



Essays on the Economics of Climate Change

Citation

Ranson, Matthew. 2012. Essays on the Economics of Climate Change. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9299646>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

©2012 - Matthew H. Ranson
All rights reserved.

Essays on the Economics of Climate Change

Abstract

This dissertation studies three aspects of the economics of climate change: how rising sea levels will affect coastal homeowners in Florida; how changes in weather will affect the prevalence of crime in the United States; and why skepticism about climate change is so common among the general public.

Chapter 1 uses housing market data to estimate the welfare costs of shoreline loss along coastal beaches in Florida. I develop a structural housing market model and use it to provide a welfare interpretation for the coefficients from a new “discontinuity matching” hedonic research design. Using housing sales data, beach width surveys, and historical beach nourishment records, I then estimate Florida homeowners’ willingness to pay for an extra foot of sand. I find that changes in beach width have little impact on housing prices, except possibly at very eroded beaches.

Chapter 2 estimates the impact of climate change on the prevalence of criminal activity in the United States. The analysis is based on monthly crime and weather data for 2,972 U.S. counties from 1960 to 2009. The results show that temperature has a strong positive effect on criminal behavior, and that between 2010 and 2099, climate change will cause an additional 35,000 murders, 216,000 cases of rape, 1.6 million aggravated assaults, 2.4 million simple assaults, 409,000 robberies, 3.1 million burglaries, 3.8 million cases of larceny, and 1.4 million cases of vehicle theft. The social cost of these climate-related crimes is between 20 and 68 billion dollars.

Chapter 3 develops a model of rational skepticism about policy-relevant scientific questions. Many policy debates have three features: first, individuals initially

disagree about some scientific question; second, new evidence about the question becomes available; and third, the evidence may be systematically biased. Under these conditions, Bayesian disagreements persist even in the face of an infinite quantity of new evidence. Furthermore, Bayesian updating based on the new evidence produces “skeptics”, in the sense that individuals whose prior beliefs conflict most with the observable evidence end up with the most extreme posterior beliefs about the degree of bias. These results provide insight into the phenomenon of climate skepticism.

Contents

Title page	i
Abstract	iii
Table of Contents	v
List of Figures	viii
List of Tables	ix
Acknowledgments	x
Introduction	xi
1 What Are the Welfare Costs of Shoreline Loss? Housing Market Evidence from a Discontinuity Matching Design	1
1.1 Introduction	2
1.2 Background	6
1.3 Theory	10
1.3.1 Hedonic Model	10
1.3.2 Discontinuity Matching Research Design	15
1.4 Data	20
1.4.1 Beach Width Data	20
1.4.2 Property Sales Data	22
1.4.3 Nourishment Data	23
1.4.4 Summary Statistics	25
1.5 Econometric Approach	30
1.5.1 Repeat Sales Panel Approach	30

1.5.2	Differences-in-Differences Approach	32
1.5.3	Discontinuity Matching Approach	36
1.6	Results	42
1.6.1	Repeat Sales Results	42
1.6.2	Differences-in-Differences Results	43
1.6.3	Discontinuity Matching Results	56
1.7	Discussion and Policy Implications	59
1.7.1	Summary of Main Results	59
1.7.2	Cost-Benefit Analysis of Beach Nourishment	60
1.7.3	Discussion	62
1.8	Conclusion	63
2	Crime, Weather, and Climate Change	65
2.1	Introduction	66
2.2	Background on Weather and Crime	69
2.3	Data	71
2.3.1	Data Sources	71
2.3.2	Summary Statistics	74
2.4	Methodology	78
2.5	Results	83
2.5.1	Weather and Crime Rates	83
2.5.2	Climate Change and Crime Rates	86
2.6	Discussion	93
2.7	Conclusion	95
2.8	Appendix	97
3	Bayesian Updating with Potentially-Biased Evidence: A Model of Climate Skepticism	105
3.1	Introduction	106

3.2	General Model	110
3.3	Application to Climate Skepticism	111
3.3.1	Model	111
3.3.2	Results	114
3.4	Discussion	117
3.5	Conclusion	119
3.6	Appendix	120
4	References	122
4.1	References for the Introduction	123
4.2	References for Chapter 1	124
4.3	References for Chapter 2	133
4.4	References for Chapter 3	138

List of Figures

1.1	Map of the Study Area	21
1.2	Location and Timing of Surveys and Nourishment Projects	28
1.3	Distribution of Beach Width, by Nourishment Frequency	29
1.4	Panel Regression Results: Capitalized Price Versus Beach Width	45
1.5	Panel Regression Results: Capitalized Price Versus Beach Width, Detail	46
1.6	Trend in Sales Price, by Year of Most Recent Nourishment	48
1.7	Differences-in-Differences Results	53
1.8	Differences-in-Differences Results by Nourishment Intensity	55
1.9	Discontinuity Matching Results: Trends in Width and Price	58
2.1	Map of the Study Region	76
2.2	Crime Rate Trends, by Climate Zone	79
2.3	Seasonal Weather Patterns, by Climate Zone	80
2.4	Seasonal Crime Rate Trends, by Climate Zone	81
2.5	The Effect of Daily Maximum Temperature on Monthly Crime	85
2.6	Distribution of Daily Weather, by Scenario	90
2.7	Present Discounted Social Cost of Climate-Related Crime, Per Person	93
2.8	Monthly Crime and Daily Temperature, by Decade	98
2.9	Monthly Crime and Daily Temperature, by Climate Zone	99
2.10	Monthly Crime and Daily Temperature, by Mean Monthly Temperature	101
2.11	Monthly Crime and Daily Temperature, by Crime Rate Quartile	102
2.12	Crime and Daily Temperature, by Month, Quarter, and Half-year	104

List of Tables

1.1	Examples of Major Beach Nourishment Projects	24
1.2	Characteristics of Beach Nourishment Projects	25
1.3	Summary Statistics	26
1.4	Housing Prices and Beach Width: Panel Results	44
1.5	Characteristics of Properties Sold Before and After Nourishment	47
1.6	Beach Width and Nourishment	49
1.7	Beach Width and Nourishment, by Nourishment Intensity	51
1.8	Housing Prices and Nourishment: Differences-in-Differences Results . .	52
1.9	Housing Prices and Nourishment: Diff-in-Diff Results by Intensity . . .	54
1.10	Beach Width Markov Regression Results	56
1.11	Discontinuity Matching Results	59
1.12	Willingness to Pay for Wider Beaches	61
2.1	Uniform Crime Reporting Offense Definitions	72
2.2	Summary Statistics, by Climate Zone	77
2.3	Maximum Daily Temperature and Monthly Crime	87
2.4	The Predicted Impact of Climate Change on Crime	91

Acknowledgments

I thank the many people who have helped me over the last six years:

My advisors: Richard Zeckhauser, for encouraging me think about research in new ways; Robert Stavins, for practical support and candid advice; and Martin Weitzman, for always emphasizing the connection between mathematics and intuition.

The many colleagues, faculty members, and friends who have taken the time to comment on these chapters. I especially thank Tatyana Deryugina, Michael Greenstone, Eduardo Montoya, Erich Muehlegger, John Pratt, and the participants at the Harvard environmental economics lunches.

The administrators who have supported my research, particularly Louisa van Baalen and Nicole Tateosian. I am thankful for funding from a Harvard Kennedy School fellowship, the Joseph Crump fellowship, the Julius B. Richmond fellowship, and a GSAS dissertation completion fellowship.

My family. I am grateful to Boots, for his loyal companionship. Most of all, I thank my wonderful wife, Marianne Cook, for her love and support.

Introduction

In the absence of concerted action by the world's governments, anthropogenic greenhouse gas emissions will cause global temperatures to increase between 3 and 8 degrees Fahrenheit over the next century. At the same time, precipitation patterns will change, snow cover will decrease, sea levels will rise, and a variety of other climatic shifts will take place (IPCC, 2007). Developing rational public policies to address these changes will require researchers to develop credible estimates of the economic impacts of climate change and to communicate this technical knowledge to policy makers and the general public. This dissertation contributes to both of these goals.

Chapter 1 of the dissertation uses housing market data to estimate the welfare costs of shoreline loss along coastal beaches in Florida. In this chapter, I develop a forward-looking structural model of a housing market in which a time-variant housing characteristic (beach width) follows a Markov process. I use this model to provide an exact welfare interpretation for the coefficients from a new "discontinuity matching" hedonic research design. Using a unique panel dataset on housing sales, beach width survey measurements, and the timing of 204 beach nourishment projects along 300 miles of Florida's coastline, I then estimate homeowners' willingness to pay for an extra foot of sand. In contrast to previous work, I find that changes in beach width have little impact on housing prices, except possibly at very eroded beaches. The results imply that the welfare costs of sea level rise may be low up to a threshold, and then increase sharply.

Chapter 2 estimates the impact of climate change on the prevalence of criminal activity in the United States. The analysis is based on a panel of monthly crime, temperature, and precipitation data for 2,972 U.S. counties over the 50-year period from 1960 to 2009. I identify the effect of weather on monthly crime by using a semi-parametric bin estimator and controlling for county-by-month and county-by-year fixed effects. The results show that temperature has a strong positive effect on criminal behavior, with little evidence of lagged impacts. Under the IPCC's A1B climate scenario, the United States will experience an additional 35,000 murders, 216,000 cases of rape, 1.6

million aggravated assaults, 2.4 million simple assaults, 409,000 robberies, 3.1 million burglaries, 3.8 million cases of larceny, and 1.4 million cases of vehicle theft, compared to the total number of offenses that would have occurred between the years 2010 and 2099 in the absence of climate change. The present discounted value of the social costs of these climate-related crimes is between 20 and 68 billion dollars.

As a complement to Chapters 1 and 2, Chapter 3 explores how people learn about policy-relevant scientific questions. Many public policy debates have three features: first, individuals initially disagree about some core scientific question; second, new evidence about the question becomes available over time; and third, there is some possibility that the evidence is systematically biased. I show that under these conditions, Bayesian disagreements persist even in the face of an infinite quantity of new evidence. Furthermore, Bayesian updating based on the new evidence produces “skeptics”, in the sense that individuals whose prior beliefs conflict most with the observable evidence end up with the most extreme posterior beliefs about the degree of bias. These results explain why skepticism about the quality of climate science is most prevalent among people who believe that climate change is unlikely to occur.

As the abstracts above illustrate, there are a variety of themes that cut across the chapters of this dissertation. In this introduction, I reflect on these themes, as well as a broader question: why has it been so difficult to learn about the economic impacts of climate change? I consider several main challenges.

Perhaps the most central difficulty is that economic climate change research involves—in an unescapable, fundamental way—out-of-sample prediction. Although economists are frequently forced to make policy recommendations based on limited data, the degree of extrapolation required for climate change is enormous. An eight degree F increase in global temperatures is completely outside the realm of historical experience: there are no recent global-wide warming episodes of a similar magnitude on which to base economic impact estimates. In the absence of such data, economists have relied on three main categories of methodologies to predict the impacts of climate change on economic outcomes: a production function approach (e.g., Adams et al, 1988; Rosenzweig and Parry, 1994), a cross-sectional “Ricardian” approach based on differences

in climate across locations (e.g., Mendelsohn, Nordhaus, and Shaw, 1994; Cragg and Kahn, 1997), and a panel approach based on differences in weather over time in a particular location (e.g., Deschenes and Greenstone, 2007, 2011; Schlenker and Roberts, 2009; Dell, Jones, and Olken, forthcoming; and Chapters 1 and 2 of this paper). Each approach has strengths and weaknesses.

The strength of the production function approach is its simplicity and modest data requirements. This methodology is based on scientific estimates of the impact of temperatures and precipitation on output from a particular economic activity. For example, to determine how climate change is likely to affect wheat, corn, and soybean production in the Western United States, Adams et al (1988) rely on agricultural science models that describe how growing season temperatures affect the productivity of each type of plant. Because this information is relatively easy to collect (e.g., by studying plants grown under experimentally-controlled greenhouse conditions), the production function methodology provides a straightforward approach for estimating climate change impacts. One weakness, however, is that it fails to account for short-run and long-run margins for adaptation. For example, in the face of drier than normal conditions, farmers might adjust irrigation practices, plant a different mix of crops, or—in the very long run—convert their agricultural land to some other use. Because these opportunities for adaptation reduce the costs of changes in climate, the production function approach will typically overestimate the direct impact of climate change. A second weakness is that this approach does not account for partial and general equilibrium market responses. For example, in years with below average agricultural production, crop prices tend to rise, thus changing the distribution of impacts and incentives across the economy.

An alternative to the production function approach is the cross-sectional “Ricardian” approach proposed by Mendelsohn, Nordhaus, and Shaw (1994), which is based on analysis of differences in economic outcomes across geographic areas with different climates. Relative to the production function approach, the key advantage of the cross-sectional approach is that it provides a method for estimating the long-run effects of climate on economic outcomes. Under the reasonable assumption that people are likely to choose factors of production that are best suited to their local climate, a regression of

economic outcomes on climate will produce estimates that reflect the long-run relationship between climate and economic outcomes. For example, Mendelsohn, Nordhaus, and Shaw (1994) use data on farmland prices and average temperatures and precipitation to estimate the welfare impacts of climate change on agricultural areas of the United States. However, although the use of cross-sectional data does address some concerns about adaptation, it also a weakness. Because there are many other environmental and economic factors—soil type, proximity to the coast, availability of mineral resources, historical settlement patterns—that are correlated with climate, it is possible for cross-sectional regressions to be biased by correlated unobservables. In some contexts, such as agriculture, it seems at least possible to control for the most important variables that affect the outcomes of interest. However, in other contexts, such as residential demand for climate, controlling for all relevant locational characteristics seems less realistic.

A third methodology that has been used to estimate the economic impacts of climate change is a panel approach based on differences in weather over time in a particular location. The advantage of this approach is that it controls for idiosyncratic differences across locations while exploiting presumably random variation in weather over time within each location. As a result, regressions based on this approach include a larger number of data points and thus provide a higher degree of precision. More importantly, when used with fixed effects methods, this panel approach controls for permanent unobservable characteristics that may otherwise bias estimates of the impacts of climate change. However, this methodology also has a serious drawback: individuals' and firms' responses to short-term variation in weather are likely to be quite different from their responses to longer-term changes in climate. Thus, although this approach can better account for short-term adaptation than the production function method, it still may not produce very accurate estimates of long-run impacts.

In addition to the difficulties created by the absence of relevant historical experience with changing climate, economic climate change research must also grapple with a second major challenge: the fact that the greatest changes in climate will occur in the distant future. The fact that large temperature increases may not occur for a century or more raises a variety of issues (see, e.g., Weitzman, 1998; Nordhaus, 2007; Weitz-

man, 2007; Wagner and Zeckhauser, 2011). Some of these relate the deep scientific and economic uncertainties that arise over long time periods. Can impact estimates based on recent experience be extrapolated to predict economic outcomes a hundred years from now? If so, what assumptions should be made about the baseline growth rate of economic activity and technological progress? Will the passage of time create new, unforeseen opportunities for adaptation? Other issues relate to the challenges that long time periods create for framing economic analysis. What discount rate should be used to compare costs and benefits incurred in the present and the distant future, and should this discount rate include ethical considerations? What empirical evidence could shed light on the welfare impacts of extreme changes in climate that—with non-zero probability—might occur? How should these low-probability disasters enter into current cost-benefit calculations? In answering these questions, it is important to be mindful of the poor historical record of attempts to predict far into the future. As I argue in Chapter 2, these uncertainties about the distant future imply that even a carefully constructed analysis of long-term climate impacts is—in a very real sense—little better than an extended “back-of-the-envelope” calculation.

A third set of issues related to estimating the economic impacts of climate change stems from the possibility that marginal impacts are likely to be heterogeneous and nonlinear. For example, although Chapter 1 of this dissertation suggests that the near-term welfare impacts of rising sea levels are not statistically distinguishable from zero, Chapter 2 provides striking evidence that climate change will cause significant changes in crime patterns across the United States. There is also heterogeneity in climate change impacts across different geographic areas: as Chapter 2 demonstrates, the social costs of future climate-related crime in the northern Plains are twice as large as the social costs in the mid-Atlantic and South. Chapter 2 also shows that there are clear non-linearities in the relationship between temperature and property crimes, with non-zero marginal effects only below a threshold at 40 degrees F. Similarly, Chapter 1 presents suggestive (although only modestly significant) evidence that the welfare costs of sea level rise may be low until the amount of remaining beach reaches a threshold level (10 or 20 feet), and then rise sharply.

A fourth set of issues that is a major focus of this dissertation is the challenge of present-

ing information about climate change research in a way that is useful to policy makers and the general public. Chapters 1 and 2 make clear why this is so problematic: the impacts of climate change will be heterogeneous and nonlinear, and the methodologies used to assess the impacts can be highly technical. Nonetheless, as Chapter 3 argues, achieving consensus about appropriate climate policy actions will require effective and transparent scientific communication of both results and methods, regardless of the complexity of the problem of climate change.

Of course, in addition to the issues described here, research on the economics of climate change must address many other challenges: how to aggregate micro-level empirical impact studies to generate predictions of macro-level outcomes; how to characterize the complex political interactions that will determine the future path of greenhouse gas emissions; how to estimate the costs of different abatement policies; how to account for impacts that are difficult to monetize, such as regional biodiversity loss or species extinction; how to evaluate the merits and drawbacks of geoengineering; and how to evaluate the returns to investment in climate change policies relative to investments in policies addressing other social problems such as poverty, education, and health. Although no single study can address all of these issues, the purpose of this dissertation is to provide new research—set in the diverse contexts of coastal housing markets, crime, and Bayesian learning—that contributes to a better understanding of the economics of climate change.

Chapter 1:

What Are the Welfare Costs of Shoreline Loss? Housing Market Evidence from a Discontinuity Matching Design

1.1 Introduction

Many sections of the U.S. coastline are severely eroding. The average long-term rate of shoreline loss along the New England and mid-Atlantic coasts is 1.6 feet per year, with much higher rates in areas such as southern Nantucket Island in Massachusetts (12 feet per year) and the southern portion of the Delmarva Peninsula in Maryland (9.5 feet per year) (Hapke et al, 2010; Woods Hole, 2000). In Florida, some segments of beach lose as much as ten to twenty feet per year (FDEP, 2000, 2001). Under current predictions of a 1.1 foot rise in average sea levels by the year 2100, these erosion rates will accelerate, and between 3,000 and 7,000 square miles of dry land could be lost (IPCC, 2007; Titus, 1989). Although it is possible to protect coastal areas from shoreline loss—through installation of hardened features such as seawalls and groins, imposing set-back and minimum-height home construction requirements, and performing periodic nourishments to place new sand onto eroded beaches—the costs are substantial. For example, the Florida Department of Environmental Protection projects that 1.1 billion dollars would be needed between 2011 and 2015 for full implementation of the state’s strategic beach management plan.

Surprisingly, given both the substantial costs of preventing coastal erosion and the serious risks posed by retreating shorelines and rising sea levels, there is little rigorous evidence on the benefits of wider beaches for coastal property owners. Existing studies suggest that the sale price of a coastal home increases between \$70 and \$8,000 per one foot increase in beach width (Gopalakrishnan et al, 2010; Landry, Keeler, and Kriesel, 2003; Pompe and Rinehart, 1995). However, since cross-sectional hedonic property value regressions suffer from well-known theoretical and econometric problems, the interpretation of these estimates is not clear. Coefficients from hedonic regressions are biased when one of the housing attributes (such as beach width) varies over time (Abelson and Markandya, 1985). Furthermore, cross-sectional hedonic regressions are vulnerable to problems with omitted variable bias (Chay and Greenstone, 2005; Kuminoff,

Parmeter, and Pope, 2010)—as might be the case if higher quality houses are built along wider sections of beach.

In this chapter, I estimate the welfare costs of shoreline loss along coastal beaches in Florida, using three distinct research designs that each solve both of the theoretical and econometric problems discussed above. These research designs are: (1) a repeat-sales regression of housing prices on beach width that controls for fixed housing characteristics and aggregate housing market shocks; (2) a differences-in-differences approach based on the sharp and substantial discontinuity in beach width caused by beach nourishment projects; and (3) a new “discontinuity matching” approach that exploits capitalized housing price differentials created by government policies that result in predictable changes in future beach width. In all three approaches, I take seriously the problem of giving a theoretical interpretation to the estimated coefficients.

This chapter makes three main contributions. First, using the empirical approaches discussed above, I develop the first panel data estimates of homeowners’ marginal willingness to pay to avoid coastal shoreline loss. My analysis is based on a unique dataset that includes 1.1 million housing sales transactions at parcels located within five kilometers of a coastal beach in sixteen Florida counties between 1983 and 2009 (the dataset includes 388 miles of coastline). I link these data to high-resolution beach width survey records at fixed monuments located approximately 1000 feet apart along the Florida coastline. Finally, I add information about the timing, location, and volume of sand for 204 beach nourishment projects. This list represents the most detailed dataset of Florida nourishment projects ever compiled.

Second, I develop a “Rosen-like” structural housing market model that provides an intuitive interpretation for the coefficients from hedonic regressions of housing prices on a time-varying neighborhood characteristic (such as beach width). When homebuyers have rational expectations and changes in the characteristic are Markovian, the model equilibrium implies that the cross-sectional relationship between housing prices and

characteristics has an exact interpretation as willingness to pay for a policy intervention that increases current amenity quality by one unit and then allows it to evolve in an unconstrained way in future periods. Unlike previous work, which has treated the coefficients from panel hedonic regressions as biased estimates of marginal willingness to pay for a *permanent* increase in amenity quality (Abelson and Markandya, 1985; DiPasquale and Wheaton, 1994; Bayer, McMillan, Murphy, and Timmins, 2008; Bishop and Murphy, 2011), I argue that it is much more accurate—and useful—to interpret these coefficients as willingness to pay for a *one-time* marginal policy intervention (such as beach nourishment).

The chapter's third contribution is to develop a new "discontinuity matching" research design that recovers consumers' marginal willingness to pay by exploiting capitalized housing price differentials caused by construction projects that result in predictable changes in future amenity quality. Recent empirical work on hedonic models has considered several sources of identifying variation in amenity quality, including unexpected shocks (Davis, 2004; Greenstone and Gallagher, 2008; Kuminoff and Pope, 2010a,b) and cross-sectional discontinuities resulting from arbitrary geographic boundaries such as school district borders (Black, 1999). In contrast, my discontinuity matching approach exploits the change in capitalized housing prices that accompanies *predictable* discontinuities in amenity quality. Typically, these discontinuities will be the result of a policy intervention, e.g., beach nourishment, construction of a new school, or completion of a public transportation project. For example, suppose that two otherwise-similar houses are located on two different beaches, one of which—by random chance—is heavily eroded this year. If the government announces that the eroded beach will be nourished next year, then prospective homebuyers would rationally expect the two beaches to have similar width next year (and in all future periods). Thus, a comparison of current prices and current beach width across these two houses will reveal the marginal rental value of living on a wider beach for one year.

The chapter establishes several empirical results. First, using semi-parametric panel

regressions, I find that beach width has only a modest effect on housing prices. According to these regressions, the difference in sales price between a house with a 200 foot wide beach and a house with 50 feet of beach is only about 2.1 percent. However, houses with less than 20 feet of beach do experience a suggestive—but only marginally significant—price discount of 6 to 14 percent. Second, my differences-in-differences regressions show beach nourishment adds a statistically significant 83 feet to the width of the average beach. However, housing prices only gain an insignificant 1.2 percent between two years before and after nourishment, and I can reject the possibility that housing prices increase by more than 4.9 percent during this period. Finally, using the discontinuity matching approach, I estimate that the yearly rental value of an extra foot of beach is approximately \$29 per household, and not statistically different from zero. Overall, the results imply that the welfare costs of sea level rise may be low up to a threshold, and then increase sharply.

