	(0.775)	(0.764)	(0.760)	(0.744)	(0.752)
VLMPA Implementation	−2.325**	−2.123*	−2.240**	−2.256**	−2.363**
	(0.968)	(0.947)	(0.917)	(0.947)	(0.962)
Donut Announcement	−0.618	0.003	0.036	0.265	0.107
	(1.189)	(1.137)	(1.049)	(1.054)	(0.897)
Donut Implementation	−1.039	−0.728	−0.978	−0.843	−1.132
	(0.834)	(1.095)	(1.006)	(1.013)	(0.942)
log(Wind Speed)	Y	Y	Y	Y	Y
log(NPP)	Y	Y	Y	Y	Y
Observations	1,020	1,020	1,020	1,020	1,020
R^2	0.772	0.772	0.779	0.781	0.778
Adjusted R^2	0.740	0.740	0.748	0.750	0.747
Note:					*p<0.1; **p<0.05; ***p<0.01

Columns (1) – (5) represent the inclusion of donuts extending 10 km, 20 km, 30 km, 40 km, and 50 km from the VLMPA, respectively. All columns include controls for the log of average wind speed and the log of average net primary productivity (NPP). VLMPA Announcement is an indicator for being in a VLMPA during the announcement period, V_{it}^A. VLMPA Implementation is an indicator for being in a VLMPA during after implementation, V_{it}^I. Donut Implementation is an indicator for being in the donut of a VLMPA during the announcement period, D_{it}^A. Donut Implementation is an indicator for being in the donut of a VLMPA after implementation, D_{it}^I. Standard errors are clustered by VLMPA-donut pairs.

Table B.3: Areas under three distinct order polynomial curves estimating the difference in annualized fishing hours

VLMPA name	3rd-order polynomial			4th-order polynomial			5th-order polynomial		
	Lower bound	Estimate	Upper bound	Lower Bound	Estimate	Upper Bound	Lower bound	Estimate	Upper bound
Ross Sea Protected Area	-10.3	-6.2	-2.1	-10.6	-6.1	-1.6	-10.7	-6.1	-1.4
Papahānaumokuākea Marine National Monument	-4.0	-2.1	-0.2	-4.1	-2.1	-0.1	-4.1	-2.1	0.0
Pacific Remote Islands Marine National Monument	-10.1	-1.2	7.7	-3.4	9	21.3	-29.8	-8.7	12.3
Pitcairn Islands Marine Reserve	-5.8	-2.6	0.6	-5.4	-2.5	0.4	-5.6	-2.5	0.6
Terres Australes Françaises	-13.4	-6.2	0.9	-14.3	-6.2	1.8	-15.2	-6.3	2.6
Phoenix Islands Protected Area	-400.3	-332.7	-265.0	-407.9	-332.7	-257.4	-414.1	-332.7	-251.3
Nazca-Desventuradas Marine Park	-2.9	-1.6	-0.2	-3.1	-1.5	0.2	-3.5	-1.6	0.3
Revillagigedo National Park	-14.7	-7.5	-0.3	-12.6	-7.1	-1.6	-12.0	-7.2	-2.4
Pooled VLMPAs	-117.1	-86.6	-56.0	-120.1	-86.9	-53.7	-123.5	-87.0	-50.6

A positive (negative) number means that annualized banned fishing effort exceeds (falls below) annualized fishing effort during the period preceding the announcement of a possible VLMPA creation. Results remain stable across all specifications for all the VLMPAs but the Pacific Remote Islands Marine National Monument, which witnesses a reversal of sign for the main area estimate depending on whether the order of the polynomial is odd or even.

Additional figures

The inclusion of all the vessels to ever fish in the VLMPAs area (Figure B.1) reveals a general displacement of fishing effort towards the border, meaning that vessels which appeared later during the 2012–2018 period have been more prone to fishing closer to the border than vessels which were tracked prior to the announcement. Unfortunately, there is no way to rule out that these vessels were fishing at those same spots prior to being tracked by AIS. Should one be willing to assume that the date when the vessels were first tracked and the date when they actually started fishing coincided, then this would reinforce the narrative that the implementation of these "no-fishing" areas has been somewhat taken seriously by the fishers, *i.e.* that fear of getting apprehended may have pushed them to move their operations closer to the boundary of the restricted area. Should that be the case, it could imply that fishers are more worried of physical interception by the managing authority than they are of being observed practicing illegal fishing effort by satellite.