This chapter builds on a growing literature on the microfoundations of hedonic models (Rosen, 1974; Roback, 1982; Bajari and Benkard, 2005; Bishop and Timmins, 2008a,b; Kuminoff and Jarrah, 2010). A few studies have cast housing choice as a dynamic utility maximization problem (DiPasquale and Wheaton, 1994; Bayer, McMillan, Murphy, and Timmins, 2008; Bishop and Murphy, 2011) or as a dynamic process with slow adjustment (Riddell, 2001; Mankiw and Weil, 1989), or have modeled neighborhood characteristics as dynamic processes (McCluskey and Rausser, 2001). Others have developed methodologies for using panel data to identify hedonic regressions, either using new discrete choice methods (Bajari et al, 2010; Kueth, Foster, and Florax, 2008) or a repeat sales methodology employing first-differencing or fixed effects (Palmquist, 1982; Mendelsohn et al, 1992). However, to the best of my knowledge, there are no hedonic studies that attempt to estimate rental prices by decomposing sales prices into current and future rental components. Although authors do recognize that the price on the left side of a hedonic equation should be interpreted as a discounted sum of rental prices (Dougherty and Van Order, 1982; Abelson and Markandya, 1985; Blackley and

Follain, 1996; Meese and Wallace, 2003; Bajari and Kahn, 2007; Diewert, Nakamura, and Nakamura, 2009) or as a sum of use and option values (Plantinga, Lubowski, and Stavins, 2002), in practice, most studies use a “static-equivalent” rental price calculated by multiplying the sales price by the discount rate (Bajari and Kahn, 2005; Gyourko and Tracy, 1991; Bishop and Murphy, 2011). The chapter is also related to more macro-oriented literatures on housing markets, for example, literatures on calculating price indices and implicit rents for owner-occupied housing (Case and Schiller, 1989; Rondinelli and Veronese, 2011), evaluating the relationship between rental prices and sales prices (Gallin, 2004), and assessing the welfare impacts of changes in housing prices (Bajari, Benkard, and Krainer, 2005).

The remainder of this chapter is organized as follows. Section 1.2 provides background on coastal shoreline retreat and beach nourishment. Section 1.3 describes a structural model of a housing market and demonstrates how it can be used to calculate willingness to pay for wider beaches. Section 1.4 describes my dataset and presents summary statistics. Section 1.5 explains the details of my econometric approach, and Section 1.6 presents my main empirical results. Section 1.7 discusses the results, and Section 1.8 concludes.

1.2 Background

The United States has more than 12,000 miles of coastline, of which a significant portion consists of sandy beaches (NationalAtlas.gov, 2011). Unlike dry land, coastal beaches are highly dynamic physical environments that experience significant seasonal and yearly changes. For example, many beaches erode during the winter, due to heavy waves, and then accrete during milder summer weather. Over longer time horizons, beaches exhibit a variety of erosional patterns that depend on natural factors such as the underwater coastal profile, dominant wave and weather patterns, and major storm events (NRC, 1995).

Because proximity to the coast provides a variety of benefits, the land along many beaches is heavily developed. Unfortunately, historical development patterns in many areas have failed to anticipate the degree to which erosion can reshape the coastline. Furthermore, some types of development, such as the dredging of coastal waterways and inlets, have contributed substantially to erosion problems. Thus, many properties that once looked out over wide coastal beaches now face serious problems with shoreline loss.

Policy responses to shoreline loss take several main forms (NRC, 1995). The first is the construction of hardened features, such as seawalls and jetties, that are intended to protect buildings and prevent sand from moving along the coast. Although these features can succeed in trapping pockets of sand, they sometimes have the perverse result of creating leeward hotspots in which erosion patterns are magnified. A second policy option is establishing legal permitting requirements that require new houses to be set back some specified distance from the beach. While this approach is workable in undeveloped areas, it has the obvious drawback of failing to address erosion problems at existing homes. A third option is abandonment and retreat. This is considered an option of last resort.

One final policy response to coastal erosion along sandy beaches is beach nourishment. In a typical nourishment project, sand from an offshore borrow area is pumped or dredged onto a beach to make it wider (NRC, 1995). Because the volume of sand required for nourishment projects is quite large—as high as several million tons—locating suitable sources of sand is a major challenge for these projects. Furthermore, the process is expensive: nourishment costs approximately one million dollars per mile of beach (USACE, 1996). Because of this high cost, localities often obtain state and federal funding for nourishment projects. For example, between 1950 and 1993, the U.S. Army Corps conducted 56 large beach nourishment projects that covered a total of 210 miles of U.S. shoreline. The cumulative federal cost share for these projects was \$881 million dollars (USACE, 1996). The NOAA Coastal Resources Center (2009) reports

that federal, state, and local organizations have spent at least \$2.5 billion dollars on 242 major beach nourishment projects since 1950.

Although the costs of policy responses such as beach nourishment are not difficult to calculate, the benefits of these policies are less clear. The central question is: what is the value of widening a particular section of beach? This question is complicated by the fact that beaches provide a variety of economic benefits to local communities, including recreational opportunities, scenic views, protection from coastal storms, and tourism revenues.

In this chapter, I focus on estimating only one component of the economic contribution of beaches: the welfare benefits to beachfront homeowners. Although there are many reasons why beaches may be valuable, their contribution to the welfare of local residents is likely to be one of the most important. Furthermore, by focusing on the economic benefits to homeowners, I am able to use a hedonic property value approach exploits the relationship between housing prices and beach quality (Rosen, 1974).

There is a small existing literature that uses such hedonic techniques to estimate the benefits of wider beaches to local homeowners. To the best of my knowledge, all of this previous work focuses on the cross-sectional relationship between beach width and housing prices.¹ The most recent of these studies is Gopalakrishnan et al (2010), who instrument for beach width using distance from the continental shelf and beach attributes such as scarps.² They find that in a cross-section of coastal properties in ten North Carolina towns, a one-foot increase in beach width is associated with a 1.1 percent increase in property values (about \$8,800). Although their empirical strategy does control for the potential endogeneity of beach nourishment decisions, it does not address the possibility that higher-quality houses are more likely to be built on wider

¹There has also been theoretical work on beach nourishment. Most notably, Smith et al (2009) discuss beach nourishment as an example of a dynamic capital accumulation problem. They show that nourishment frequency depends on whether sand erodes at a rate greater or less than the discount rate.

²Other than Gopalakrishnan et al (2010), I am aware of no other studies of the benefits of beach width that use quasi-experimental methods.

beaches.

This cross-sectional literature also includes a variety of earlier studies. For example, Landry, Keeler, and Kriesel (2003) estimate the benefits of beach nourishment for Tybee Island, Georgia, using the cross-sectional relationship between beach width and property values, for 318 properties sold between 1990 and 1997. They find that a one meter increase in beach width increases property values by \$213. Using similar techniques, Pompe and Rinehart (1995) find that increasing beach width by one foot increases beachfront property values by \$558 to \$754 in the Grand Strand area of North Carolina. Properties half a mile inland also benefit by \$165 to \$254. Other authors have also used hedonic techniques to evaluate the benefits of beach nourishment, but their empirical strategies are less rigorous. For example, Edwards and Gable (1991) use a hedonic model to estimate the value of proximity to a beach in South Kingstown, Rhode Island, and then calculate the benefits of beach nourishment by assuming that nourishment prevents beaches from becoming unusable. Parsons and Powell (2001) use a similar methodology to evaluate the benefits of beach nourishment in Delaware. Woglom (2003) develops a simulation model of beach nourishment, using parameter estimates from earlier studies.

Additionally, a few related studies estimate the value of proximity to the coastline, but do not directly analyze beach width. For example, Milon, Gressel, and Mulkey (1984) use cross-sectional regressions to estimate how the value of a home depends on its distance from the mean high water (MHW) mark, for homes in Apalachicola Bay, Florida. Other similar studies include Bin et al (2008), Parsons and Wu (1991), and Brown and Pollakowski (1977). Kriesel, Randal, and Lichtkoppler (1993) consider the value of erosion protection in the Great Lakes. Bell and Leeworthy (1990) and Bin et al (2005) use a travel cost approach to estimate the recreational benefits of beach days, but do not consider the impact of beach width on willingness to pay.

As discussed in the introduction, the interpretation of estimates from this existing body

of work faces several challenges, given potential theoretical and empirical problems. When neighborhood attributes (such as beach width) vary over time, a regression of housing prices on the time-variant attribute does not identify true marginal willingness to pay for a permanent increase in attribute quality (Abelson and Markandya, 1985). Furthermore, since cross-sectional hedonic regressions are vulnerable to problems with omitted variable bias (Chay and Greenstone, 2005; Kuminoff, Parmeter, and Pope, 2010), it is entirely possible that previous work has found a positive relationship between housing prices and beach width because higher quality houses are built along wider sections of beach. Thus, I devote the remainder of this chapter to developing several research designs for estimating homeowners' willingness to pay for wider beaches that solve both the theoretical and econometric issues in previous work.

1.3 Theory

1.3.1 Hedonic Model

In this section, I develop a simple structural model of a housing market, based on Rosen (1974), that explicitly considers the time dimension of housing choice. The model has two purposes. First, it provides a welfare interpretation for the coefficients from panel hedonic regressions of housing prices on time-varying neighborhood characteristics. Second, the model suggests a new “discontinuity matching” research design that can be used to recover homeowners' implied valuations of neighborhood characteristics. This procedure exploits capitalized housing price differentials caused by policy interventions that lead to predictable improvements in future neighborhood amenity quality (e.g., interventions such as beach nourishment, construction of a new school, or completion of a public transportation project).

The model differs from other recent multi-period hedonic models (DiPasquale and Wheaton, 1994; Bayer, McMillan, Murphy, and Timmins, 2008; Bishop and Murphy,

2011) in several ways. First, my model is not dynamic, in the sense that all choices are made in the first period, with no possibility of re-optimization or sorting in future periods. However, because these initial choices do reflect consumers' expectations about how housing characteristics will change in the future, the model still allows for an "as-if" dynamics that captures much of the intuition and theoretical content of a fully-dynamic model. Second, the model excludes transaction costs. Although real-estate fees and moving costs do contribute substantially to the cost of purchasing a house, inclusion of these transaction costs would complicate the model without changing its fundamental conclusions. Third, consumers in my model have advance knowledge about policy interventions that improve the quality of neighborhood characteristics. This stands in contrast to other recent work, in which policy interventions are modeled as unpredictable shocks (Davis, 2004; Greenstone and Gallagher, 2008; Kuminoff and Pope, 2010a,b).

The model is as follows. Suppose that there are many heterogeneous consumers, indexed $1, \dots, j, \dots, J$, each of whom has preferences θ_j and receives fixed income \bar{y}_j each period. Both θ_j and \bar{y}_j can vary across individuals, but are constant over time. There are also many houses, indexed $1, \dots, i, \dots, I$, each of which has a time-invariant characteristic π_i that represents permanent housing quality (e.g., a composite index measuring the number of bedrooms, square footage, and ceiling height) and a time-variant characteristic w_{it} . Although w_{it} could represent any housing or neighborhood amenity that changes over time, for expositional purposes, suppose that all houses are located on the coast, and that w_{it} measures the width of beach between house i and the high-tide mark. The evolution of w_{it} over time (due to erosion and accretion of sand) is described by the following assumption:

Assumption 1. Markov Amenity Quality: *The time-variant amenity w_{it} follows a Markov process. Furthermore, all houses share the same Markov transition probabilities for amenity quality, given by the transition function $T(w', w)$:*

$$\Pr(w_{i,t+1} = w' | w_{it} = w) = T(w', w) \quad (1.1)$$

Immediately before period 1 begins, each consumer takes out a loan, based on her future income, and uses the loaned money to purchase a house at the market price $p(w_1, \pi)$. The model timing is such that at the time the consumer purchases a house, she knows with certainty the amenity value in period 1 and has rational expectations about the amenity values in period 2 onwards. Let r be the competitive interest rate. Then, each period, the consumer pays the mortgage payment $r \cdot p(w_1, \pi)$ and uses her remaining income that period to purchase c units of a “composite” good with unit price 1, where the composite good represents a mixture of any other goods and services that the consumer finds desirable: food, entertainment, transportation, etc. A consumer’s utility from consuming c units of the composite good and owning a house with current characteristics w and π for one time period is given by $u(w, \pi, c; \theta_j)$. The shared pure rate of time preference is ρ .

Let Γ represent the set of combinations of characteristics $\{(w_{1,1}, \pi_1), \dots, (w_{i1}, \pi_i), \dots, (w_{I1}, \pi_I)\}$ of all houses in period 1. For analytical convenience, I formulate each consumer’s maximization problem as a choice of characteristics, rather than as a choice of discrete houses. In other words, rather than choosing a house i from the set of available houses $\{1, \dots, I\}$, each consumer chooses a combination of characteristics (w_1, π) from the set of available characteristics Γ . Thus, dropping the subscripts i that index houses, consumer j ’s maximization problem is:

$$\begin{aligned} \max_{(w_1, \pi), c \in \{\Gamma, \mathbb{R}_+\}} E \left[\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t u(w_t, \pi, c; \theta_j) \right] \\ \text{s.t.} \quad \begin{cases} \sum_{t=1}^{\infty} \left(\frac{1}{1+r} \right)^t \bar{y}_j \geq p(w_1, \pi) + \sum_{t=1}^{\infty} \left(\frac{1}{1+r} \right)^t c \\ \Pr(w_{t+1} = w' | w_t = w) = T(w', w) \end{cases} \end{aligned} \quad (1.2)$$

Equation (1.2) is a straightforward expected utility maximization problem. The consumer chooses beach width, permanent housing quality, and the composite good in such a way as to maximize the expectation of the present discounted flow of future

utility, while still satisfying the intertemporal budget constraint that the present discounted sum of future income must be greater than or equal to the price of the selected house plus the present discounted sum of future expenditures on the composite good. Note here the analytical value of Assumption 1.1. Even though consumers are forward-looking, the “memorylessness” property of Markov processes allows the price of house i to be represented as a function of only two variables: w_{i1} and π_i . In the beach width example, the assumption implies that once a prospective homebuyer observes the current width of the beach in front of house i , information about beach width in previous periods provides no additional information about whether the beach is likely to erode or accrete in the future. Thus, Assumption 1.1 collapses a vector of past and current measurements and future beliefs about beach width into a single metric: current beach width. This greatly simplifies the formulation and solution of the consumer’s choice problem.

Under some mild regularity conditions (e.g., that the market is sufficiently thick that there is no need to consider corner solutions caused by gaps in the continuum of housing characteristics), the solution to (1.2) is characterized by the following first-order conditions:

$$\frac{\partial}{\partial w_1} E \left[\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t u(w_t, \pi, \bar{y}_j - rp(w_1, \pi); \theta_j) \right] = 0 \quad (1.3)$$

$$\frac{\partial}{\partial \pi} E \left[\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t u(w_t, \pi, \bar{y}_j - rp(w_1, \pi); \theta_j) \right] = 0 \quad (1.4)$$

Let u_j represent the utility function for consumer j with preferences θ_j . Interchanging the order of the differentiation and expectation operators in Equation (1.3) leads to the following result:

Theorem 1. Welfare Interpretation of Panel Hedonic Regressions: *Suppose that in equilibrium, consumer j purchases house i . Consider a counterfactual marginal policy intervention that would increase the initial (period 1) quality of the time-varying amenity w_{i1} at house i by one unit, and then allow it to evolve freely in future periods according to the Markov process*

described in Assumption 1. Consumer j 's willingness to pay for this intervention is given by the derivative of equilibrium housing prices with respect to current amenity quality:

$$\left. \frac{\partial p}{\partial w_1} \right|_{w_{i1}} = \frac{\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[\frac{\partial u_j}{\partial w} \Big|_{w_t} \frac{\partial w_t}{\partial w_1} \Big|_{w_{i1}} \right]}{\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[r \frac{\partial u_j}{\partial c} \Big|_{w_t} \right]} \quad (1.5)$$

This theorem, which generalizes Rosen's (1974) result, provides a welfare interpretation for the empirical relationship between housing prices and amenity quality.³ The left-hand side of Equation (1.5) is the derivative of housing price with respect to the time-varying amenity (e.g., beach width). This derivative can be directly estimated as the coefficient from a regression of housing prices on beach width. The right-hand side of Equation (1.5) represents consumer j 's willingness to pay—by giving up some of the composite good each time period—to achieve the increase in expected utility caused by starting at an initial amenity level that is one unit higher. Thus, unlike Rosen's original theorem, which interprets the relationship between housing prices and amenity quality as willingness to pay for a permanent one-unit increase in amenity quality, Equation (1.5) expresses willingness to pay for a policy intervention that increases initial amenity quality by one unit and then allows it to evolve in an unconstrained way in future periods, according to the Markov transition function specified in Assumption 1.

The value of Theorem 1 is that it provides an exact welfare interpretation for the coefficients from hedonic regressions with time-varying characteristics. It is well known from previous work that panel regressions produce biased estimates of marginal willingness to pay for a permanent increase in amenity quality (Abelson and Markandya, 1985; DiPasquale and Wheaton, 1994; Bayer, McMillan, Murphy, and Timmins, 2008;

³In his seminal 1974 paper, Rosen develops a one-period housing market model in which heterogeneous consumers purchase houses of different quality. He argues that in this static equilibrium, the relationship between price and amenity quality reflects the marginal consumer's willingness to pay for a marginal increase in the amenity. In my notation, his conclusion can be written as: $\frac{\partial p}{\partial w} = \frac{\partial u_j}{\partial w} / \frac{\partial u_j}{\partial c}$. However, because his model has only one period, it is not applicable to situations in which w varies over time.

Bishop and Murphy, 2011). However, when the time-varying characteristic follows a Markov process, Theorem 1 provides a useful economic interpretation for such regressions. For example, in the context of this chapter, Theorem 1 implies that the coefficient from a regression of housing prices on beach width can be interpreted as homeowners' willingness to pay for a one-time beach nourishment project that widens the beach by one foot in the current year.

1.3.2 Discontinuity Matching Research Design

Theorem 1 provides a valuable interpretation for the relationship between observed housing prices and consumers' preferences. However, in many circumstances, it is desirable to know consumers' exact marginal willingness to pay for a *permanent* increase in amenity quality. Thus, in this section, I use the structural housing model from the previous section to motivate a new "discontinuity-matching" research design that can be used to estimate consumers' marginal willingness to pay for a guaranteed, immediate, one-period, one-unit increase in amenity quality. This estimate can be then scaled, using the discount rate, to generate an estimate of MWTP for a permanent increase in amenity quality.

The core idea of the discontinuity matching design is to exploit capitalized housing price differentials caused by predictable discontinuities in future neighborhood amenity quality. Typically, these discontinuities will be the result of a policy intervention, e.g., beach nourishment, construction of a new school, or completion of a public transportation project. Because the price of a house reflects the capitalized value of the future flow of utility from owning that house, predictable improvements in future amenity quality will be reflected in current prices. Thus, current price differences—between matched sets of houses that are expected to have similar post-intervention (post-discontinuity) amenity values—reflect only current differences in amenity quality.

For additional intuition, consider the following example. Imagine two otherwise-identical houses located on two different sections of beach. Suppose that both beaches have similar rates of erosion, but that due to random fluctuations, one of the beaches is heavily eroded this year, and other is not. Now, suppose that the government announces that next year, the heavily eroded section of beach will be nourished. As a result, potential homebuyers believe that the two sections of beach will be approximately the same width next year. Because the houses located on these beaches are otherwise identical, and because homebuyers have identical expectations about the widths of the beaches in year 2 and onwards, any difference in the current-year sales price of the two houses must be attributable to the current difference in beach width. Thus, after controlling for consumers' beliefs about beach width next year, the relationship between current housing prices and current beach width has an exact interpretation as the marginal rental value of an extra year of improved beach width.

To prove a formal version of this argument, I use the structural housing model from the previous section to model the consequences of a policy intervention that causes a predictable future discontinuity in amenity quality. As before, at the start of period 1, amenity quality w_{i1} at each house is given. However, between periods 1 and 2, a policy intervention equalizes amenity quality across all houses (or at least, equalizes the probability distribution of amenity quality), regardless of each house's period 1 quality. Then, in periods 3 onward, the amenity at each house evolves as a Markov process.

A more formal statement of this assumption is as follows:

Assumption 2. Period 2 State Is Independent of Period 1 State: *In period 1, amenity quality w_{i1} at each house i is given. In period 2, amenity quality is determined by a policy intervention. Let $F_{i2}(w)$ be the post-intervention cumulative distribution function of amenity quality w_{i2} at house i at time $t = 2$. Then:*

$$F_{i2}(w|w_{i1}) = F_2(w) \quad \forall i, w$$

where $F_2(w)$ is a CDF that is shared by all houses in period 2 and is known by consumers when they purchase houses at the beginning of period 1. In periods 3 and onward, amenity quality at each house evolves independently according to the Markov process described in Assumption 1.

The key feature of Assumption 2 is that because of the policy intervention, period 1 amenity quality does not affect period 2 amenity quality. For example, for the coastal houses described in the previous section, Assumption 2 implies that all sections of beach are nourished to the same design width at the beginning of period 2, regardless their width in period 1. Then, in periods 3 onwards, each section of beach erodes and accretes independently according to a common Markov transition function. Thus, from the perspective of a consumer purchasing a house at the beginning of period 1, all houses have the same expected amenity quality in periods 2 onward.

Before proving the main result, I adopt one additional simplifying assumption: consumers' utility functions are quasilinear in the composite good c (or equivalently, quasilinear in income). Formally:

Assumption 3. Quasilinear Utility: Consumer j 's utility function is quasilinear in the composite good c :

$$u(w, \pi, c; \theta_j) \equiv v(w, \pi; \theta_j) + k_j c \quad (1.6)$$

where $v(\cdot; \theta_j)$ may take any functional form and k_j is a constant representing consumer j 's marginal utility of consumption for the composite good.

For notational convenience, let u_j represent the utility function for consumer j with preferences θ_j . I now state the main theoretical result:

Theorem 2. Discontinuity Matching Theorem: Suppose that in equilibrium, consumer j purchases house i . Under the conditions described in Assumptions 2 and 3, consumer j 's willingness to pay for a certain one-unit increase in amenity quality for period 1 only is given

by the derivative of equilibrium housing prices with respect to period 1 amenity quality:

$$\left. \frac{\partial p}{\partial w_1} \right|_{w_{i1}} = \frac{\left. \frac{\partial u_j}{\partial w} \right|_{w_{i1}}}{\frac{r}{\rho} \cdot \frac{\partial u_j}{\partial c}} \quad (1.7)$$

Proof. Theorem 1 states that:

$$\left. \frac{\partial p}{\partial w_1} \right|_{w_{i1}} = \frac{\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[\left. \frac{\partial u_j}{\partial w} \right|_{w_t} \frac{\partial w_t}{\partial w_1} \right]_{w_{i1}}}{\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[r \frac{\partial u_j}{\partial c} \right]_{w_t}} \quad (1.8)$$

By Assumption 3, marginal utility of consumption is constant, which implies that Equation (1.8) can be rewritten as:

$$\left. \frac{\partial p}{\partial w_1} \right|_{w_{i1}} = \frac{\sum_{t=1}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[\left. \frac{\partial u_j}{\partial w} \right|_{w_t} \frac{\partial w_t}{\partial w_1} \right]_{w_{i1}}}{\frac{r}{\rho} \cdot \frac{\partial u_j}{\partial c}} \quad (1.9)$$

Decomposing the right-hand side into current and future terms shows that:

$$\left. \frac{\partial p}{\partial w_1} \right|_{w_{i1}} = \frac{\left. \frac{\partial u_j}{\partial w} \right|_{w_{i1}}}{\frac{r}{\rho} \cdot \frac{\partial u_j}{\partial c}} + \sum_{t=2}^{\infty} \left(\frac{1}{1+\rho} \right)^t E \left[\left. \frac{\partial u_j}{\partial w} \right|_{w_t} \cdot \frac{\partial w_t}{\partial w_1} \right]_{w_{i1}} \quad (1.10)$$

This equation makes it clear that in equilibrium, the cost of purchasing an extra unit of the amenity depends on the (certain) marginal utility gained in period 1, plus the discounted expected marginal utility gained in future time periods.

Now note that Assumption 2 implies that $\frac{\partial w_t}{\partial w_1} = 0$ for all $t \geq 2$. In other words, changes in period 1 amenity quality have no effect on amenity quality in periods 2 onward. As a result, the future term in Equation (1.10) evaluates to zero and the equation simplifies to the result shown in Equation (1.7). \square

Theorem 2 is the key theoretical contribution of this chapter. It shows that marginal willingness to pay for a certain, immediate, one-unit, one-period increase in the amenity can be calculated directly from the cross-sectional relationship between housing prices and amenity quality in the year before a policy intervention occurs. This

fact is of great empirical importance. Because housing prices and amenity quality are directly observable, the left-hand side of Equation (1.7) can be used to estimate a parameter that has an exact interpretation as marginal willingness to pay in a theoretically-consistent hedonic model. This estimate of marginal willingness to pay for a temporary increase in amenity quality can then be scaled up, using the discount rate, to estimate marginal willingness to pay for a permanent change in amenity quality.

The logic underlying Theorem 2 rests on the assumption that the policy intervention causes all houses to have the same expected amenity quality in period 2, regardless of their period 1 quality. To achieve this equalization, between periods 1 and 2 some houses experience a large discontinuity in amenity quality, and some experience a smaller discontinuity. Thus, although pre-intervention amenity quality differs across houses, all houses have identical predicted post-discontinuity quality. It is this “discontinuity matching” on the basis of expected period 2 amenity quality that justifies interpreting price differences between houses as a measure of willingness to pay for marginal improvements in period 1 amenity quality.

Like many theoretical results, the practical value of Theorem 2 depends on whether or not its predicates—particularly Assumption 2—are true in relevant real-world situations. I argue that as long as it is possible to control for idiosyncratic time shocks and fixed housing characteristics, then Assumption 2 will be true in a number of practically useful situations. For example, in the empirical section of this chapter, I use discontinuity matching to estimate homeowners’ willingness to pay for a home located on a wider section of beach, using the sharp and highly-predictable change in beach width caused by nourishment projects. By matching sections of beach that are predicted to have similar widths next period, this approach allows a comparison of the current period price of houses located on narrow sections of beach (that receive nourishment next period) against the price of houses located on wide sections of beach (that are not nourished next period). However, the methodology could also be applied in many other contexts. For example, it could be used to estimate marginal willingness to pay

for construction of new subway or bus lines, based on the differential discontinuity in public transportation access caused by the opening of the new station or line.

1.4 Data

To support my analysis of the welfare costs of coastal shoreline loss, I have constructed a unique dataset on housing sales transactions, beach width, and nourishment projects, for properties located within 5 kilometers of the beach in sixteen coastal Florida counties. The dataset covers the period from 1983 to 2009, and includes 388 miles of Florida's coastline.

The sixteen counties included in the analysis are: Bay, Brevard, Broward, Charlotte, Duval, Escambia, Lee, Manatee, Martin, Miami-Dade, Palm Beach, Pinellas, Sarasota, St. Johns, St. Lucie, and Volusia. Figure 1.1 shows the a map of these counties, and indicates coastal areas where beach width survey data are available. As the map shows, these counties are primarily located on Florida's Atlantic coast and southern Gulf coast. These counties were chosen because they collectively cover the majority of segments of shoreline designated as "critically-eroded" by the Florida Department of Environmental Protection (FDEP, 2011b).

Below, I describe each of the data sources in more detail.

1.4.1 Beach Width Data

I have compiled data on beach width from a database of coastal surveys maintained by the Florida Department of Enviromental Protection (FDEP, 2008a). Each record in the database represents the distance from a fixed coastal survey monument to the mean high water (MHW) mark along a particular segment of beach beach.⁴ The survey monuments are spaced approximately 1,000 feet apart along much of the Florida coastline;

⁴The mean high water mark represents the location on the beach reached by the average high tide.

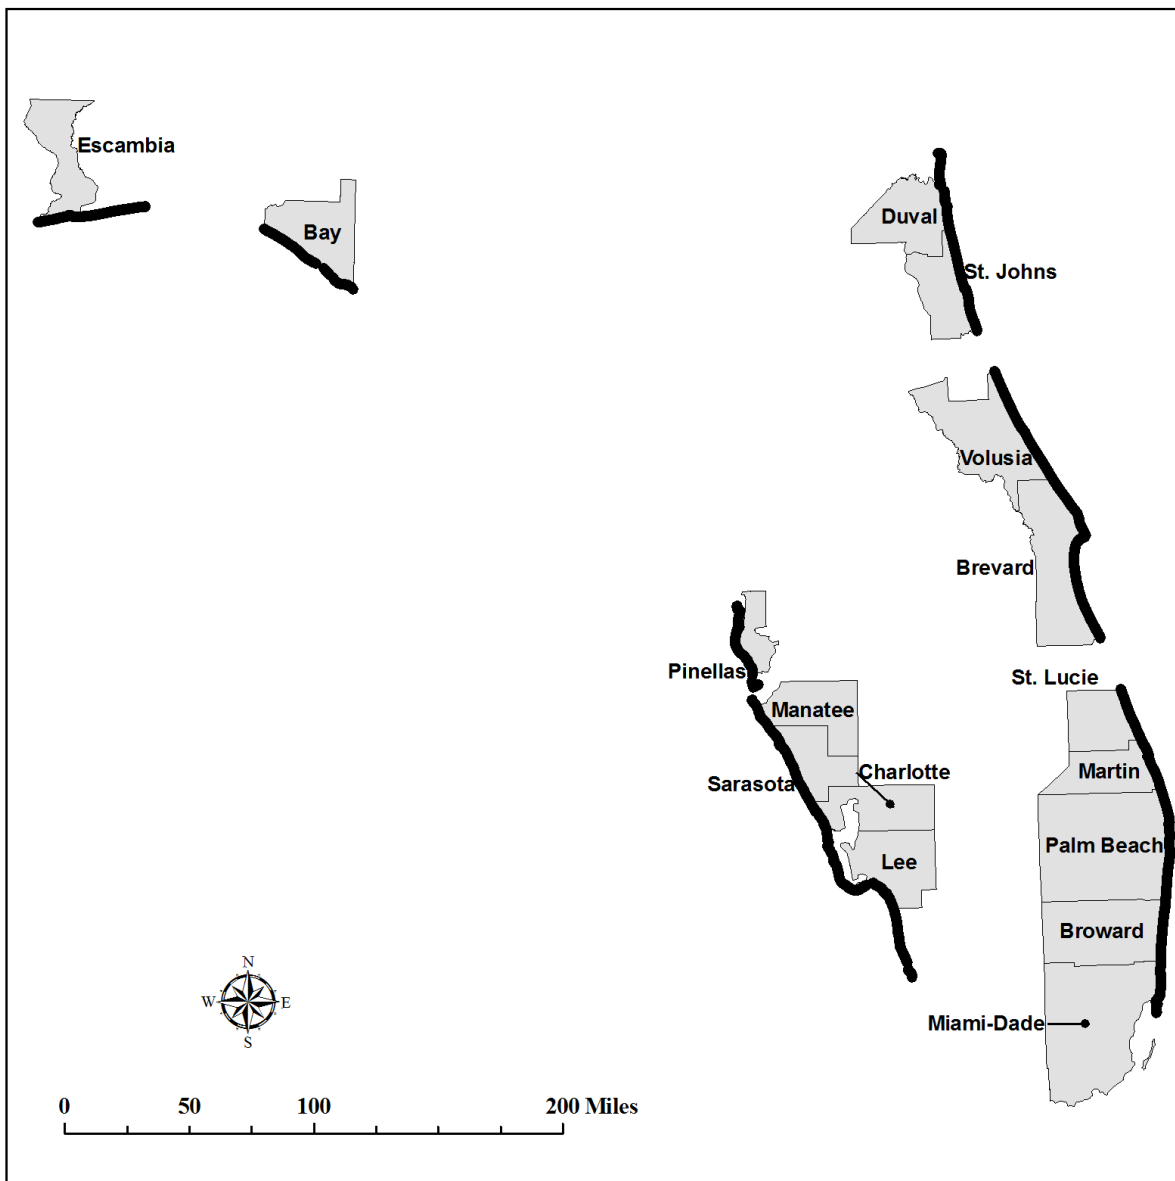


Figure 1.1: Map of the Study Area

Note: The thick solid line denotes areas where coastal mean high water (MHW) line survey data is available from the Florida Department of Environmental Protection. Grey shaded areas represent the sixteen in-sample individual counties.

thus, these data have a high level of geographic resolution. In most locations, surveys are available every few years between 1983 and 2009, with better coverage in recent years.

Because the survey monuments are not located at a consistent distance from the upper end of the beach, the MHW data do not represent absolute beach width. To calculate absolute beach width, I use georeferenced aerial photography data (FDEP, 2011a) to manually geocode the location of the upper end of the beach near each survey monument. For consistency in determining the upper end of the beach, I use the location of the most seaward manmade structure near each survey monument. I then calculate beach width as the distance from the upper end of the beach to the survey monument, plus the distance from the survey monument to the MHW mark, with a trigonometric adjustment to account for the angle at which the MHW survey measurements were taken.

1.4.2 Property Sales Data

I have collected data on approximately 1.1 million qualified housing sales transactions in the sixteen in-sample counties. These data are taken from electronic records maintained by the property appraisers' offices in each county. The data include sales information, such as sale date, sale price, and sale qualification, as well as property characteristics, including building type, acreage, and construction date.

I use two methods to geocode the location of each property. When possible, I link each property's parcel ID number to detailed GIS parcels maps obtained from the Florida Department of Revenue (FDEP, 2011). In cases when there is no match, I geocode the property's street address using ESRI's ArcMap Business Analyst, using address point data when available and otherwise using an offset of 70 feet from the address location along the street centerline. I then link each geocoded address to the nearest GIS parcel, and use that parcel for subsequent analysis.

For each parcel, I then calculate the distance to the coast. I also assign each parcel-by-year observation a beach width measurement, based on beach width at the nearest FDEP survey monument.

1.4.3 Nourishment Data

I have constructed a list of beach nourishment projects using data from several sources. My primary source is a database compiled by researchers at Western Carolina University (WCU, 2008). This database lists 361 beach nourishment projects that took place in Florida between 1944 and 2006. I supplement this data with information from Florida State University's beach erosion control database, a list of ongoing nourishment projects maintained by the Florida Department of Environmental Protection, and a variety of other sources (e.g., FSU, 2008b; FDEP, 2009). After merging these data and eliminating duplicate and out-of-sample records, I generate a dataset that represents the near-universe of nourishment projects that took place between 1983 and 2009 in the 16 in-sample counties.

Table 1.1 provides detailed information about a subset of the largest nourishment projects included in the analysis. The table shows the set of variables that I have been able to collect for each nourishment project, including the year the project was completed, the project location, the cost of the project, the volume of sand deposited, and the length of beach nourished. Unfortunately, not all variables—particularly costs—are available for all nourishment projects.

Table 1.2 describes the characteristics of the 204 projects included in my analysis. The table shows that the average project placed 700,000 cubic yards of sand (roughly 900,000 tons) onto a 3.8 mile segment of beach, for an average nourishment intensity of approximately 47 cy/ft. The average cost of a nourishment project was \$6.6 million. The table also summarizes the characteristics of a set of “major” projects with nourishment intensities of at least 25 cy/ft. As expected, these projects had higher volumes,

Table 1.1: Examples of Major Beach Nourishment Projects

County	Beach	Year	Volume (cy, millions)	Length (miles)	Intensity (cy / ft)	Cost (millions)
Bay	Panama City Beach	1999	9.1	18.3	98	41.9
Brevard	Cape Canaveral-Cocoa Beach	2001	3.1	9.4	64	26.8
Brevard	South Reach Beach	2002	1.2	3.1	67	.
Broward	Broward County Segment III	2006	1.5	9.3	40	3.1
Broward	Hollywood-Hallandale	1991	1.1	5.1	40	14.3
Broward	Pompano Beach-Lauderdale	1983	1.9	6.0	66	15.0
Dade	Sunny Isles	1988	1.3	2.6	88	27.9
Duval	Amelia Island	2002	1.9	5.1	96	9.4
Duval	Mayport-Kathryn Abby Hanna	1985	1.3	2.2	102	4.0
Escambia	Pensacola Beach	2003	4.2	8.2	96	23.0
Escambia	Pensacola Beach	2006	3.5	8.2	79	10.4
Escambia	Perdido Key	1985	2.4	1.2	520	9.9
Escambia	Perdido Key	1990	5.4	5.1	211	20.3
Lee	Captiva Island	1989	1.6	5.5	61	10.9
Lee	Captiva Island	1996	1.1	8.8	34	.
Lee	Captiva Island	2006	1.4	9.8	39	.
Manatee	Anna Maria Key	1993	2.3	4.6	92	19.3
Manatee	Anna Maria Key	2002	1.9	5.3	65	9.8
Martin	Jupiter Island	1983	1.0	7.8	26	5.1
Martin	Jupiter Island	1987	2.2	7.8	57	6.5
Martin	Jupiter Island	1996	1.7	5.5	64	9.7
Martin	Jupiter Island	2002	3.0	7.7	77	1.8
Martin	Jupiter Island	2007	2.3	7.3	62	.
Palm Beach	Boca Raton North	1988	1.1	1.5	129	6.3
Palm Beach	Delray Beach	1984	1.3	2.9	80	8.0
Palm Beach	Delray Beach	1992	1.2	1.6	123	7.3
Palm Beach	Delray Beach	2002	1.1	1.6	118	4.6
Palm Beach	Juno Beach	2001	1.0	3.4	72	13.2
Palm Beach	Midtown Beach	2003	1.4	10.2	93	.
Palm Beach	Palm Beach	1984	1.3	15.3	33	.
Palm Beach	Palm Beach	1992	1.2	15.1	30	6.0
Palm Beach	Phipps Ocean Park	2006	1.0	1.6	121	.
Pinellas	Sand Key	2006	1.3	6.9	34	23.4
Pinellas	Sand Key Phase II	1990	1.3	2.6	87	2.4
Pinellas	Sand Key Phase IV	1999	1.6	6.7	43	22.9
Sarasota	Longboat Key - multicounty	1993	3.1	12.8	63	27.3
St Johns	Anastasia - St Augustine	2003	4.4	4.3	217	.
St Johns	Anastasia - St Augustine	2005	2.8	2.9	185	14.1
St Johns	Duval County	1995	1.2	5.9	37	.
St Lucie	Martin County	1996	1.3	4.3	59	11.6

Note: This table presents basic characteristics for the 40 largest (by volume of sand placed on the beach) of the 204 beach nourishment projects included in the main analysis. Note that length is estimated from the length of shoreline covered by the FDEP survey markers affected by the nourishment project.

Table 1.2: Characteristics of Beach Nourishment Projects

	Mean	Std Dev	Minimum	Maximum	Observations
All Nourishments					
Year Completed	1996	7	1983	2007	204
Volume (cy, millions)	0.7	1.0	0.0	9.1	204
Length (miles)	3.8	3.8	0.2	19.4	204
Intensity (cy/ft)	47	51	0	520	204
Cost (millions)	6.6	8.1	0.0	41.9	127
Sales	62	130	0	1,473	204
Major Nourishments					
Year Completed	1997	7	1983	2007	94
Volume (cy, millions)	1.2	1.3	0.2	9.1	94
Length (miles)	4.2	3.5	1.0	18.3	94
Intensity (cy/ft)	75	60	25	520	94
Cost (millions)	9.8	8.6	0.9	41.9	68
Sales	72	172	0	1,473	94

Note: This table summarizes the characteristics of the set of beach nourishment projects included in the main analysis. The top panel includes all projects; the bottom panel includes only “major” nourishment projects covering 1 mile or more of beach, with nourishment intensity greater than or equal to 25 cubic yards per foot. In both panels, the sales variable represents the number of sales that occurred at properties within 20 meters of a nourished section of beach, in the year in which the beach was nourished. The observations variable represents the number of nourishment projects for which the selected variable is available.

covered more shoreline, and had higher costs.

1.4.4 Summary Statistics

Table 1.3 summarizes basic characteristics of the properties and beaches included in the main analysis, for the subset of properties located within 20 meters of a beach. The results are disaggregated into three categories based on the number of nourishments that took place at each FDEP marker between 1980 and 2010. The top two panels shows that housing and beach characteristics in nourished areas differ significantly from housing and beach characteristics in areas that are never nourished.

The third panel in Table 1.3 describes the number of sales that occur in each five year period between 1980 and 2010, for properties located within 20 meters of the beach. The panel shows that there are 41,187 sales at FDEP survey monuments that are never

Table 1.3: Summary Statistics

	Number of Major Nourishments per Monument		
	0	1 or 2	3 or more
Housing Characteristics			
Single Family	0.10 (0.00)	0.03*** (0.00)	0.09 (0.00)
Condo	0.90 (0.00)	0.97*** (0.00)	0.91 (0.00)
Vacant	0.18 (0.00)	0.02*** (0.00)	0.03*** (0.00)
Parcel Acreage	0.59 (0.01)	0.29*** (0.01)	1.35*** (0.02)
Housing Area (sq ft)	1,593 (12)	1,369*** (6)	1,701*** (24)
Bedrooms	1.94 (0.01)	1.62*** (0.01)	2.68*** (0.04)
Bathrooms	1.99 (0.01)	1.78*** (0.01)	2.45*** (0.04)
Year Renovated	1983.1 (0.1)	1982.0*** (0.1)	1999.2*** (0.2)
Year Built	1980.4 (0.1)	1978.3*** (0.1)	1981.4*** (0.2)
Structure Quality (1-6)	3.03 (0.01)	3.20*** (0.01)	3.21*** (0.02)
Brick Construction	0.12 (0.00)	0.09*** (0.00)	0.10*** (0.00)
Features Appraised Value	2,327 (219)	719*** (46)	5,209*** (402)
Distance to Beach (m)	1 (0)	1*** (0)	1*** (0)
Sales per Parcel	2.59 (0.01)	2.63*** (0.01)	2.77*** (0.01)
Sale Price (000s)	987 (15)	1,059*** (14)	538*** (15)
Beach Characteristics			
Beach Width (ft)	195.3 (1.0)	235.1*** (0.7)	197.0 (0.9)
Std Dev Beach Width (ft)	33.11 (0.36)	42.31*** (0.13)	56.84*** (0.41)
Nourishments	0.00 (0.00)	0.98*** (0.00)	2.69*** (0.01)
Cumulative Intensity (cy/ft)	0.00 (0.00)	56.58*** (0.29)	143.18*** (0.82)
Sales, by Time Period			
1983 to 1984	1,968	2,580	759
1985 to 1989	6,456	7,719	1,966
1990 to 1994	6,314	10,733	1,998
1995 to 1999	8,529	13,403	2,515
2000 to 2004	11,495	18,730	2,888
2005 to 2009	6,425	12,592	1,677
All years	41,187	65,757	11,803
Geographic Units			
Parcels	15,888	24,998	4,259
FDEP Survey Monuments	679	454	171
Six mile zones	51	45	18

Note: The “Housing Characteristics” panel presents mean unweighted parcel characteristics, with standard deviations in parentheses. The “Beach Characteristics” panel presents mean unweighted FDEP survey monument characteristics, with standard deviations in parentheses. The “Sales, by Time Period” panel presents the number of sales that occurred in each five-year period between 1980 and 2010. The “Geographic Units” panel describes the number of geographic units included in the analysis. The columns represent the number of major nourishments (with intensity >25 cy/ft) at each FDEP survey monument. All t-tests are relative to the 0 nourishments group. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

nourished, 65,757 sales at monuments with one or two nourishments, and 11,803 sales at monuments with three or more nourishments. The number of sales increases over most of the sample, reaching a peak in 2000-2004 (at the height of the housing bubble) and then falling off in more recent years.

The fourth panel in Table 1.3 shows the number of geographic units represented by properties within 20 meters of the beach. There are 45,145 such properties, located near 1,304 FDEP survey monuments in 67 six-mile-long “zones”. I have defined these zones as contiguous sections of coastline that span six miles each. Because the beach width data show substantial spatial correlation between adjacent FDEP monuments (which are only 1,000 feet apart on average), I cluster and weight all regression results in this chapter at the level of these zones.

An important constraint on my analysis is the availability of survey data on beach width. Figure 1.2 shows the availability of beach width data, by year and location along the coastline. As the figure indicates, data on beach width is available in only about a third of the survey monument-by-year combinations. This missing data is serious concern for my panel and discontinuity matching regressions. I deal with this problem in two ways. In the repeat sales panel regressions, I adopt the assumption that the data are missing at random after controlling for property and time fixed effects. In the discontinuity matching approach, I use imputed width data.⁵ However, note that because the differences-in-differences regressions do not require data on beach width, they are not affected by the missing data issue. The availability of this unbiased approach alleviates concerns about data problems in the other two approaches.

Figure 1.2 also shows the location and timing of beach nourishment projects. Several patterns are evident from the figure. First, many segments of beach are nourished at somewhat regular intervals, ranging from three to ten years. Second, nourishments

⁵I impute missing beach width data using an expectation-maximization algorithm that employs a Kalman filter followed by smoothing. Details about the imputation procedure are available upon request.