Figure B.1: Displacement of fishing effort, all vessels

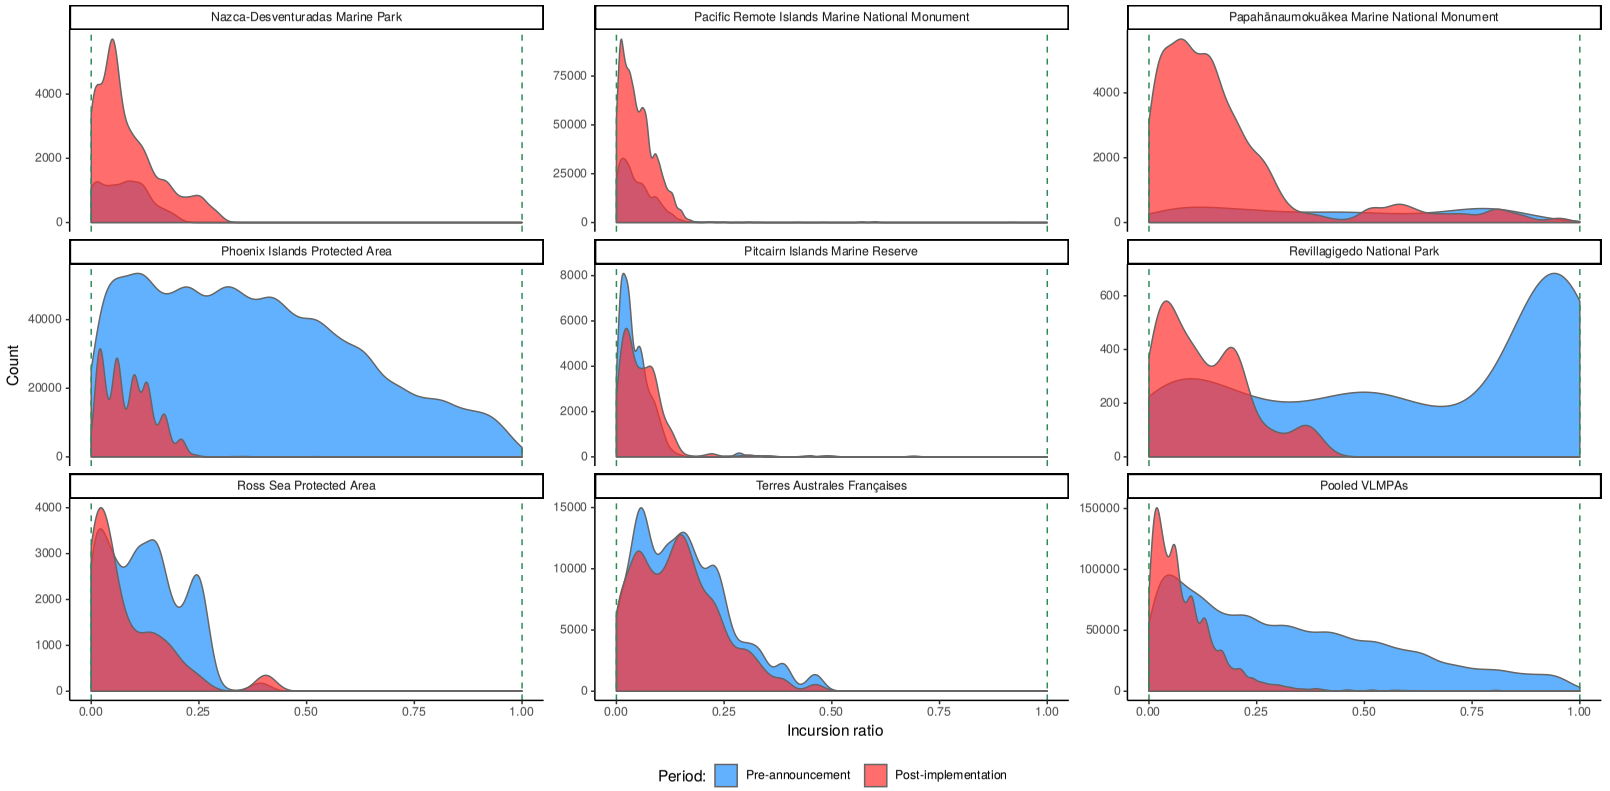

Comparison of the annualized count of fishing effort inside the VLMPAs prior to the announcement of their creation and after their implementation. All vessels are included, regardless of the date of their first appearance in the GFW database.