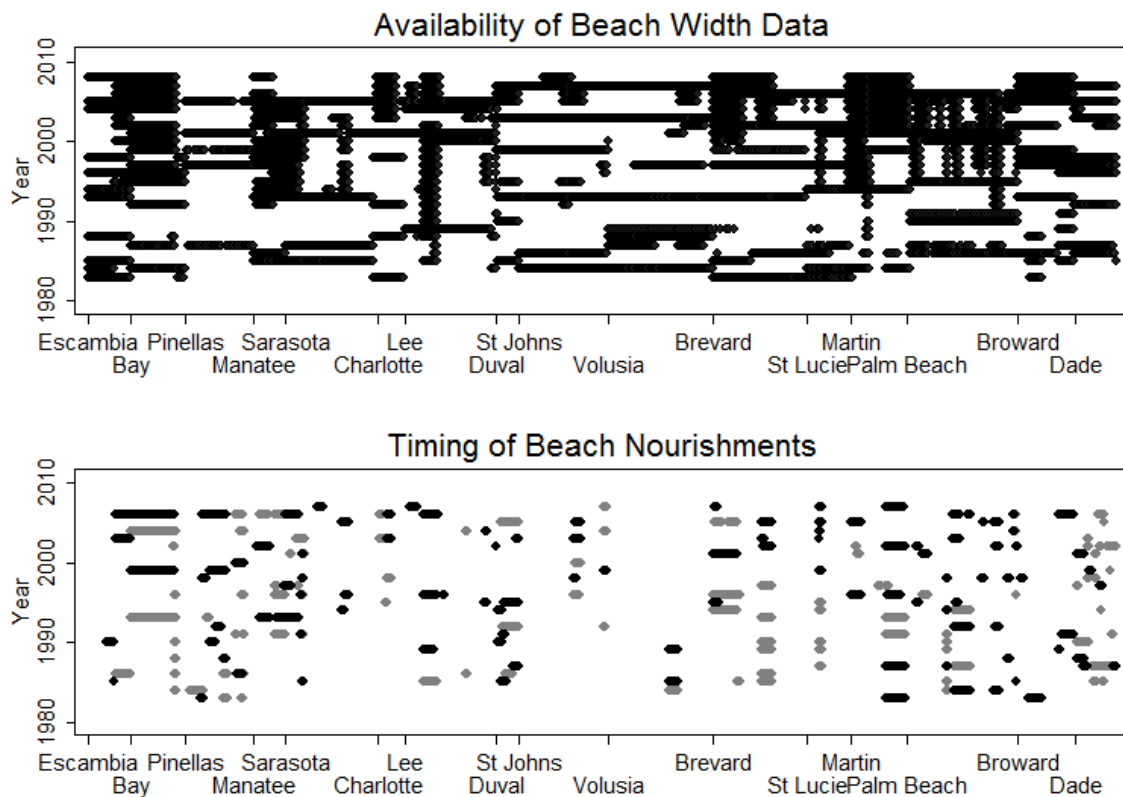


Figure 1.2: Location and Timing of Surveys and Nourishment Projects

Note: The upper panel shows the availability of beach width survey data, by year and location along the coast. The lower panel shows the timing of beach nourishment projects, by year and location. Major beach nourishment projects are shown in black; minor projects are shown in grey. Each point on the x-axis represents a unique coastal survey monument (in most areas, monuments are spaced approximately 1000 feet apart along the coast). To construct this figure, monuments were ordered from north to south along Florida's Gulf coast, and then from north to south along Florida's Atlantic coast. The x-axis should not be interpreted as a measure of absolute distance, only of relative location. As shown in Figure 1.1, survey coverage of Florida's coastline is incomplete.

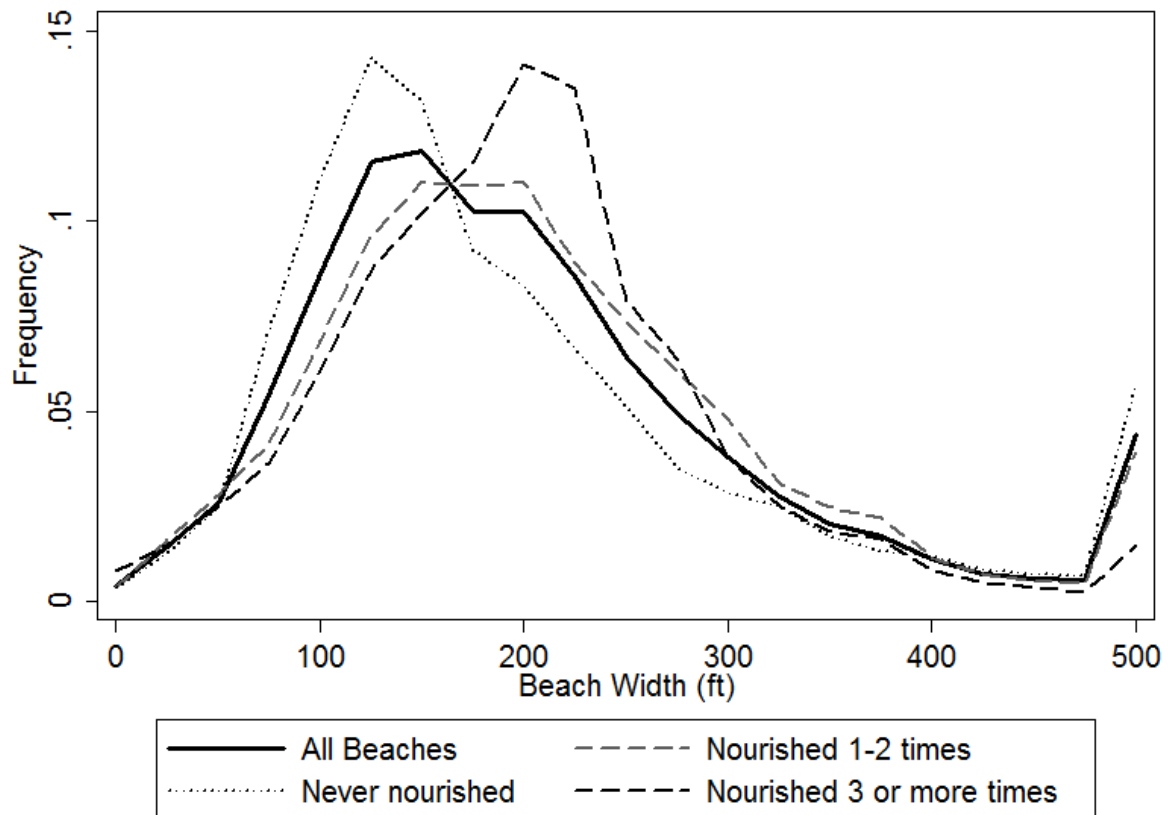


Figure 1.3: Distribution of Beach Width, by Nourishment Frequency

Note: This figure shows histograms of beach width, for all beaches, and by nourishment frequency. Beach widths greater than 500 ft are shown at 500 ft. Each observation represents a unique survey monument and year. The figure includes data from 1983 to 2009 for all in-sample counties.

only occur at certain segments of coastline.

Finally, Figure 1.3 shows a histogram of the distribution of beach width across all monument-by-year observations. The figure shows that the modal beach width is approximately 140 feet, and that the distribution of beach width has a long right tail. As the figure indicates, most beaches are between 25 and 400 feet wide.

1.5 Econometric Approach

In this chapter I take three distinct econometric approaches to estimating homeowners' willingness to pay for wider beaches: (1) a repeat sales panel approach; (2) a differences-in-differences approach; and (3) a new “discontinuity matching” approach. In the following subsections, I explain how each research design uses housing market data to estimate $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$, the partial derivative of the housing price function at time t with respect to beach width at time t . Under the assumption that the structural model described in Section 3 captures the main features of housing markets in coastal Florida, I use Theorems 1 and 2 to provide direct welfare interpretations for the estimated coefficients.

1.5.1 Repeat Sales Panel Approach

I begin by using a repeat sales panel approach to develop an estimator for $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$. By Theorem 1 from Section 3, this partial derivative is equal to the marginal consumer's willingness to pay for a one-time policy intervention that would add one foot of sand to the beach in front of house i at time t , and then allow beach width to evolve freely in future periods. To motivate the estimator for this derivative, I adopt the following notation.

Let each beachfront house i belong to a neighborhood n , where n represents closest FDEP survey monument (so that the average neighborhood includes roughly 1000 feet of shoreline). Let p_{nit} represent the sale price of house i in neighborhood n in year t , and let w_{nt} describe beach width in neighborhood n in year t . As before, π_i represents unobservable permanent characteristics of house i . Additionally, let τ_t capture aggregate housing market shocks in year t , and let ϵ_{nit} be an normally distributed, zero-mean error term, with variance σ^2 , that captures other sources of variation in price.

Following common practice in the hedonics literature, I assume that the price function

$p_{nit} \equiv p(w_{nt}, \tau_t, \pi_i, \epsilon_{nit})$ takes a log-linear functional form. However, to allow for more flexibility in the relationship between price and beach width, I use a semi-parametric specification for w_{nt} . Let the binary variables $w_{nt}^0, w_{nt}^{50}, \dots, w_{nt}^j, \dots, w_{nt}^{500}$ indicate whether beach width in neighborhood n in year t is between j feet and $j+50$ feet. I then assume that the price function can be written as:

$$\log p_{nit} = \sum_j \beta_j w_{nt}^j + \tau_t + \pi_i + \epsilon_{nit} \quad (1.11)$$

Equation (1.11) suggests calculating marginal willingness to pay for a policy intervention using the following approximations to the partial derivative of price with respect to width:

$$\begin{aligned} E \left[\frac{\partial p}{\partial w_t} \Big|_{w_t=w^j} \right] &\approx E \left[\frac{p(w^J, \tau_t, \pi_i, \epsilon_{nit}) - p(w^0, \tau_t, \pi_i, \epsilon_{nit})}{w^J - w^0} \right] \\ &= \frac{(e^{\beta_J} - e^{\beta_0}) e^{\tau_t + \pi_i + (\sigma^2/2)}}{w^J - w^0} \end{aligned} \quad (1.12)$$

where with some mild abuse of notation, I let $w^j \equiv j$ indicate the lower end of the j th beach width bin.

To estimate willingness to pay using Equation (1.12), I first use panel data on repeat housing sales to estimate Equation (1.11) using ordinary least squares. The identifying assumption is that conditional on the year and house fixed effects, variation in beach width is orthogonal to any other housing price determinants captured in the error term. Because beach width at adjacent survey monuments is highly correlated, I cluster standard errors by six-mile-long sections of beach. Additionally, to improve the efficiency of the estimates, I weight each observation by the sum of the inverse of the total number of housing sales in its neighborhood. Then, in a second step, I substitute the estimated coefficients from Equation (1.11) into Equation (1.12) and calculate the marginal willingness to pay estimate $\widehat{\frac{\partial p}{\partial w_t}} \Big|_{w_t}$.

1.5.2 Differences-in-Differences Approach

One potential concern about the repeat-sales approach is that within particular neighborhoods, long-term changes in beach width may be correlated with changes in long-term determinants of housing prices. For example, as neighborhoods become wealthier, they may make more frequent investments in beach nourishment. In the panel model from the previous section, this would create correlation between the beach width bin variables and the error term.

To address this possibility, I implement a differences-in-differences regression approach based on abrupt changes in beach width caused by nourishment projects. Again, the goal is to estimate the derivative $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$. The advantage of a differences-in-differences approach is that the factors that influence the decision to nourish a particular stretch of beach—such as neighborhood wealth and political influence—are likely to evolve slowly over time. Thus, after inclusion of appropriate fixed effects, the sudden increase in beach width caused by beach nourishment is likely to be orthogonal to other determinants of housing prices.

Unfortunately, the disadvantage of a differences-in-differences approach is that it fits less well into the structural model that provides a welfare interpretation for the derivative $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$. Theorem 1 depends on the assumption that changes in beach width can be modeled as a Markov process. To argue that beach nourishment fits this assumption requires that (1) homebuyers have no advance knowledge of beach nourishment projects, and (2) the decision to nourish beaches is random, conditional on beach width in the previous period. Neither assumption is a perfect description of reality. However, for the purposes of this section, I proceed as if both assumptions hold. Then, in the following section, I describe a discontinuity matching design that is better able to estimate the welfare effects of deliberate policy interventions.

I implement the differences-in-differences approach in two steps. First, I divide the coastline into one-mile “neighborhoods” and then use houses with repeat sales to de-

velop a housing price index for each of these neighborhoods. Second, I run differences-in-differences regressions on a dataset that includes five years of forward and backward lagged data for each neighborhood-by-year observation.

The following subsections describe these two steps in more detail.

Step 1: Neighborhood Housing Price Indices

Before running the differences-in-differences regressions, I estimate a separate housing price index for each one-mile beachfront neighborhood. Developing price indices serves two purposes. First, it removes variation in prices that can be attributed to fixed idiosyncratic characteristics of individual houses. Second, it transforms an unbalanced panel of housing sales transactions, which contains only two or three sales per house, into a balanced panel of price indices with an observation for almost every year and neighborhood combination. Because there is only modest variation in beach width across individual FDEP survey monuments within one-mile sections of coastline, this transformation sacrifices little information about the relationship between housing prices and beach width. Using a price index rather raw housing sales data is necessary because I include only five years of forward and backward lagged data for each neighborhood-by-year observation.

To estimate the price indices, I model the price function $p_{nit} \equiv p(P_{nt}, \pi_i, \epsilon_{nit})$ as depending on three sets of parameters: fixed effects P_{nt} that capture all sources of neighborhood-by-year variation in housing prices, house fixed effects π_i that reflect fixed characteristics of house i , and a zero-mean error term ϵ_{nit} . The fixed effects P_{nt} represent the price index for neighborhood n in year t . I assume that these parameters enter the price function according to the following log-linear functional form:

$$\log p_{nit} = \pi_i + \tau_{nt} + \epsilon_{nit} \quad (1.13)$$

where I define $\tau_{nt} \equiv \log(P_{nt})$. I estimate this equation using OLS, and then use the estimated coefficients to calculate the price index $P_{nt} = \exp(\tau_{nt})$ for each neighborhood

and year unit.

It is important to note that any subsequent regression of this price index on other housing and neighborhood variables will be consistent only if: (i) the regression includes neighborhood fixed effects, and (ii) none of the independent variables vary at the level of individual houses.

Step 2: Differences-in-Differences Regressions

The differences-in-differences research design uses changes in housing prices at beaches that are not nourished to generate a counterfactual for changes in housing prices at beaches that are nourished. This comparison depends on the identifying assumption that in the absence of nourishment, the trend in prices at nourished beaches would have been similar to the trend in prices at unnourished beaches.

To run differences-in-differences regressions, I construct a dataset that includes five years of pre-data and four years of post-data for each neighborhood-by-year observation. Let t represent the base year for each set of observations, and let l represent time elapsed since the base year. Thus, each neighborhood-by-year observation appears ten times, as $(n, t, l) = (n, t + 5, -5), \dots, (n, t - 1, -1), (n, t, 0), \dots, (n, t - 4, 4)$. I then assume that the housing price index function $P(N_{nt}, \tau_{tl}, \pi_{nt})$ can be modeled as:

$$P_{ntl} = \sum_{s=-5}^4 \beta_s N_{nt} \cdot \mathbf{1}\{l = s\} + \tau_{tl} + \pi_{nt} + \epsilon_{ntl} \quad (1.14)$$

In this equation, N_{nt} is a dummy variable that indicates that the beach near neighborhood n was nourished in the base year t . The indicator $\mathbf{1}\{l = s\}$ takes value 1 if $l = s$, and zero otherwise. The year-by-elapsed year fixed effects τ_{tl} capture aggregate housing market shocks, and the neighborhood-by-base year fixed effects π_{nt} capture fixed neighborhood characteristics of neighborhood n for the set of ten elapsed observations with common base year t . Because of the possibility that adjacent neighborhoods do not represent independent observations, I cluster standard errors by six-mile-long sec-

tions of beach.

I also develop similar a differences-in-differences model of beach width. Using the same dataset of pre- and post-data, I model the beach width $w_{ntl} \equiv w(N_{nt}, \lambda_{tl}, \kappa_{nt})$ as follows:

$$w_{ntl} = \sum_{s=-5}^4 \phi_s N_{nt} \cdot \mathbf{1}\{l = s\} + \lambda_{tl} + \kappa_{nt} + \zeta_{ntl} \quad (1.15)$$

In this equation, w_{ntl} represents beach width in neighborhood n in year t , lagged l years. The binary variable N_{nt} indicates that the beach in neighborhood n was nourished in year t , and the indicator $\mathbf{1}\{l = s\}$ takes value 1 if $l = s$, and zero otherwise. The base year-by-elapsed year fixed effects λ_{tl} capture aggregate shocks to beach width, and the neighborhood-by-base year fixed effects κ_{nt} capture fixed differences in beach width between neighborhoods, for sets of elapsed observations with the same base year. The zero-mean error term ζ_{ntl} captures other sources of variation in beach width.

The rationale for constructing this “stacked” dataset is that because beach width approximately follows a random walk, the fixed-effects approach embedded in a differences-in-differences design is not really an appropriate model. However, by limiting each fixed effect to include only ten years of data, I minimize inaccuracies by modeling changes over only a short period of time. Note that because I estimate separate coefficients for each forward and backward lag of beach nourishment, and because I cluster standard errors by neighborhood and base year, duplicating observations in this way does not raise concerns about overestimating the precision of the coefficients.

To use the differences-in-differences results to calculate the derivative $\frac{\partial p_t}{\partial w_t} \big|_{w_{it}}$, I first estimate Equations (1.14) and (1.15) using OLS. I then approximate the average derivative of price with respect to beach width as:

$$\begin{aligned} E \left[\frac{\partial p}{\partial w_t} \big|_{w_t} \right] &\approx E \left[\frac{P(N_{nt} = 1, \tau_{t0}, \pi_{nt}) - P(N_{nt} = 0, \tau_{t0}, \pi_{nt})}{w(N_{nt} = 1, \lambda_{t0}, \kappa_{nt}) - w(N_{nt} = 0, \lambda_{t0}, \kappa_{nt})} \right] \\ &= \frac{\beta_0}{\phi_0} \end{aligned} \quad (1.16)$$

Under the assumptions discussed above and in Theorem 1, this derivative represents

willingness to pay for a one-time beach nourishment project that adds one foot of sand to the beach in the current year.

In addition to the main differences-in-differences specification discussed above, I also run alternative regressions in which I break down the impacts of beach nourishment based on the intensity of each project (measured in cubic yards of sand added per foot of coastline). I divide nourishments into four categories: 1 to 24 cy/ft, 25 to 50 cy/ft, 50 to 74 cy/ft, and 75 or more cy/ft. I then run differences-in-differences regressions for price and width that include separate sets of nourishment variables for nourishments in each of these categories. Using these coefficients, I then generate several alternative estimates of the derivative $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$.

1.5.3 Discontinuity Matching Approach

The repeat sales and differences-in-differences designs from the previous sections both have drawbacks. Although Theorem 1 provides a strong theoretical foundation for the repeat sales design, this approach is vulnerable to omitted variables that change over time within neighborhoods. In contrast, although the differences-in-differences design is a more robust empirical approach, using Theorem (1) to give a welfare interpretation to coefficients from differences-in-differences regressions requires somewhat unrealistic assumptions.

To overcome the limitations of these two approaches, in this section I describe a new “discontinuity matching” research design, motivated by Theorem 2, that uses predictable discontinuities in beach width to identify marginal willingness to pay. The goal of this procedure is to develop an estimate of the derivative $\frac{\partial p_t}{\partial w_t} \Big|_{w_{it}}$ from Theorem 2. Recall that under the conditions specified in Theorem 2, this derivative expresses the marginal consumer’s willingness to pay for an immediate, one-period, marginal increase in beach width. For this result to hold, prospective homebuyers must be aware that next period, the government will implement a beach nourishment project that will

cause a discontinuity in beach width along some sections of coastline. I argue that this condition is a reasonable approximation of reality.

Implementing the discontinuity matching research design requires three steps. First, I divide beachfront houses into mile-long neighborhoods, and then estimate a housing price index for each neighborhood-by-year cell. This step removes cross-sectional variation in price that can be attributed to fixed characteristics of individual houses. Second, I develop a simple rational model of homebuyers' beliefs about the effect of nourishment projects on beach width, and then estimate the parameters of this model using historical data on beach width. Third, I use a nearest-neighbor matching procedure to identify sets of neighborhood-by-year units that are predicted to have similar beach width in the following year (based on the empirical belief model). Within each matched set, all neighborhoods share the same next-period predicted beach width. However, some neighborhoods reach that next-period beach width by being nourished; some reach it without nourishment. Thus, I use the matching procedure to estimate the treatment effect of *next-period* nourishment on the *current* width and price index of these two sets of neighborhoods, and then use these treatment effects to calculate the pre-intervention cross-sectional derivative between price and beach width. Under Theorem 2, this derivative can be interpreted as the marginal rental value of widening the beach by one foot for one year.

I now discuss these three steps in more detail.

Step 1: Estimate Neighborhood Housing Price Indices

I begin the discontinuity matching procedure by estimating a separate housing price index for each one-mile-long beachfront neighborhood. As in the difference-in-differences approach, developing price indices removes variation in prices that can be attributed to fixed characteristics of individual houses, and transforms an unbalanced panel of housing sales into a balanced panel of neighborhood price indices. Because

there is relatively little variation in beach width within these one-mile neighborhoods, this transformation sacrifices little information about the relationship between housing prices and beach width.

To estimate the neighborhood price indices, I follow the procedure described previously in Section 1.5.2.

Step 2: Model Homebuyers' Beliefs about Beach Width

The second step of the discontinuity matching procedure is to develop and estimate a model of homebuyers' beliefs about the evolution of beach width over time. Although there are a variety of possible belief structures, I assume a simple model in which a consumer who buys a house knows current beach width, as well as whether the beach will be nourished the following year, and makes rational predictions based on this information.⁶ More formally, let a consumer's information set at time t include current width w_{nt} and a next-period nourishment indicator $N_{n,t+1}$, for every neighborhood n . The consumer believes that the evolution of beach width over time follows an AR(1) process:

$$w_{nt} = \kappa + \alpha w_{n,t-1} + \phi N_{nt} + \zeta_{nt} \quad (1.17)$$

in which current beach width w_{nt} depends on a constant term κ , previous-year beach width $w_{n,t-1}$, a current nourishment indicator N_{nt} , and a normally-distributed zero-mean error term ζ_{nt} . I assume the consumer's information set includes the parameters κ , α , and ϕ that govern this system.

Equation (1.17) is a simplistic model of the true transition function for beach width. Sand erosion and accretion are complex processes that are governed by complicated longshore sediment transport equations that include variables such as sand grain size, shore profile, and prevailing currents (Van Wellen, Chadwick, and Mason, 2000). How-

⁶Poor et al (2001) show that objective measures of environmental quality perform as well as subjective measures in explain home prices.

ever, because consumers are unlikely to have a sophisticated understanding of the factors that influence beach width, Equation (1.17) may nonetheless represent an appropriate model of consumers' beliefs. Furthermore, this reduced form model is consistent with the necessary conditions stated in Theorem 2—in particular, that beach width may be affected by a policy intervention but otherwise follows a Markov process.

To generate an estimate of consumers' beliefs about next-period beach width for every neighborhood-by-year observation, I estimate Equation (1.17) using OLS, using panel data on beach width and the timing and location of nourishment projects. I then use the estimated coefficients to predict beach width in the following period ($\hat{w}_{n,t+1}$), as follows:

$$\hat{w}_{n,t+1} \equiv \kappa + \alpha w_{nt} + \phi N_{n,t+1} \quad (1.18)$$

Step 3: Estimate Willingness To Pay Using Nearest Neighbor Matching

The goal of the discontinuity matching procedure is to develop an estimate of the cross-sectional partial derivative of current housing prices with respect to current beach width, for matched neighborhoods with beaches that are predicted to have the same beach width in the following period. Under Theorem 2, this derivative has an interpretation as marginal willingness to pay for an extra year of improved beach width. To explain the nearest neighbor matching procedure that I use to estimate this derivative, I adopt the Rubin potential outcomes framework (Rubin, 1974; Rosenbaum and Rubin, 1983).

Suppose that at the beginning of year t , consumers believe that the beach near neighborhood n will be $\hat{w}_{n,t+1}$ feet wide in year $t+1$. Based on the model of beliefs presented in Equation (1.17), there are only two ways that this could happen: either the beach is nourished at the beginning of year $t+1$, or it is not. Let $N_{n,t+1}$ be a binary treatment variable that takes value 1 if beach nourishment occurs at neighborhood n in year $t+1$, and 0 otherwise. Now, let $w_{nt}(1)$ and $w_{nt}(0)$ denote the two potential outcomes for

beach width at beach n in year t , depending on the value of the treatment $N_{n,t+1}$. For example, $w_{nt}(1)$ denotes beach width in neighborhood n and year t when the beach in this neighborhood is nourished in year $t + 1$ (i.e., $N_{n,t+1} = 1$). Similarly, let $P_{nt}(1)$ and $P_{nt}(0)$ be the two potential outcomes for the housing price index in neighborhood n in year t , where again the treatment $N_{n,t+1} \in \{0, 1\}$ indicates whether nourishment occurs in neighborhood n in year $t + 1$.

Casting the problem in this potential outcomes terminology is counterintuitive, given that the actual observed outcomes w_{nt} and P_{nt} are determined *before* the beach nourishment project takes place in year $t + 1$. Furthermore, given that next-period nourishment is more likely at beaches that are currently eroded, it would appear that the treatment $N_{n,t+1}$ is a function of the observed outcome w_{nt} , not vice versa. The resolution of this apparent contradiction relies on the memorylessness property of Markov models. To understand the intuition, adopt for the moment the perspective of an observer under a “veil of ignorance”, to whom only $\hat{w}_{n,t+1}$ is observable. For this observer, learning the two potential outcomes that beach width could have taken in period t would add no additional information about whether the beach is nourished in year $t + 1$. In other words, conditional only on predicted next period width, whether the beach was nourished in year t is as good as randomly assigned—which means that the outcomes can be viewed in a quasi-experimental framework.

More formally, consider the two key conditions required for a matching estimator to generate the same consistent inference as a randomized experiment: (i) the probability that any particular unit is assigned to the treatment group must be greater than zero and less than one; and (ii) conditional on observable covariates, the treatment must be independent of the potential outcomes (Abadie and Imbens, 2011). Condition (i) is clearly satisfied by limiting the sample of neighborhoods to those with at least one beach nourishment in one year. The argument that condition (ii) is satisfied requires an assumption about the characteristics of neighborhoods that receive beach nourishment. In particular, it requires that there are no omitted variables that are correlated

with both beach width and housing prices. For the remainder of the analysis, I assume that this assumption—that conditional on the covariate $\hat{w}_{n,t+1}$, the treatment $N_{n,t+1}$ is independent of the potential outcomes $w_{nt}(1)$ and $w_{nt}(0)$ —is true.

Given the potential outcomes defined above, I now follow Abadie and Imbens (2011) and develop nearest neighbor matching estimators for the treatment effects of next-period nourishment on current beach width and current housing price indices. Consider first the population average treatment effect for the subpopulation of treated units (PATT) for current beach width:

$$PATT_w = E[w_{nt}(1) - w_{nt}(0) | N_{n,t+1} = 1] \quad (1.19)$$

Because both $w_t(1)$ and $w_t(0)$ cannot both be observed for the same unit, I use a nearest neighbor matching strategy to estimate the unobserved outcomes. This procedure selects, for each treated unit, the set of k untreated units that are most similar on the basis of a vector of matching variables \mathbf{X}_{nt} that includes $\hat{w}_{n,t+1}$. The comparison of similarity between a treated unit and an untreated unit is based on a distance metric calculated as the norm of the vector of differences between the vectors of matching variables for the treated and untreated unit. Let Ω_{nt} denote the set of k untreated units that are nearest neighbors for treated observation nt . I estimate the unobserved potential outcomes $w_{nt}(0)$ as:

$$\hat{w}_{nt}(0) = \sum_{m \in \Omega_{nt}} \frac{1}{k} w_{mt} \quad (1.20)$$

for units nt for which $N_{n,t+1} = 1$. I then use this estimated potential width outcome to create a matching estimator for the PATT:

$$\widehat{PATT}_w = \frac{1}{q} \cdot \sum_{nt: N_{n,t+1}=1} (w_{nt}(1) - \hat{w}_{nt}(0)) \quad (1.21)$$

where q represents the number of treated observations, i.e., observations nt for which $N_{n,t+1} = 1$.

I use a similar procedure to estimate the PATT for the price index. I define the PATT as:

$$PATT_P = E[P_{nt}(1) - P_{nt}(0) | N_{n,t+1} = 1] \quad (1.22)$$

I then estimate the unobserved potential price outcomes for treated units as:

$$\hat{\hat{P}}_{nt}(0) = \sum_{m \in \Omega_{nt}} \frac{1}{k} P_{mt} \quad (1.23)$$

and the PATT as:

$$\widehat{PATT}_P = \frac{1}{q} \cdot \sum_{nt: N_{n,t+1}=1} (P_{nt}(1) - \hat{\hat{P}}_{nt}(0)) \quad (1.24)$$

Now return to the problem of estimating the pre-nourishment cross-sectional relationship between housing prices and beach width. A natural way to estimate the derivative of price with respect to width, evaluated at the potential width outcome $w_t(0)$ in year t , is:

$$\begin{aligned} \left. \frac{\partial p}{\partial w_t} \right|_{w_t(0)} &= E \left[\frac{P_t(1) - P_t(0)}{w_t(1) - w_t(0)} \right] \\ &= \frac{\widehat{PATT}_P}{\widehat{PATT}_w} \end{aligned} \quad (1.25)$$

1.6 Results

In this section I present the main empirical results from my analysis. I begin with the results from the repeat sales research design.

1.6.1 Repeat Sales Results

To establish a baseline against which to compare later results, Columns (1) and (2) of Table 1.4 present the results from “conventional” OLS hedonic regressions of log sales price on housing characteristics and semi-parametric beach width dummy variables. Column (1) presents results for all properties; Column (2) presents results for only properties with repeat sales. These regression control for year-by-county and year-by-housing type (condo vs single family) fixed effects, but do not include property fixed effects. The coefficients show that prices increase with beach width, at least for beaches

less than 300 feet wide. For example, Column (2) shows that a house located on a beach that is 0 to 49 feet wide has a sale price that is 5.4 percent lower than a house located on a beach that is 200 to 249 feet wide.

Columns (3) and (4) of Table 1.4 show results from similar regressions based on a repeat-sales approach. Unlike the conventional OLS regressions, the repeat sales regressions include property fixed effects that control for idiosyncratic characteristics of each individual house or condo. These regressions show a much more modest relationship between sales price and beach width. A house in the 0-49 foot category sells for a statistically insignificant 1.9 percent less than a house in the 200-249 foot category.

Figure 1.4 compares the conventional and repeat-sales coefficients from Columns (2) and (3) of Table 1.4. The figure emphasizes the fact that the conventional results overstate the relationship between price and width, relative to the repeat sales approach. Although both approaches suggest that houses with less than 100 feet of beach experience a modest price discount, the magnitude of this discount is much greater in the conventional coefficients.

Figure 1.5 presents an alternative repeat-sales specification, based on a more detailed set of 10-foot beach width bins. The figure confirms the general patterns from the main repeat-sales analysis. However, the point estimates from this figure show that homes located with less than 20 feet of beach sell for a price discount of between 6 and 14 percent, compared to houses with 200 feet of beach. This discount is highly suggestive, but only marginally statistically significant.

1.6.2 Differences-in-Differences Results

Assessment of Research Design

Before presenting the main differences-in-differences results, I present evidence on the appropriateness of the research design. One important question is whether the tim-

Table 1.4: Housing Prices and Beach Width: Panel Results

	(1)	(2)	(3)	(4)
Width: 0 ft	-0.045* (0.022)	-0.054* (0.024)	-0.019 (0.020)	-0.022 (0.019)
Width: 50 ft	-0.029 (0.017)	-0.032 (0.018)	-0.021 (0.013)	-0.007 (0.013)
Width: 100 ft	0.007 (0.014)	-0.004 (0.015)	-0.023* (0.009)	-0.020 (0.010)
Width: 150 ft	-0.015 (0.013)	-0.019 (0.011)	-0.014 (0.008)	-0.010 (0.008)
Width: 250 ft	0.004 (0.015)	0.018 (0.017)	0.015 (0.013)	0.011 (0.012)
Width: 300 ft	0.066* (0.030)	0.042 (0.030)	-0.012 (0.026)	-0.024 (0.028)
Width: 350 ft	0.014 (0.017)	0.010 (0.019)	-0.020 (0.035)	-0.037 (0.034)
Width: 400 ft	0.014 (0.032)	-0.005 (0.026)	0.011 (0.032)	0.002 (0.033)
Width: 450 ft	0.007 (0.049)	0.025 (0.066)	-0.009 (0.042)	-0.045 (0.046)
Width: 500 ft	-0.005 (0.049)	-0.014 (0.049)	-0.025 (0.044)	-0.044 (0.045)
Non-vacant	0.198*** (0.052)	0.130 (0.073)	0.175** (0.053)	0.137* (0.054)
Renovated	0.105** (0.034)	0.150*** (0.040)	-0.020 (0.029)	0.013 (0.034)
Log Acreage	0.095* (0.045)	0.087 (0.050)		
Total Area (sq ft)	0.000*** (0.000)	0.000*** (0.000)		
Bedrooms	0.103*** (0.012)	0.098*** (0.015)		
Bathrooms	-0.004 (0.018)	-0.007 (0.022)		
Effective Year Built	0.006*** (0.001)	0.007*** (0.002)		
Actual Year Built	0.002* (0.001)	0.001 (0.001)		
Structure Quality (1-6)	0.007 (0.012)	0.007 (0.015)		
Brick Construction	-0.093*** (0.022)	-0.104*** (0.027)		
Extra Features Value	-0.000 (0.000)	-0.000 (0.000)		
Sales	57,199	23,547	23,547	23,547
Clusters	64	64	64	64
R-squared	0.779	0.792	0.629	0.582
Year-Area FE	No	No	No	Yes
Year-County FE	Yes	Yes	Yes	No
Year-Housing Type FE	Yes	Yes	Yes	No
County FE	Yes	Yes	No	No
Parcel FE	No	No	Yes	Yes
Repeat Sales Only	No	Yes	Yes	Yes

Note: This table presents coefficients and standard errors from a regression of log sales price on beach width. Column (1) shows a conventional hedonic analysis based on all available sales transactions, and Column (2) shows a conventional hedonic analysis based only on sales at houses with two or more sales. Columns (3) and (4) show results from a fixed-effects repeat-sales analysis corresponding to equation (1.11). In all specifications, each observation represents a unique housing sale between 1983 and 2009, for which beach width data was available, at a house within 20 meters of the beach. Standard errors are clustered by six-mile zone, and regressions are weighted at the same level. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

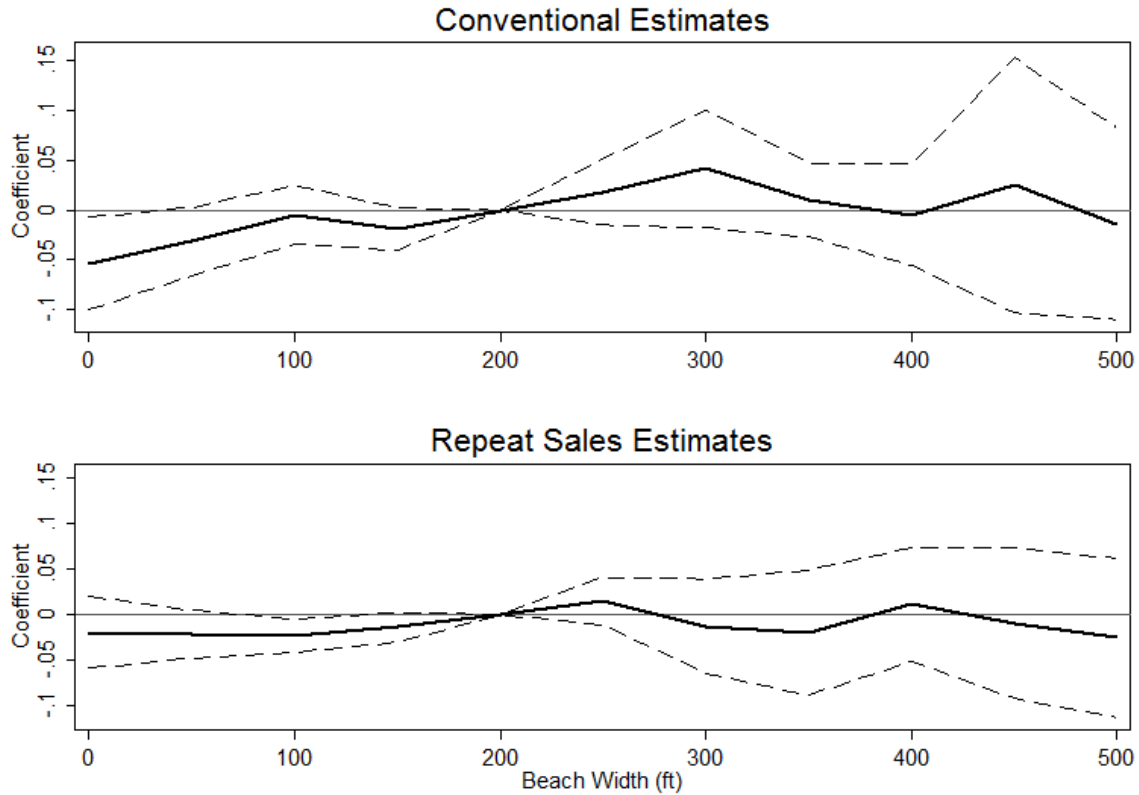


Figure 1.4: Panel Regression Results: Capitalized Price Versus Beach Width

Note: This figure plots coefficients and 95 percent confidence intervals from Table 1.4. The top panel presents the conventional OLS hedonic estimates from Column (2) of Table 1.4, and bottom panel presents the repeat-sales estimates from Column (3). The dependent variable in both specifications is log sales price. The omitted width category is 200 feet. The figure is based on the set of housing transactions at properties located within 20 meters of the beach.

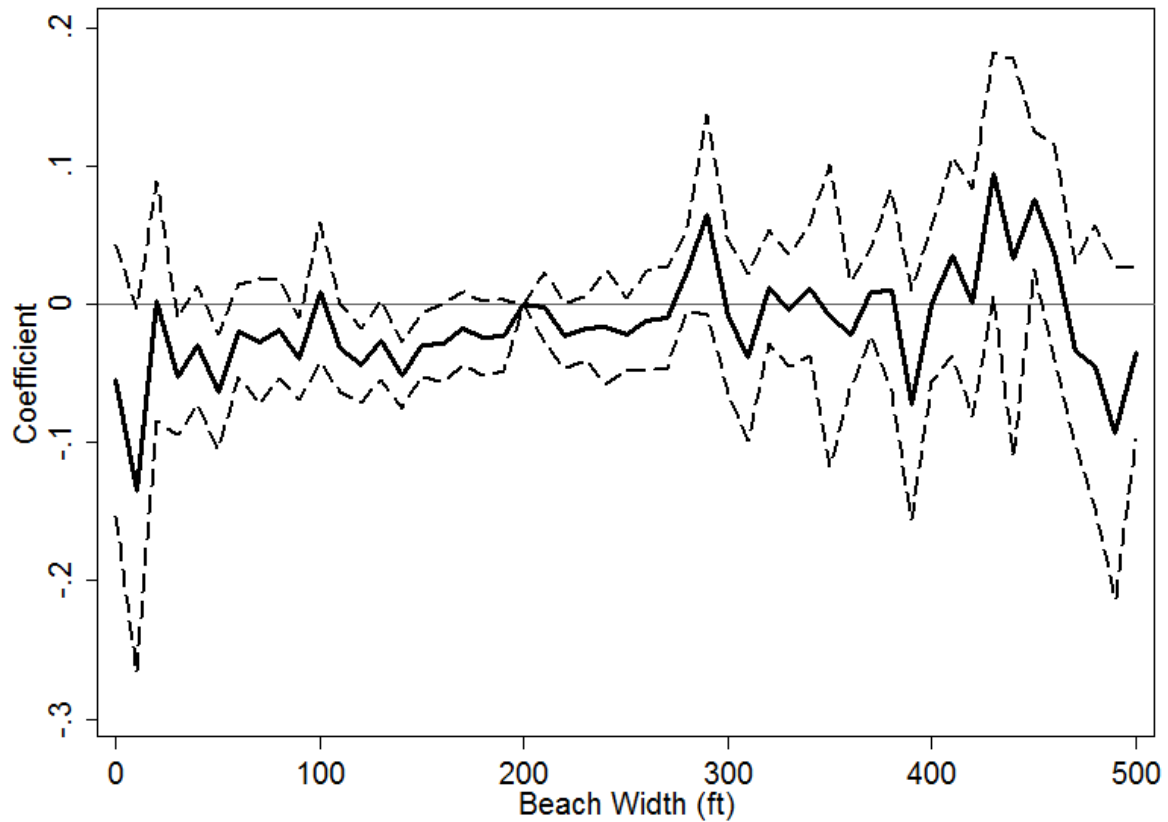


Figure 1.5: Panel Regression Results: Capitalized Price Versus Beach Width, Detail

Note: This figure plots coefficients and 95 percent confidence intervals for a repeat sales regression similar to Equation (1.11), but with the beach width variable divided into 10 foot bins. The dependent variable is log sales price, and the figure is based on the set of housing transactions at properties located within 20 meters of the beach. The omitted width category is 200 feet.

Table 1.5: Characteristics of Properties Sold Before and After Nourishment

	Two Years Before	One Year Before	Year of Nourishment	One Year After
Condo	0.68 (0.03)	0.69 (0.03)	0.69 (0.03)	0.69 (0.03)
Parcel Acreage	0.63 (0.07)	0.60 (0.08)	0.59 (0.07)	0.66 (0.08)
Housing Area (sq ft)	2,555 (161)	2,424 (145)	2,385 (152)	2,698 (194)
Bedrooms	1.56 (0.10)	1.46 (0.10)	1.42 (0.09)	1.47 (0.09)
Bathrooms	1.22 (0.09)	1.13 (0.09)	1.11 (0.09)	1.14 (0.09)
Year Renovated	1987.0 (0.8)	1986.5 (0.9)	1986.9 (0.8)	1986.7 (0.9)
Year Built	1979.8 (0.8)	1979.1 (0.9)	1978.8 (0.9)	1979.9 (0.8)
Structure Quality (1-6)	3.41 (0.07)	3.35 (0.07)	3.33 (0.07)	3.33 (0.07)
Brick Construction	0.21 (0.02)	0.22 (0.02)	0.21 (0.02)	0.22 (0.02)
Features Appraised Value	15,556 (2,951)	9,782 (1,715)	24,802 (12,771)	17,789* (3,406)
Sale Price (000s)	685 (132)	745 (160)	844 (169)	779 (144)
Beach Width (ft)	199 (9)	180 (7)	261*** (7)	250*** (7)
Major Nourishments	0.02** (0.01)	0.00 (0.00)	1.00 (0.00)	0.00 (0.00)
Intensity (cy/ft)	2.97* (1.24)	0.34 (0.13)	56.70*** (1.62)	0.34 (0.14)
neighborhoods	294	303	317	317

Note: This table shows means and standard errors of the characteristics of properties sold in the years before and after beach nourishment, for properties located within 20 meters of the beach. The table includes pre and post data for each nourishment project with intensity greater than 25 cy/ft. Each observation represents the mean characteristics for a one-mile neighborhood in a particular year. The t-tests are based on unpaired, two-tailed comparisons, relative to the one year before nourishment category. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

ing of beach nourishment projects is related to the characteristics of parcels that are sold. For the differences-in-differences design to be valid, nourishment must not be correlated with unobservable parcel characteristics. Although this criterion is fundamentally untestable, it is more likely to be true if nourishment is uncorrelated with observable property characteristics.

Table 1.5 summarizes the characteristics of parcels that are sold in the two years before and after nourishment. The table shows that housing characteristics are strongly balanced before and after nourishment. Across a variety of characteristics—including vacancy, parcel acreage, living area, year built, and number of bedrooms and bathrooms—there are no statistically significant differences in pre- and post- characteristics. The primary characteristics that do show a significant relationship with nourishment timing are beach width, nourishment status, and nourishment intensity.

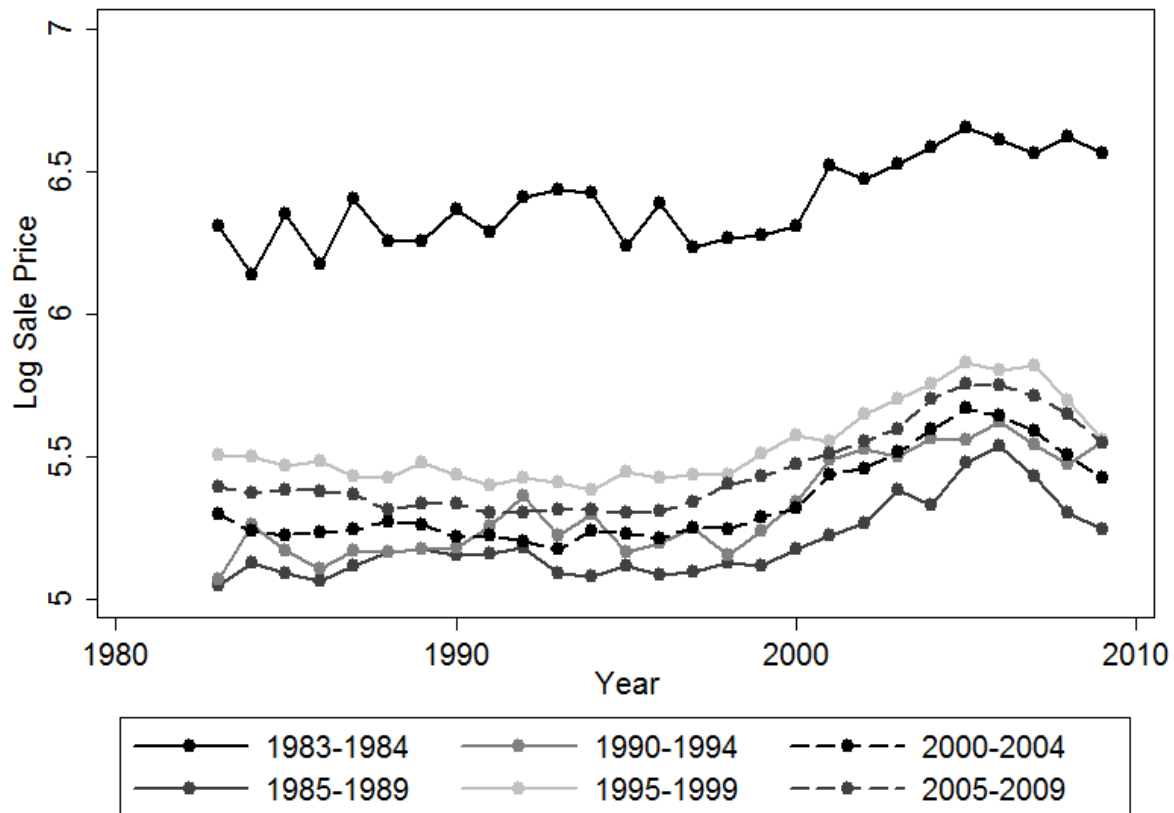


Figure 1.6: Trend in Sales Price, by Year of Most Recent Nourishment

Note: The figure plots mean log sale price over time for properties located within 20 meters of the beach, by the five-year period in which the beach near each survey monument was last nourished. Means are calculated across all one-mile neighborhood-by-year observations. The figure excludes areas that were not nourished between 1983 and 2009.