Figure B.2: Evolution of the number of vessels by country

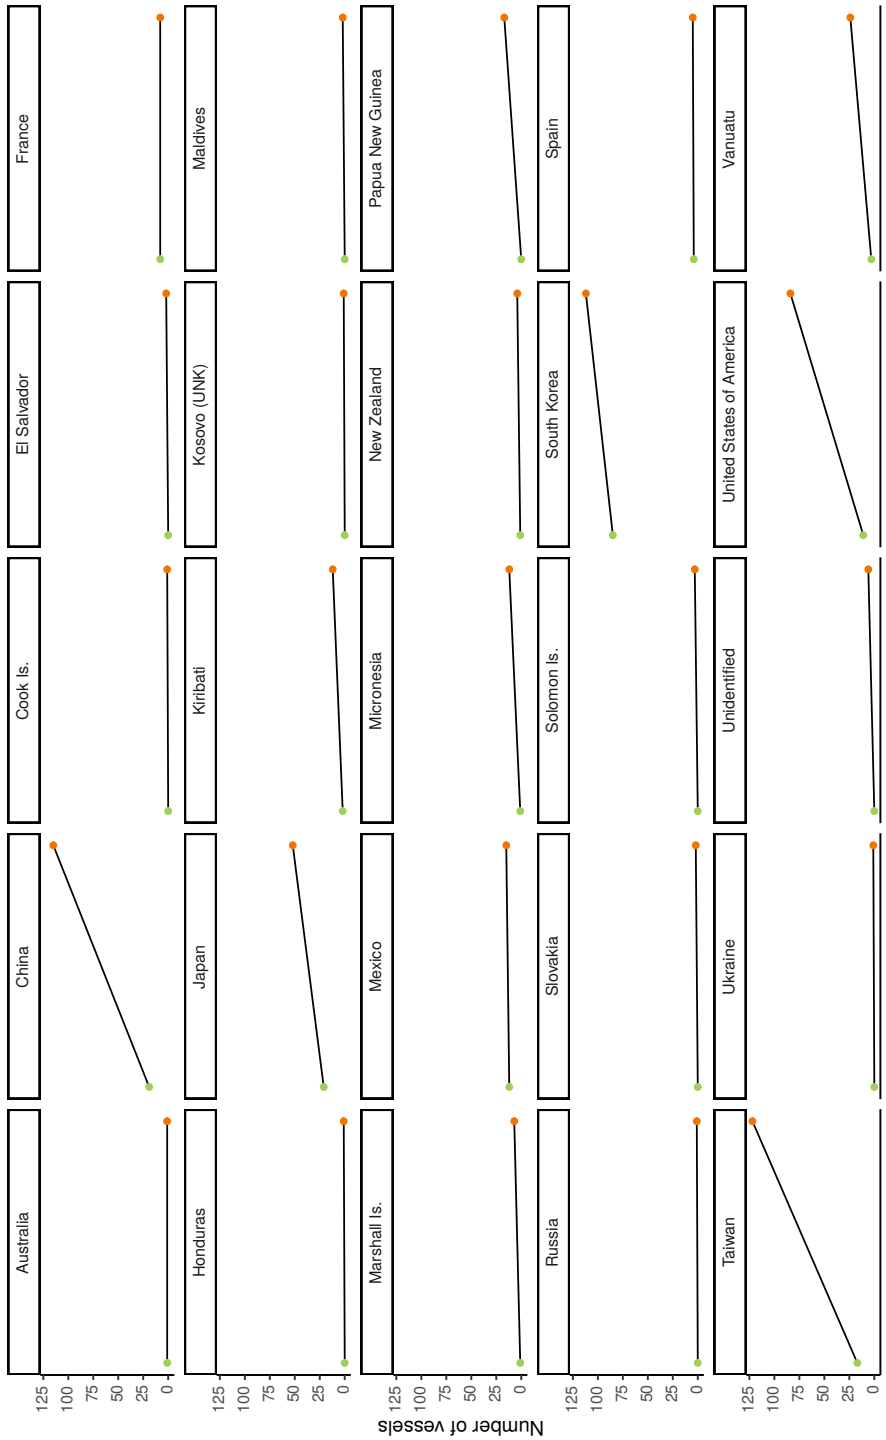

Number of vessels practicing banned fishing in any of the 8 VLMPAs. *Restricted vessels* refer to the vessels which were tracked prior to the announcement that a protected area would be created. *All* refer to all the vessels to ever practice banned fishing, regardless of the date of their first appearance in the database.

B.2 VLMPA Creation Timelines

Ross Sea Protected Area

Announced by the Commission for the Conservation of Marine Living Resources (CCAMLR) on October 27, 2016, the Ross Sea protected area came into force on December 1, 2017 and has become the largest MPA in the world. The VLMPA's no-take zone covers three separate areas that cover 1.12 million km^2, or 72% of the 1.55 million km^2.

Papahānaumokuākea Marine National monument

The Papahānaumokuākea Marine National monument was originally established on June 15, 2006 by the Bush administration. Efforts to expand the monument started on January 29, 2015 when a group of prominent Native Hawaiians wrote to president Barack Obama asking the federal government to expand its protections around the Northwestern Hawaiian Islands. The monument was expanded from its original size of 362,073 km^2 to 1,508,870 km^2 on August 26, 2016 by presidential proclamation. It is now the fourth largest MPA in the world. It is a fully "no-take" area.

Pacific Remote Islands Marine National Monument

The Pacific Remote Islands Marine National Monument was originally established by presidential proclamation on January 6, 2009. On May 20, 2014 of a scientific report promoting the expansion of the existing monument was published. On September 25, 2014 president Barack Obama expanded the Pacific Remote Islands Marine National Monument from 215,000 km^2 to 1,270,000 km^2, making it the sixth largest MPA to date.

Terres Australes Françaises