Another important test of a differences-in-differences design is whether the control group has a similar pre-intervention trend to the treatment group. Figure 1.6 plots the aggregate trend in the neighborhood housing price indices for six groups, based on the most recent five-year period in which the beach was nourished. The figure shows that homes in these different categories do experience very similar price trends. However, the data are somewhat noisy, suggesting that it may be important to control for other sources of variation, such as county-by-year price trends.

A final criterion for the research design is whether the intervention has a meaningful effect, i.e., whether nourishment increases beach width. Table 1.6 presents the results

Table 1.6: Beach Width and Nourishment

	(1)	
Nourishment, 5 years before	-4.08	(7.35)
Nourishment, 4 years before	-22.75**	(6.60)
Nourishment, 3 years before	-12.58*	(6.11)
Nourishment, 2 years before	-19.54**	(6.27)
Nourishment, 1 year before	-38.28***	(6.77)
Nourishment	45.43***	(4.83)
Nourishment, 1 year after	31.32***	(4.35)
Nourishment, 2 years after	17.82***	(4.25)
Nourishment, 3 years after	17.86**	(6.02)
Nourishment, 4 years after	21.49**	(7.65)
Observations	34,243	
Clusters	61	
R-squared	0.149	
Elapsed Year-Base Year FE	Yes	
Elapsed Year-Location FE	Yes	

Note: This table presents coefficients and standard errors from the differences-in-differences regression corresponding to equation (1.15). The dependent variable is beach width, in feet. Each observation represents a unique one-mile neighborhood and year between 1983 and 2009, for which beach width data was available. Standard errors are clustered by six-mile zone, and regressions are weighted at the same level. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

from estimating Equation (1.15), which compares the width of beaches that are nourished and unnourished in a particular year. The table shows that nourishment causes a sharp, highly significant increase in width. Compared to the previous year, the average section of beach gains 83 feet in width in the year of nourishment (with a 95 percent confidence interval of 59 to 97 feet). Overall, the coefficients show an intuitive pattern: the beach erodes in the years before nourishment, gains considerable width in the year of nourishment, and then continues eroding.

As a further test of the nourishment intervention, Table 1.7 shows the results from estimating the change in beach width as a function nourishment intensity (cubic yards of sand placed per foot of beach). The table shows results for three intensity categories, 25 to 49 cy/ft, 50 to 74 cy/ft and >75 cy/ft, relative to the omitted category of no nourishment (results for the 1-24 cy/ft category are similar but not shown). The coefficients show that changes in width are strongly increasing in nourishment intensity. For example, nourishments in the 25-49 cy/ft category cause the beach to increase by

an additional 57 feet, whereas nourishments in the >75 cy/ft category cause the beach to increase by an additional 121 feet.

Main Differences-in-Differences Results

Table 1.8 shows the results of estimating Equation (1.14), which compares the housing price index for one-mile neighborhoods located near beaches that are nourished or unnourished in a particular year. To allow for the possibility that beach width may be more important for houses located closer to the beach, the table presents separate results for properties in three different distance categories: 0 to 19 meters from the beach, 20 to 799 meters from the beach, and 800 to 5,000 meters from the beach.

The regression results in all three categories reveal a consistent pattern: nourishment projects have no effect on sales prices. For example, for parcels located directly on the beach (the 0-19 meters category), the coefficients imply that housing prices increase by a statistically insignificant 1.2 percent between two years before nourishment and the year after of nourishment, with a 95 percent confidence interval of -2.5 percent to +4.9 percent. To illustrate the results, Figure 1.7 plots changes in beach width and changes in housing prices (in the 0-19 m category), relative to the number of years elapsed since nourishment. The figure emphasizes the fact that although nourishment causes a sharp change in beach width, it has no immediate effect on housing prices.

Table 1.9 presents the results of an alternative differences-in-differences specification that compares the change in prices accompanying high-intensity nourishments against the change in prices accompanying low-intensity nourishments. The results again reveal that housing prices do not respond to the timing of nourishment projects. As shown in Figure 1.8, even though higher intensity nourishments cause a greater increase in beach width, these changes in beach width are not reflected in housing prices, which show no obvious relationship to the timing of nourishment.

Table 1.7: Beach Width and Nourishment, by Nourishment Intensity

	(1)	
25-49 cy/ft, 5 years before	1	(10)
25-49 cy/ft, 4 years before	-6	(7)
25-49 cy/ft, 3 years before	-1	(7)
25-49 cy/ft, 2 years before	-4	(7)
25-49 cy/ft, 1 year before	-20*	(9)
25-49 cy/ft	37***	(5)
25-49 cy/ft, 1 year after	20**	(6)
25-49 cy/ft, 2 years after	3	(6)
25-49 cy/ft, 3 years after	3	(7)
25-49 cy/ft, 4 years after	25**	(8)
50-74 cy/ft, 5 years before	-13	(10)
50-74 cy/ft, 4 years before	-28*	(11)
50-74 cy/ft, 3 years before	-13	(9)
50-74 cy/ft, 2 years before	-38***	(6)
50-74 cy/ft, 1 year before	-53***	(10)
50-74 cy/ft	40***	(8)
50-74 cy/ft, 1 year after	22**	(7)
50-74 cy/ft, 2 years after	14*	(5)
50-74 cy/ft, 3 years after	12	(7)
50-74 cy/ft, 4 years after	8	(11)
75-up cy/ft, 5 years before	-8	(11)
75-up cy/ft, 4 years before	-39***	(11)
75-up cy/ft, 3 years before	-29*	(14)
75-up cy/ft, 2 years before	-27	(22)
75-up cy/ft, 1 year before	-54***	(10)
75-up cy/ft	67***	(8)
75-up cy/ft, 1 year after	57***	(6)
75-up cy/ft, 2 years after	42***	(7)
75-up cy/ft, 3 years after	36*	(14)
75-up cy/ft, 4 years after	32	(17)
Observations	34,243	
Clusters	61	
R-squared	0.157	
Elapsed Year-Base Year FE	Yes	
Elapsed Year-Location FE	Yes	

Note: This table presents coefficients and standard errors from differences-in-differences regression corresponding to equation (1.15), where the effects of nourishment are disaggregated by intensity into four categories: 1-24 cy/ft, 25-49 cy/ft, 50-74 cy/ft, and ≥ 75 cy/ft. The dependent variable is beach width, in feet. Each observation represents a unique one-mile neighborhood and year between 1983 and 2009, for which beach width data was available. Standard errors are clustered by six-mile zone, and regressions are weighted at the same level. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

Table 1.8: Housing Prices and Nourishment: Differences-in-Differences Results

	0-19 m	20-799 m	800-5000 m
Nourishment, 5 years before	-0.019* (0.008)	-0.014 (0.008)	-0.019 (0.012)
Nourishment, 4 years before	-0.011 (0.010)	-0.017* (0.007)	-0.018** (0.006)
Nourishment, 3 years before	-0.011 (0.008)	-0.013* (0.006)	0.009 (0.020)
Nourishment, 2 years before	-0.019 (0.014)	-0.015 (0.008)	0.008 (0.017)
Nourishment, 1 year before	-0.023* (0.010)	-0.005 (0.009)	-0.008 (0.008)
Nourishment	-0.016 (0.011)	-0.011 (0.007)	0.001 (0.012)
Nourishment, 1 year after	-0.007 (0.012)	-0.017** (0.006)	0.010 (0.014)
Nourishment, 2 years after	0.000 (0.010)	-0.017* (0.007)	0.000 (0.009)
Nourishment, 3 years after	0.004 (0.012)	-0.001 (0.007)	0.005 (0.011)
Nourishment, 4 years after	0.011 (0.008)	0.002 (0.007)	0.010 (0.011)
Observations	65,071	81,462	82,006
Clusters	61	67	65
R-squared	0.534	0.659	0.468
Elapsed Year-Base Year FE	Yes	Yes	Yes
Elapsed Year-Location FE	Yes	Yes	Yes

Note: This table presents coefficients and standard errors from the differences-in-differences regression corresponding to equation (1.14). The dependent variable is the log of the sale price index for properties located within 20 meters of the beach. Each observation represents a unique one-mile neighborhood and year between 1983 and 2009. Standard errors are clustered by six-mile zone, and regressions are weighted at the same level. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

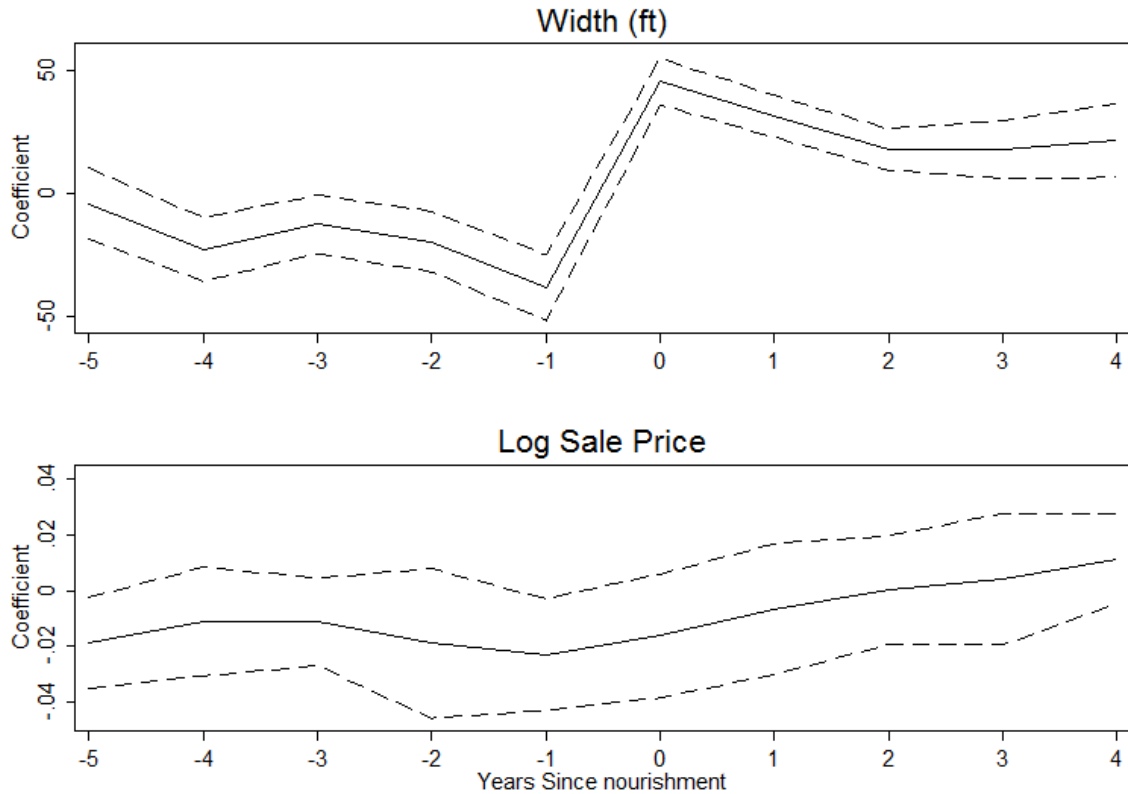


Figure 1.7: Differences-in-Differences Results

Note: The top panel plots 95 percent confidence intervals for the differences-in-differences coefficients from Table 1.6. The dependent variable is beach width in feet. The bottom panel plots confidence intervals for the differences-in-differences coefficients from Table 1.8. The dependent variable is log sales price. The shared x-axis represents the number of years since nourishment, where nourishment occurs in year 0. The figure is based on the set of properties located within 20 meters of the upper end of the beach.

Table 1.9: Housing Prices and Nourishment: Diff-in-Diff Results by Intensity

	0-19 m		20-799 m		800-5000 m	
25-49 cy/ft, 5 years before	-0.025	(0.013)	-0.012	(0.010)	-0.009	(0.007)
25-49 cy/ft, 4 years before	-0.024*	(0.011)	-0.017	(0.014)	-0.018*	(0.008)
25-49 cy/ft, 3 years before	-0.031**	(0.010)	-0.023**	(0.008)	-0.022*	(0.009)
25-49 cy/ft, 2 years before	-0.011	(0.013)	0.002	(0.011)	-0.016	(0.010)
25-49 cy/ft, 1 year before	-0.017	(0.012)	0.002	(0.010)	-0.004	(0.008)
25-49 cy/ft	-0.011	(0.020)	-0.015	(0.012)	-0.021*	(0.010)
25-49 cy/ft, 1 year after	-0.003	(0.019)	-0.018*	(0.008)	-0.011	(0.008)
25-49 cy/ft, 2 years after	-0.000	(0.015)	-0.017	(0.010)	-0.001	(0.008)
25-49 cy/ft, 3 years after	0.012	(0.017)	0.008	(0.013)	-0.011	(0.008)
25-49 cy/ft, 4 years after	0.020	(0.012)	0.017	(0.013)	0.004	(0.008)
50-74 cy/ft, 5 years before	-0.004	(0.020)	0.001	(0.009)	0.005	(0.009)
50-74 cy/ft, 4 years before	-0.023	(0.018)	-0.006	(0.008)	-0.015	(0.010)
50-74 cy/ft, 3 years before	0.006	(0.012)	0.002	(0.010)	-0.014	(0.009)
50-74 cy/ft, 2 years before	-0.035	(0.022)	-0.031*	(0.015)	-0.006	(0.008)
50-74 cy/ft, 1 year before	-0.039*	(0.016)	-0.003	(0.020)	-0.007	(0.009)
50-74 cy/ft	-0.023	(0.012)	-0.005	(0.011)	-0.019*	(0.010)
50-74 cy/ft, 1 year after	-0.023	(0.015)	-0.013	(0.011)	-0.016	(0.014)
50-74 cy/ft, 2 years after	0.006	(0.017)	-0.015	(0.011)	-0.029	(0.014)
50-74 cy/ft, 3 years after	-0.007	(0.014)	0.000	(0.011)	-0.008	(0.009)
50-74 cy/ft, 4 years after	0.008	(0.015)	-0.005	(0.012)	-0.013	(0.016)
75-up cy/ft, 5 years before	-0.018	(0.017)	-0.009	(0.013)	-0.004	(0.020)
75-up cy/ft, 4 years before	0.037	(0.025)	-0.023	(0.014)	-0.000	(0.012)
75-up cy/ft, 3 years before	0.005	(0.023)	-0.005	(0.011)	0.010	(0.014)
75-up cy/ft, 2 years before	0.027	(0.021)	-0.013	(0.016)	-0.001	(0.014)
75-up cy/ft, 1 year before	-0.005	(0.019)	0.002	(0.009)	0.010	(0.018)
75-up cy/ft	0.001	(0.019)	0.004	(0.016)	0.020*	(0.008)
75-up cy/ft, 1 year after	-0.004	(0.021)	-0.014	(0.013)	0.015	(0.011)
75-up cy/ft, 2 years after	-0.024*	(0.011)	-0.017	(0.011)	-0.001	(0.011)
75-up cy/ft, 3 years after	-0.008	(0.026)	-0.019*	(0.009)	-0.012	(0.016)
75-up cy/ft, 4 years after	-0.007	(0.017)	-0.024	(0.018)	-0.013	(0.020)
Observations	65,071		81,462		82,006	
Clusters	61		67		65	
R-squared	0.530		0.660		0.496	
Elapsed Year-Base Year FE	Yes		Yes		Yes	
Elapsed Year-Location FE	Yes		Yes		Yes	

Note: This table presents coefficients and standard errors from the differences-in-differences regression corresponding to equation (1.14), where the effects of nourishment are disaggregated by intensity into four categories: 1-24 cy/ft, 25-49 cy/ft, 50-74 cy/ft, and ≥ 75 cy/ft.. The dependent variable is the log of the sale price index for properties located within 20 meters of the beach. Each observation represents a unique one-mile neighborhood and year between 1983 and 2009. Standard errors are clustered by six-mile zone, and regressions are weighted at the same level. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

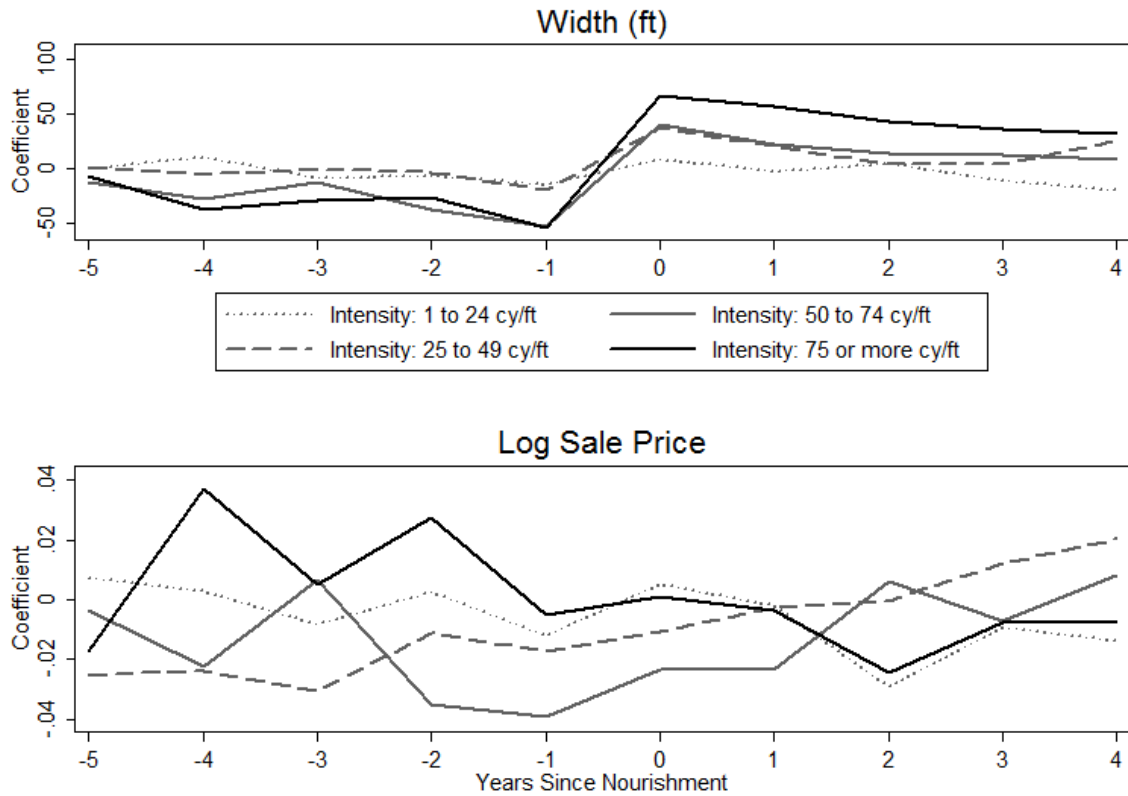


Figure 1.8: Differences-in-Differences Results by Nourishment Intensity

Note: The top panel plots the differences-in-differences coefficients from Table 1.7. The dependent variable is beach width in feet. The bottom panel plots the differences-in-differences coefficients from Table 1.9. The dependent variable is log sales price index for houses located within 20 meters of the beach. The shared x-axis represents the number of years since nourishment, where nourishment occurs in year 0. The intensity categories represent the quantity of sand placed on the beach, in cubic yards, per foot of shoreline. Both panels present coefficients relative to sections of beach that are not nourished in year 0.

Table 1.10: Beach Width Markov Regression Results

	(1)	(2A)	(2B)	(2C)	(2D)	(2E)
Width, year -1	0.91*** (0.02)	0.93*** (0.01)	0.91*** (0.03)	0.92*** (0.03)	0.82*** (0.03)	0.88*** (0.02)
Major Nourishment	66.44*** (6.62)	68.00*** (16.80)	87.00*** (10.16)	59.53*** (4.69)	41.14 (17.21)	37.29** (7.97)
Constant	16.62*** (3.21)	15.12*** (2.68)	15.51** (5.01)	15.98** (4.65)	40.38** (6.70)	27.08* (7.50)
Observations	4,342	1,534	1,820	572	156	260
Clusters	50	33	30	12	4	6
R-squared	0.850	0.880	0.834	0.840	0.689	0.810

Note: This table presents the coefficients from estimating equation (1.17), with standard errors in parentheses. The dependent variable is current beach width, in feet. Column (1) includes all observations. Columns (2A) through (2E) include the set of beaches that are nourished 1, 2, 3, 4, or 5 times, respectively, during the period from 1983 to 2009. In all models, standard errors are clustered by six-mile zones. * denotes $p < .05$; ** denotes $p < .01$; *** denotes $p < .001$.

1.6.3 Discontinuity Matching Results

The discontinuity matching procedure includes two substantive components: developing a rational model of beliefs about future beach width, and using a nearest neighbor matching procedure to calculate the treatment effect of future beach nourishment on current beach width and housing prices.

Table 1.10 shows the results of estimating a model of beliefs about beach width corresponding to Equation (1.17). Column (1) represents the simplest possible AR(1) specification, in which beach width is modeled as a function of lagged beach width, a binary nourishment variable, and a constant. The results indicate that beach width is highly autocorrelated, with a coefficient of .91 on lagged width, and that nourishment causes a 66 foot increase in beach width. The coefficients are very precisely estimated.

Columns (2A) through (2E) show the results of estimating separate results for beaches with 1, 2, 3, 4, or 5 nourishments, respectively, during the period from 1983 to 2009. The rationale for estimating separate regressions is that beaches that erode at faster rates may also be nourished more frequently. The regressions results confirm this hypothesis: the coefficient on lagged width decreases from .93 and .91 at beaches with 1

or 2 nourishments to .82 and .88 at beaches with 4 or 5 nourishments. Furthermore, the effects of nourishment on beach width are larger at beaches that are nourished less frequently.

Figure (1.9) illustrates the results from the matching stage of the discontinuity matching procedure. The top panel shows how beach width evolves over time for two groups of neighborhood-by-year observations: neighborhoods that are nourished (“treated units”), and similar unnourished neighborhoods that are predicted to have similar beach width in the following year (“control units”). The panel shows that the treated and control neighborhoods have reasonably similar beach width in the year of nourishment, as would be expected, given that these units are matched on predicted width for that year. In subsequent years, the treated and control groups also show similar trends, which is encouraging and suggests that the Markov assumption underlying the discontinuity matching approach is valid. However, in the year before nourishment, the treated neighborhoods show a substantial decline in beach width relative to the control group. From the perspective of the discontinuity matching procedure, this decline is the treatment effect of nourishment on pre-nourishment beach width.

The second panel of Figure (1.9) shows similar results for the effect of nourishment on pre-nourishment housing prices. Again, the panel shows that the treatment and control groups have similar values of the price index in the year of nourishment, and experience similar post-nourishment price trends. Unlike the width results, however, the panel shows that there is no discernable effect of nourishment on pre-nourishment housing prices. In other words, the treatment effect of predictable future beach nourishments on current housing prices is statistically indistinguishable from zero.

Table 1.11 presents estimates of the population average treatment effects on the treated from Equations (1.21) and (1.24). These estimates confirm the qualitative results from Figure (1.9). Although next-period nourishment is associated with a highly significant 60 foot decrease in current beach width, it is linked to an insignificant decrease of \$1,729

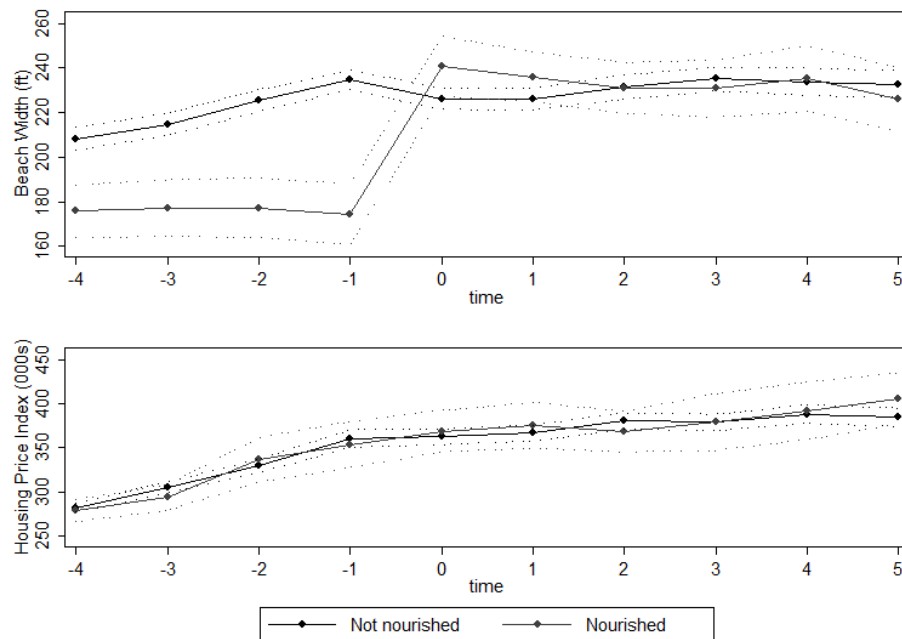


Figure 1.9: Discontinuity Matching Results: Trends in Width and Price

Note: This figure shows trends in the outcome variable for the treatment and control groups for the discontinuity matching procedure. In both panels, the treatment variable is nourishment in elapsed year 0, and the matching variable is predicted width in elapsed year 0, based on information available in year -1. The upper panel shows actual width and the lower panel shows the actual price index for properties located within 20 meters of the beach.

Table 1.11: Discontinuity Matching Results

Outcome	Coefficient	Std Err	95% Confidence Interval	
PATT, width	-60	4	-67	-53
PATT, sale price	-1,729	8,808	-18,992	15,535
$\partial P / \partial w$	29	.	.	.

Note: This table presents nearest neighbor matching results corresponding to Equations (1.21) and (1.24). Each observation represents a unique one-mile neighborhood and year. Unobserved potential outcomes for nourished beaches are estimated using the 6 closest unnourished observations. The vector of match variables includes predicted next-period beach width, year, and the total number of nourishments occurring at each beach. Matching is exact on the year variable, and the standard errors are robust. The table also shows the result of estimating Equation (1.25). To avoid being misleading, standard errors and confidence intervals are not shown for this estimate, which represents the ratio of two random variables with an unknown degree of correlation. The figure is based on the set of housing transactions at properties located within 20 meters of the beach.

in the current housing price index. Combining these two treatment effects suggests that the average homeowner is willing to pay \$29 to rent an extra foot of beach for one year. Because this estimate represent the ratio of two random variables with possibly correlated distributions, I do not calculate a standard error or a confidence interval. However, the significant uncertainties in the treatment effect for price suggest it would be difficult to reject the null hypothesis that willingness to pay is \$0.

1.7 Discussion and Policy Implications

1.7.1 Summary of Main Results

The empirical results from the previous section are striking. Using three different research designs, I find consistent evidence that homeowners place little value on the width of coastal beaches. The repeat-sales regressions suggest a statistically significant positive effect of beach width on prices, but the effect is small, and only holds for houses located on beaches that are less than 250 feet wide. Furthermore, because these panel regressions do not control for factors that change over time within neighborhoods (e.g., wealth, political influence), there is a possibility that the direction of causality runs from prices to beach width, not vice versa. In contrast, based on the

differences-in-differences regressions and the discontinuity matching results, I am unable to reject the null hypothesis that homeowners' willingness to pay to avoid coastal shoreline loss is zero.⁷ Since these two research designs are based on sharp variation in beach width caused by nourishment projects, they are more likely to identify the true causal effect of width on price.

However, the results also suggest a second important finding: the marginal benefits of an extra foot of beach may be highly nonlinear. The repeat-sales results show that the relationship between housing prices and beach width is only positive for houses with less than 250 feet of beach. Furthermore, houses with extremely eroded beaches sell for a substantial discount. In particular, the point estimates indicate that houses with less than 20 feet of remaining beach have sales prices that are 6 to 14 percent lower than houses with 200 feet of beach. Because the repeat sales design has weaknesses, and because I observe relatively few homes located on beaches with such high levels of erosion, this conclusion should not be overemphasized. Nonetheless, it suggests the possibility that the welfare costs of sea level rise may be low up to a threshold, and then increase sharply.

1.7.2 Cost-Benefit Analysis of Beach Nourishment

Setting aside any nonlinearities in the marginal benefits schedule, and ignoring the fact that only the repeat-sales coefficients are significantly different than zero, consider the following back-of-the-envelope calculations based on a literal interpretation of the point estimates from the results section. Table 1.12 presents the results of a simple break-even cost benefit analysis of the decision to nourish a section of beach. As the table shows, the differences-in-differences and repeat-sales point estimates indicate that

⁷This finding is particularly striking in light of the large amount of money that has been spent on beach nourishment over the last 50 years (U.S. ACE, 1996). However, historically, beach nourishment has been heavily subsidized by the U.S. and Florida governments, with local communities paying less than half of the actual cost of beach nourishment (NRC, 1995; U.S. ACE, 1996). Thus, the decision to nourish a particular beach may not be a valid measure of a local community's revealed willingness to pay for wider beaches.

Table 1.12: Willingness to Pay for Wider Beaches

Research Design	WTP metric	$\partial P/\partial w$	Value of 70 foot nourishment	Break-even HH per mile
Repeat Sales, <250 feet	One foot nourishment	\$68	\$4,760	210
Differences-in-Differences	One foot nourishment	\$42	\$2,927	342
Discontinuity Matching	One foot rental	\$29	\$13,685	73

Note: The table summarizes point estimates of willingness to pay for wider beaches based on the three research designs. The “WTP metric” column describes the correct welfare interpretation of each estimate of $\partial P/\partial w$. These calculations are based on a \$500,000 home. The “Value of 70 foot nourishment” column presents an estimate of the value per household of a 70 foot beach nourishment. For the discontinuity matching estimates, this nourishment is assumed to have a lifetime of 10 years, during which beach width decays following the model presented in Table 1.10. The “Break-even HH per mile” column describes the number of households that would have to be located on a one-mile section of beach in order for the marginal benefits of nourishment to exceed the \$1,000,000 nourishment cost per mile. Note that this break-even analysis does not consider other contributions that beach nourishment makes to local economies. Also note that the results apply only to properties located within 20 meters of the beach. Finally, note that these calculations are based on point estimates, and that 95 percent confidence intervals are consistent with a broader range of possible willingness-to-pay values.

willingness to pay for a one-foot beach nourishment project is between \$42 and \$68 (based on a \$500,000 home). This implies that a project that adds 70 feet of width to a beach would generate benefits of \$2,927 to \$4,760 per household. Since the cost of beach nourishment is roughly \$1,000,000 per mile, there would have to be between 210 and 342 beachfront homes per mile in order for the project to generate positive marginal benefits. Note, of course, that this simple break-even analysis ignores any benefits that beach nourishment may generate for non-beachfront properties or for the local economy (e.g., revenues from tourism).⁸

Point estimates based on the discontinuity matching procedure produce somewhat larger results (again, with the caveat that the numbers are not statistically distinguishable from zero). Assuming a decay rate of .91 and a constant of 16.5, the Markov regression results from Table 1.10 suggest that the average beach has a stable equilibrium of approximately 183 feet of width. For a nourishment project that adds 70 feet to the beach, the overall contribution to beach width during the first ten years (a typical nourishment interval) is 474 foot-years. At a marginal value of \$29 per foot-year,

⁸To put these numbers in perspective, heavily developed areas (e.g., Miami Beach) with many high rise buildings might have over a thousand condos and apartments per mile of beach. Less developed areas might have fewer than fifty single-family homes.

this implies that willingness to pay for a nourishment project is \$13,685 per household. A comparison against the \$1,000,000 per mile cost of nourishment implies that there would have to be 73 beachfront houses per mile of beach in order for a project to generate positive marginal benefits, based solely on benefits to coastal homeowners.

1.7.3 Discussion

Overall, my results paint a picture that is both pessimistic and optimistic. On one hand, the results suggest that because homeowners have relatively low willingness to pay for wider beaches, the welfare costs of shoreline loss—and more broadly, of sea level rise—may not be as serious as believed. In particular, my results show that for a typical section of beach in the normal 100 to 400 foot range, changes in beach width have little impact on property values. However, the results also hint at the possibility that damages from shoreline loss are nonlinear. In particular, because houses located on very eroded beaches appear to experience substantial price discounts relative to homes on wider beaches, it appears possible that there may exist some threshold below which shoreline loss does have serious welfare effects. Because the statistical evidence for this claim is weak, further research is needed.

Additionally, it is important to interpret the overall results of this study in context. The general lack of price effects implies that changes in beach width do not affect coastal homeowners' appraisals of the recreational and use benefits from owning a beachfront home. However, because there are a variety of housing market inefficiencies that could weaken the relationship between prices and beach quality, it is possible that these revealed preference estimates are not an accurate measure of actual benefits. For example, homebuyers may choose homes based on long-term beach width, without taking advantage of the arbitrage opportunity to purchase "undervalued" homes that are experiencing wider than usual beach width. Alternatively, beach width may not be a salient characteristic at the time of home purchase, even though homebuyers

would in fact derive more use benefits from a wider beach. Third, there may be negative externalities, such as crowding by members of the beachgoing public, that cancel out the benefits of wider beaches. Fourth, homebuyers may have unrealistic or overly optimistic beliefs about changes in beach width (e.g., even though I know that many beaches are eroding, I think the beach in front of my new home will accrete instead). Finally, the general equilibrium effects of widespread shoreline loss may be quite different than the partial equilibrium effects of specific beach nourishment projects—so that even if a beach nourishment project in a particular location has little effect on housing prices, the loss of shoreline along the entire Florida coast (as might be caused by sea level rise) could still have substantial price effects.

In respect to whether the results imply that homeowners do not value the storm protection benefits provided by wider beaches, the results are even less clear. At least two alternative explanations are possible. First, homebuyers may suffer from moral hazard, either because they have purchased insurance that reimburses them for storm-related damage, or because they believe that state and federal disaster relief programs will cover their losses. Second, homebuyers may suffer from myopia, in which the storm protection benefits from a wider beach are not a salient attribute at the time of purchase (Berger et al, 2009).

1.8 Conclusion

In this chapter, I have developed a new “discontinuity matching” research design for estimating homeowners’ marginal willingness to pay for time-variant neighborhood characteristics. I use this design, as well as traditional panel and differences-in-differences approaches, to estimate the welfare costs of shoreline loss along coastal beaches. In contrast to previous research that suggests that homeowners are willing to pay a substantial premium to live near wider beaches, I find that changes in beach width have little effect on the sale price of beachfront homes, except at very eroded

beaches. The results suggest that policy interventions to prevent shoreline loss are most valuable near homes that are directly threatened by the ocean.

Chapter 2:
Crime, Weather, and Climate Change

2.1 Introduction

The short-term effects of weather on crime are well documented. Previous work has shown that presumably-random variation in daily and weekly temperatures affects the incidence of both violent and non-violent offenses, with higher temperatures leading to higher levels of criminal activity (Brunsdon et al, 2009; Bushman, Wang, and Anderson, 2005; Cohn, 1990). However, despite the strength of this relationship, there is little evidence on how weather affects patterns of criminal behavior over longer time scales. In particular, there is great uncertainty about how climate change is likely to affect the incidence of crime.

The Intergovernmental Panel on Climate Change predicts that global temperatures are likely to rise by about 5 degrees Fahrenheit (2.8 degrees Celsius) by the year 2099, compared to baseline temperatures during the period from 1980 to 1999 (IPCC, 2007). Studies of the short-term relationship between crime and weather suggest that such a change in temperatures could have dramatic effects on crime patterns. However, given that crime rates exhibit negative serial correlation over the scale of days to weeks (Jacob, Lefgren, and Moretti, 2007), the long-term impacts of climate change on crime may be considerably smaller than the short-term impacts. The only two previous studies of the effects of climate change on crime have used highly aggregate data and found mixed results (Anderson, Bushman, and Groom, 1997; Rotton and Cohn, 2003). To address this gap in the literature, in this chapter I use an unusually long and rich panel dataset to estimate the historical relationship between weather and crime. I then use this historical relationship to predict how climate change will impact the future prevalence of criminal activity in the United States, based on existing simulations of future weather under the IPCC's A1B scenario.⁹

⁹All climate projections cited in this chapter are based on the IPCC's A1B scenario. This scenario represents a future world with high rates of economic growth and substantial convergence between developing and developed economies, where rapid technological change is based on a balance of fossil-fuel intensive and non-fossil sources of energy (IPCC, 2000). A1B is a "middle-of-the-road" scenario that tends to produce emissions and climate results that are intermediate between high emissions scenarios such as A1FI and low emissions scenarios such as B1.

To support my analysis, I have constructed a panel dataset that includes monthly crime and weather data for 2,972 U.S. counties for the period from 1960 to 2009. My data on criminal activity is drawn from the U.S. Federal Bureau of Investigation's Uniform Crime Reporting (UCR) data. These data, which are based on monthly reports from 17,000 U.S. law enforcement agencies, tabulate offenses in nine major categories: murder, manslaughter, rape, aggravated assault, simple assault, robbery, burglary, larceny, and vehicle theft. I merge this data on crime rates with historical weather data from the U.S. National Climatic Data Center's Global Historical Climatology Network Daily (GHCN-Daily) dataset. The GHCN-Daily weather data include temperature and precipitation records from 75,000 weather stations worldwide that have been subjected to a set of quality assurance checks. After combining these two data sources, I generate a dataset with 1.46 million unique county-by-year-by-month observations.

To identify the effect of daily weather on monthly crime, I use a semi-parametric weather bin estimator (Deschenes and Greenstone, 2011) and control for county-by-month and county-by-year fixed effects. The weather bin variables measure the number of days per month spent in each of ten maximum daily temperature bins (<10 degrees F, 10-20 F, ..., 80-90 F, ≥ 90 F) and five daily precipitation bins (0 mm, 1-4 mm, 5-14 mm, 15-29 mm, and ≥ 30 mm). I regress monthly crime rates on these bin variables, controlling for extensive fixed effects that capture both average crime levels in each year-by-county set of observations and average monthly patterns of crime and weather within each county. Finally, I use the results from these regressions to predict crime rates under the weather patterns likely to be experienced in each decade between 2010 and 2099, based on projections of future U.S. climate drawn from two general circulation models.

My analysis makes two main contributions. First, I document a striking relationship between monthly weather patterns and crime rates. Across a variety of offenses, higher temperatures cause more crime. For most categories of violent crimes, this relationship appears approximately linear through the entire range of temperatures experienced

in in-sample counties. However, for property crimes such as burglary and larceny, the relationship between temperature and crime is highly non-linear, with a kink at approximately 40 degrees F. Above this cutoff, changes in temperature have little effect on crime rates.

Second, I develop the first detailed predictions of how climate change will affect patterns of criminal activity in the United States. My results suggest that climate change will cause crime rates to increase substantially. Under the IPCC's A1B climate scenario, the United States will experience an additional 35,000 murders, 216,000 cases of rape, 1.6 million aggravated assaults, 2.4 million simple assaults, 409,000 robberies, 3.1 million burglaries, 3.8 million cases of larceny, and 1.4 million cases of vehicle theft, compared to the total number of offenses that would have occurred between the years 2010 and 2099 in the absence of climate change.¹⁰ The present discounted value of the social costs of these climate-related crimes is between 20 and 68 billion dollars.

I am aware of only two previous empirical studies of the effects of climate change on crime in the United States: Anderson, Bushman, and Groom (1997), who study the relationship between annual average crimes rates and temperatures, using data for the United States as a whole; and Rotton and Cohn (2003), who perform a similar analysis based on state-level annual averages.¹¹ In contrast, my analysis is based on monthly crime data and daily weather data for 2,972 U.S. counties, and thus is much more likely to capture important aspects of the relationship between weather and crime. Additionally, as I describe in more detail below, reporting inconsistencies in the FBI's crime data add considerable measurement error to interannual comparisons of crime rates. Thus, by focusing on month-to-month changes in crime rates within a particular county and year, my analysis solves measurement error issues that have plagued previous work.

¹⁰For comparison, I assume that the total baseline number of crimes that will occur in the United States between 2010 and 2099 will be: 980,000 murders, 5.7 million cases of rape, 52 million aggravated assaults, 189 million simple assaults, 25 million robberies, 135 million burglaries, 429 million cases of larceny, and 72 million cases of vehicle theft. These totals are based on the assumption that crime rates during the next century will be similar to actual crime rates between 2000 and 2009.

¹¹In addition, Simister and Cooper (2005) provide graphical evidence on seasonal variation in U.S. crime rates.

The remainder of this chapter is organized as follows. Section 2.2 provides background on the relationship between weather and crime. Section 2.3 describes the primary data sources, and Section 2.4 discusses my empirical methodology. Section 2.5 presents my main findings on the relationship between climate change and crime. Section 2.6 discusses the results and Section 2.7 concludes.

2.2 Background on Weather and Crime

Researchers have proposed several hypotheses that explain why weather might affect crime (Cohn, 1990; Agnew, 2012). The first—that weather is a variable in the production function for crime—draws on Gary Becker’s canonical model of crime, in which individuals make decisions about whether to commit criminal acts based on rational consideration of the costs and benefits (Becker, 1968). In this model, weather conditions are an input that affects both the probability of successfully completing a crime and the probability of escaping undetected afterward (Jacob, Lefgren, and Moretti, 2007). For example, pleasant evening weather may increase the number of opportunities for mugging, and dark, rainy nights may increase the probability of successfully burglarizing a house without being detected.

A second explanation draws on a social interaction theory of crime. Glaeser, Sacerdote, and Scheinkman (1996) propose that the frequency of criminal acts is driven in large part by social interactions that occur during day-to-day life. Applied to weather, such a hypothesis implies that weather conditions that foster social interactions are likely to increase crime rates (Rotton and Cohn, 2003). For example, mild weather that encourages people to go shopping would also have the effect of increasing the frequency of property crimes such as larceny.

A third possible explanation draws on theories in which external conditions directly affect human judgment in ways that cause heightened aggression and loss of control (Baumeister and Heatherton, 1996; Card and Dahl, 2011). Experimental evidence

strongly suggests that ambient temperatures affect aggression (Anderson, 1989). For example, Baron and Bell (1976) assigned male subjects to receive a positive or negative evaluation from a confederate, and then gave them the opportunity to retaliate with an electric shock. They found that retaliation was highest when the experiment took place in a room with a high ambient temperature (92-95 degrees F), and that retaliation was still heightened even at more moderate temperatures (82-85 degrees F). Such studies imply that weather may directly influence people's psychological propensity to commit violent criminal acts.

Although using empirical data to distinguish between these hypotheses is difficult, there is considerable evidence that weather does affect criminal behavior (Cohn, 1990). Previous research on this topic has typically taken one of two empirical approaches. First, some studies have focused on measuring the short-term relationship between weather and crime, using hourly, daily, or weekly microdata (Bushman, Wang, and Anderson, 2005; Cohn and Rotton, 2000). For example, Brunsdon et al (2009) measure the impact of weather on disorderly conduct using hourly data on police calls in an urban area of the United Kingdom. They find that disorderly conduct increases with temperature and humidity but is unaffected by precipitation. However, interpreting these types of studies in the context of climate change is complicated by negative serial correlation in crime. In a large study using weekly data on crime and temperatures in 116 U.S. jurisdictions for the period 1996 to 2001, Jacob, Lefgren, and Moretti (2007) find that although rates of violent crime and property crime are elevated during weeks with hot weather, the effect is offset somewhat by lower than usual crime rates in the following weeks. This result suggests that understanding the cumulative impacts of climate change on crime may require working with data at a more aggregate time scale (e.g., months). Simister and Cooper (2005) conduct such an analysis for Los Angeles, using monthly assault and weather data from 1988 to 2002. Based on regressions that include linear and quadratic effects, they find evidence of a strong linear relationship

between assault and temperature.¹²

The second main empirical approach in the literature is to use yearly data to measure how weather affects crime at the national or state levels. For example, authors have examined the time series relationship between annual average crimes rates and average temperatures, for the United States as a whole (Anderson, Bushman, and Groom, 1997) and for a panel of states (Rotton and Cohn, 2003). These studies have found mixed results, possibly due to the a lack of geographic and spatial resolution in the crime and weather data. Another issue with this work is that U.S. aggregate crime statistics suffer from known quality issues, with different data sources implying considerably different trends in crime rates in the 1970s and 1980s (Levitt, 2004). As a result, analyses based on such geographically-aggregate annual data may face econometric issues with measurement error.

2.3 Data

2.3.1 Data Sources

The analysis for this chapter is based on an unusually long and rich panel dataset of monthly crime rates and weather for 2,972 counties in the 49 continental states (including the District of Columbia). The dataset covers the 50-year period from 1960 to 2009, and contains 1.46 million unique county-by-year-by-month observations. It is based on two primary sources: Uniform Crime Reporting data from the U.S. Federal Bureau of Investigation (FBI, 2011a), and Global Historical Climatology Network Daily weather data from the National Climatic Data Center (NCDC Climate Services Branch, 2011).

The FBI's Uniform Crime Reporting (UCR) data are the longest continuously-collected historical record of criminal activity in the United States. These data are based on

¹²Simister (2002) and Simister and Van de Vliert (2005) find evidence of an approximately linear relationship between temperature and murder in a similar analysis using national-level monthly data on weather and murder for Pakistan.

Table 2.1: Uniform Crime Reporting Offense Definitions

Offense	Definition
Murder	The willful (nonnegligent) killing of one human being by another.
Manslaughter	The killing of another person through gross negligence.
Rape	The carnal knowledge of a female forcibly and against her will.
Aggravated Assault	An unlawful attack by one person upon another for the purpose of inflicting severe or aggravated bodily injury.
Simple Assault	Assaults which do not involve the use of a firearm, knife, cutting instrument, or other dangerous weapon and in which the victim did not sustain serious or aggravated injuries.
Robbery	The taking or attempting to take anything of value from the care, custody, or control of a person or persons by force or threat of force or violence or by putting the victim in fear.
Burglary	The unlawful entry of a structure to commit a felony or a theft.
Larceny	The unlawful taking, carrying, leading, or riding away of property from the possession or constructive possession of another.
Vehicle Theft	The theft or attempted theft of a motor vehicle.