On October 3, 2006, an inter-ministerial decree established the original natural reserve in the territorial waters of the French Southern Territories. On June 14 and June 15, 2016, the scientific council in charge of monitoring the reserve discussed the possibility of extending the coverage of the reserve (TAAF, 2017). On December 12, 2016, the French government published a decree extending the existing MPA by 672,969 km^2, setting up "no-take" zones, covering more than 120,000 km^2. Enforcement started on March 31, 2017, upon publication of the prefectoral order.

Pitcairn Islands Marine Reserve

On March 18, 2015, the BBC reported the British government's intention to create a marine reserve covering 834,000 km^2 around the Pitcairn Islands, once monitoring and enforcement

of the reserve could be funded (BBC News, 2015). The ordinance implementing the area was then promulgated on September 12, 2016.

Phoenix Islands Protected Area

The area was established in 2006. For years however, the government of Kiribati claimed that the Phoenix Islands Protected Area was a "no-take area" despite evidence that it was the most fished MPA in the world. Eventually, the cabinet of president Anote Tong voted on January 29, 2014 to actually close the reserve by January 1, 2015. News of the vote only broke on May 9, 2014 (Science News, 2014), since no public announcement was made at the time of the vote.

Nazca-Desventuradas Marine Park

In February 2013, the National Geographic and Oceana launched the "Pristine Seas" expedition in the Desventuradas islands, off the coast of Chile. The expedition led to the publication on February 23, 2013 of a report calling for the creation of an MPA (National Geographic Society & Oceana, 2013). On August 24, 2016, the Chilean government officially ordered the creation of the Nazca-Desventuradas Marine Park, a 300,000 km^2 "no-take" reserve.

Revillagigedo National Park

On July 17, 2016, the United Nations Educational, Scientific and Cultural Organization (UNESCO) recognized the Revillagigedo islands as a World Heritage site (UNESCO, 2016), a first step toward the sanctuarization of the area. On November 24, 2017, Mexican president Enrique Peña Nieto officially created of the Revillagigedo National Park, a "no-take" protected area spanning 150,000 km^2.