Source: FBI (2004).

monthly reports from approximately 17,000 local, county, city, university, state, and tribal law enforcement agencies. Although participation is voluntary and has increased over time, in 2010 the UCR data covered law enforcement agencies representing 97.4 percent of the U.S. population (FBI, 2011b). The data submitted by each agency each month include the number of reported offenses of murder, manslaughter, rape, aggravated assault, simple assault, robbery, burglary, larceny, and vehicle theft. Table 2.1 summarizes the definitions of each offense. In cases when a crime falls into more than one category, the FBI uses a “heirarchy rule” to assign the crime to the most serious offense category (FBI, 2004).

A central challenge in using the agency-level UCR data to construct monthly county-level crime rate time series is that the number of reported crimes in the data increases dramatically through the 1960s and 1970s, due both to changes in the number of agencies reporting and to more comprehensive reporting by individual agencies. Because of these limitations, developing a county-level time series that is consistent across years would involve considerable researcher judgment about non-reporting bias (and would most likely require excluding a large number of observations from the analysis). Although previous research on criminal behavior has made use of such annual UCR data

(e.g., Levitt, 1996), in this chapter I take a different approach in which I construct a time series that is consistent only across months within each county-by-year group of observations. To build this time series, I first drop any agency-by-year records in which an agency reported less than twelve months of data for that year.¹³ I then sum the total number of reported crimes by all remaining agencies in each county, by category of crime, to generate a county total for each month and year. Finally, using county population data from the U.S. Census (U.S. Census Bureau, 1978, 2004, 2011), I calculate the monthly crime rate per 100,000 persons, for each county-by-month observation. As I discuss below in the Methodology section, the fact that the number of reporting agencies differs across years within each county does not affect my regressions results, since I identify the effect of weather on crime using only variation in month-to-month weather and crime within a particular county and year (for which the set of reporting agencies is identical).

The second major component of my dataset is daily weather data taken from the U.S. National Climatic Data Center's (NCDC) Global Historical Climatology Network Daily data. The GHCN-Daily dataset is a compilation of weather station records drawn from a variety of sources, and includes about 75,000 weather stations worldwide (NCDC Climate Service Branch, 2011). The weather variables that I extract from the dataset are daily maximum temperature and daily precipitation. Unlike some other sources of weather data (e.g., the NCDC's Global Summary of the Day), the GHCN-Daily data are subjected a set of quality assurance reviews that include checking for weather data that are duplicated, weather data that exceed physical or climatological limits, consecutive datapoints that show excessive persistence or gaps, and data with inconsistencies internally or across neighboring stations.

Because the GHCN-Daily data report weather at a set of weather stations that are spaced irregularly across the United States, I use the station data to generate county

¹³I also drop agency-by-year records in which the agency reported data on a quarterly, bi-yearly, or yearly basis, rather than monthly. Most of these cases are agencies located in Florida or Alabama.

weather as follows. First, I create a set of grid points covering the entire United States, spaced approximately 5 miles apart. I then calculate the distance from each grid point to each weather station. Next, I estimate a county-level temperature signal using all stations within 50 miles of any grid point within a county. Finally, I adjust the absolute value of this signal so that it is equal to the average temperature reported at the stations closest to each county gridpoint. I calculate county-level precipitation using a similar procedure.

After combining the county-level crime and weather data, I take several final steps to clean the crime data. First, I drop all county-by-year records in which U.S. Census estimates indicate that the county had a population of fewer than 1,000 persons. Second, I drop all county-by-year records in which zero crimes were reported in all months, or in which weather data are missing for at least one month. Third, I eliminate outliers (almost all of which appear to be reporting errors) by dropping county-by-year observations in which the crime rate in any month is greater than twice the value of the 99th percentile crime rate for the entire sample. Finally, to minimize problems with heteroskedasticity in the data, I drop counties in which the mean crime rate is above the 99th percentile or below the 1st percentile for the entire sample. The resulting dataset includes 2,972 in-sample counties (out of the universe of 3,143 counties), with a total of 1.46 million unique county-by-year-by-month observations.

2.3.2 Summary Statistics

This section of the chapter presents summary statistics on crime and weather patterns in the United States. To illustrate how these patterns vary geographically, I divide the United States into four climate zones and then assign each county to a climate zone based on its long-term mean annual maximum daily temperature. The zones are <55 degrees F, 55 to 64 degrees F, 65 to 74 degrees F, and ≥ 75 degrees F. Panel (a) of Figure 2.1 shows a map of the climate zones. As expected, northern areas of the United States

are more likely to have cooler climates. For comparison, Panel (b) of the figure shows a map of county-level annual crime rates per 100,000 persons, for all crimes. The panel shows that crime rates are highest along the Eastern Seaboard, in the West, and in areas bordering the Great Lakes. However, there is no obvious cross-sectional relationship between the temperature zones and crime rates.¹⁴

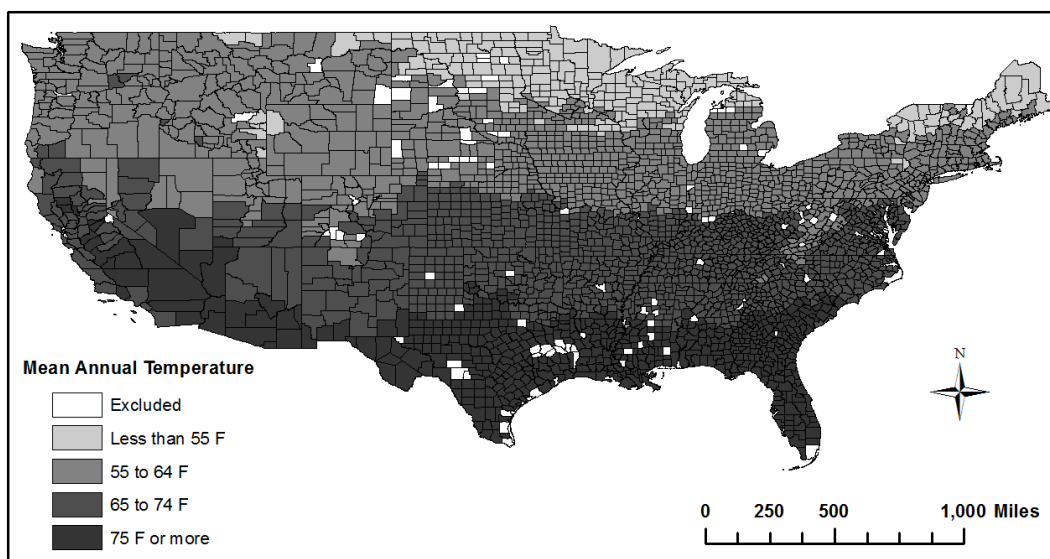
Table 2.2 summarizes basic characteristics of the crime and weather datasets, by climate zone. The first panel presents mean annual crime rates per 100,000 persons, by type of offense. The panel shows that some categories of crime, such as murder, manslaughter, rape, and robbery, are relatively uncommon. The three categories with the highest rates are larceny, burglary, and simple assault.

The second panel in Table 2.2 describes the annual distribution of daily temperatures and precipitation for in-sample counties. Unlike crime rates, these data show substantial variation across climate zones. For example, although counties in the coolest climate zone (<55 degrees F) have an average of only five days per year in which the maximum temperature exceeds 90 degrees F, counties in the warmest climate zone (≥ 75 degrees F) typically have 85 days per year with temperatures above 90 degrees F.

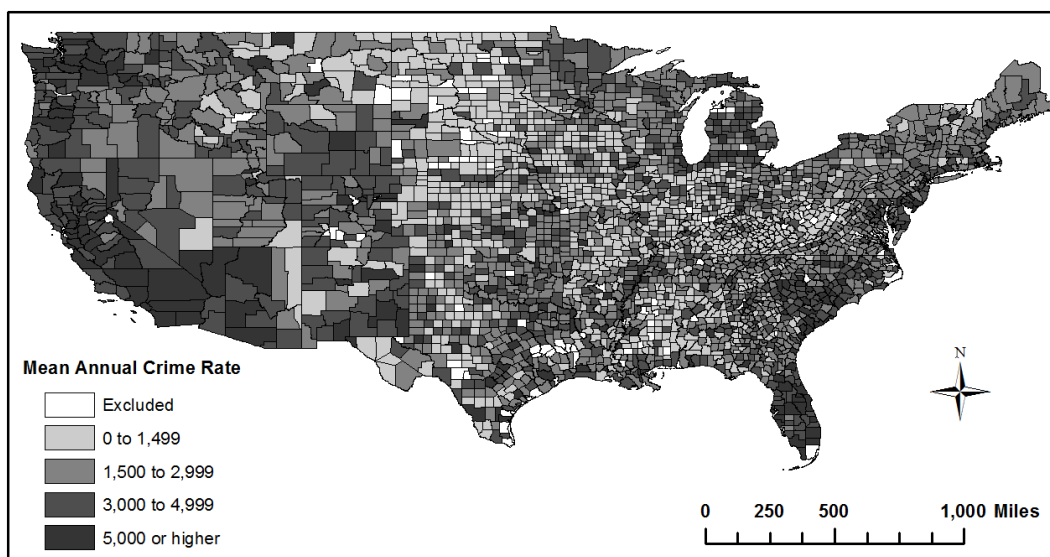
The final panel in Table 2.2 describes county socioeconomic characteristics. The panel shows that counties in cooler climate zones have fewer minorities and are more likely to be rural.

To illustrate the advantages and challenges of using the UCR crime reporting data, Figure 2.2 presents the time trend in crime rates for the nine major categories of offenses, by climate zone. Several main patterns are obvious from the data. First, crime rates increase dramatically between 1960 and 1980, in some cases by several hundred percent. Given the rapid and monotonic nature of this increase, it seems likely that it

¹⁴Given the many socioeconomic variables that influence crime, the absence of a strong visual cross-sectional relationship between temperatures and crime does not necessarily indicate the lack of a causal relationship. A cross-sectional analysis in the spirit of Mendelsohn, Nordhaus, and Shaw (1994) would have to control for other first-order determinants of crime (e.g., population density).



(a) Mean Annual Maximum Daily Temperature (F)



(b) Annual Crime Rate per 100,000 Persons (All Crimes)

Figure 2.1: Map of the Study Region

Note: Both panels show maps of all in-sample counties in the United States. The top panel depicts the mean annual maximum daily temperature, by county. The bottom panel depicts the annual number of all crimes per 100,000 persons, by county. All statistics are based on data from 1960-2009.

Table 2.2: Summary Statistics, by Climate Zone

	Mean Annual Maximum Daily Temperature			
	<55 F	55-64 F	65-74 F	≥75 F
Monthly Crime Rate (per 100,000 persons)				
Murder	0.1 (1.2)	0.2 (1.3)	0.4 (1.9)	0.6 (2.2)
Manslaughter	0.03 (0.53)	0.03 (0.50)	0.02 (0.45)	0.02 (0.42)
Rape	1.2 (3.5)	1.2 (3.2)	1.3 (3.3)	1.6 (3.5)
Aggravated Assault	6 (13)	9 (14)	15 (20)	20 (22)
Simple Assault	26 (40)	31 (39)	34 (48)	41 (52)
Robbery	1 (3)	2 (5)	3 (6)	4 (8)
Burglary	43 (48)	44 (45)	50 (46)	61 (53)
Larceny	111 (94)	117 (98)	107 (97)	125 (112)
Vehicle Theft	9 (12)	11 (15)	11 (15)	13 (18)
Annual Number of Days in Weather Bin				
Max Temp: <10 F	13 (11)	3 (4)	0 (1)	0 (0)
Max Temp: 10-19 F	20 (8)	7 (6)	1 (2)	0 (0)
Max Temp: 20-29 F	37 (9)	20 (11)	5 (5)	0 (1)
Max Temp: 30-39 F	52 (12)	44 (14)	18 (11)	3 (4)
Max Temp: 40-49 F	41 (11)	49 (14)	35 (12)	12 (8)
Max Temp: 50-59 F	40 (9)	49 (16)	50 (11)	31 (13)
Max Temp: 60-69 F	49 (11)	52 (13)	59 (13)	54 (13)
Max Temp: 70-79 F	64 (11)	64 (14)	68 (14)	77 (14)
Max Temp: 80-89 F	43 (14)	63 (18)	86 (18)	101 (28)
Max Temp: ≥90 F	6 (8)	15 (14)	43 (25)	86 (29)
Precip: 0 mm	179 (47)	165 (44)	196 (40)	215 (46)
Precip: 1-4 mm	143 (37)	149 (34)	111 (29)	95 (30)
Precip: 5-14 mm	32 (11)	37 (14)	36 (12)	33 (14)
Precip: 15-29 mm	9 (4)	11 (6)	15 (7)	15 (7)
Precip: ≥30 mm	2 (2)	3 (3)	6 (4)	8 (5)
County Characteristics				
Population	36,289 (51,843)	96,286 (246,799)	71,616 (316,927)	86,742 (237,757)
Pct White	97 (8)	96 (6)	87 (16)	78 (18)
Pct Female	50 (1)	51 (1)	51 (2)	51 (2)
Pct Ages 0-4	7 (2)	7 (1)	7 (1)	8 (1)
Pct Ages 5-19	25 (5)	25 (4)	25 (4)	26 (5)
Pct Ages 65-up	15 (4)	14 (4)	13 (4)	13 (5)
Pct Metro Center	2 (15)	7 (26)	6 (24)	4 (20)
Pct Metropolitan	14 (35)	24 (43)	22 (41)	27 (44)
Pct Urban	54 (50)	50 (50)	50 (50)	56 (50)
Pct Rural	30 (46)	19 (40)	22 (41)	13 (33)
Counties	209	1,092	1,141	530
Complete County Years	9,183	48,288	45,544	18,954
County Month Obs.	110,196	579,456	546,528	227,448

Note: The table shows mean crime rates, weather conditions, and socioeconomic characteristics for all in-sample counties for the years 1960-2009. Numbers in parentheses indicate standard deviations. Results are presented separately for counties in each of four climate zones, based on mean annual maximum daily temperature.

is driven by increased reporting of crimes, rather than by changes in underlying criminal behavior. Second, trends across climate zones appear broadly similar, although there is some heterogeneity in absolute levels. Finally, there is strong evidence of high frequency variation in crime rates due to seasonality.

Figures 2.3 and 2.4 present additional evidence on seasonality in the data. Figure 2.3 shows the mean value of daily maximum temperature and daily precipitation, by climate zone and month. The figure shows strong seasonal patterns in all climate zones, for all variables. Seasonal variation is largest in the coolest climate zone (<55 degrees F), where the mean temperature difference between January and July is 60 degrees. For comparison, the seasonal temperature difference between January and July in the warmest climate zone (≥ 75 degrees F) is about 35 degrees.

Figure 2.4 presents similar graphs illustrating how crime rates vary by climate zone and month. The figure shows that all categories of crime show evidence of seasonality, although the degree of seasonal variation varies widely across crimes. A few categories of crime, particularly murder and robbery, show only modest seasonal variation. Other categories, such as rape, assault, and non-violent property crimes, exhibit strong seasonality. Additionally, the relationship between seasonality and crime rates varies across climate zones and type of crimes. For example, larceny and burglary show more pronounced seasonal variation in cooler climate zones, whereas robbery shows somewhat more seasonality in warmer climates.

2.4 Methodology

The summary statistics from the previous section show a strong correlation between monthly weather and crime rates. In this section I develop a causal econometric model of this relationship. Specifically, I model crime in month m of year y in county i as

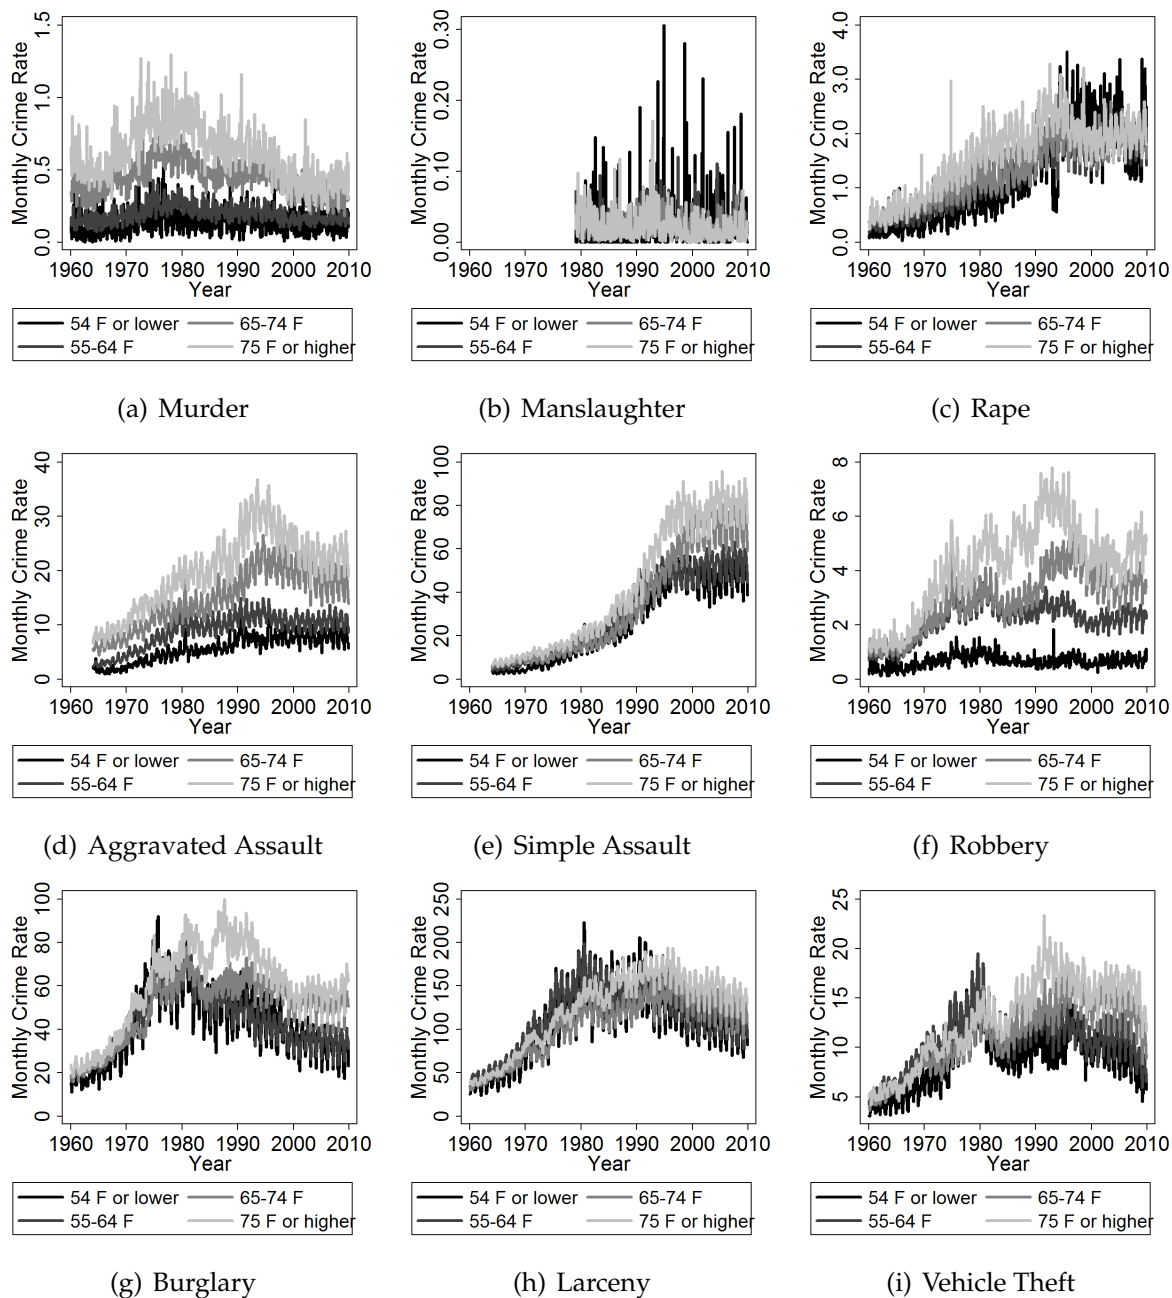


Figure 2.2: Crime Rate Trends, by Climate Zone

Note: Each panel shows the mean crime rate across counties within each climate zone, by year and month. The crime rate variables represent the monthly number of crimes per 100,000 persons.

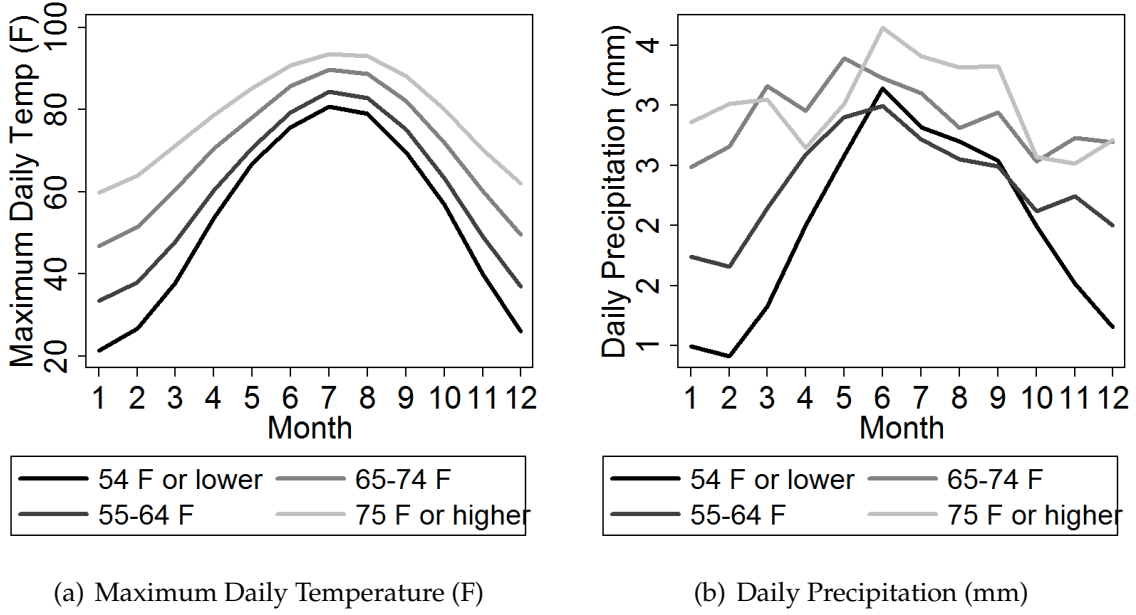


Figure 2.3: Seasonal Weather Patterns, by Climate Zone

Note: Each panel shows mean weather across counties within each climate zone, by month, for the period from 1960 to 2009.

follows:

$$\begin{aligned}
 C_{iym} = & \sum_{j=1}^{10} \alpha_0^j T_{iym}^j + \sum_{k=1}^5 \beta_0^k P_{iym}^k \\
 & + \sum_{j=1}^{10} \alpha_1^j T_{iym-1}^j + \sum_{k=1}^5 \beta_1^k P_{iym-1}^k \\
 & + \phi_{im} + \theta_{iy} + \epsilon_{iym}
 \end{aligned} \tag{2.1}$$

In this equation, C_{iym} represents the monthly crime rate per 100,000 residents, ϕ_{im} is a county-by-month fixed effect, θ_{iy} is a county-by-year fixed effect, and ϵ_{iym} is a zero-mean error term. Following Deschenes and Greenstone (2011), I model the daily distribution of temperatures within a month using ten bin variables: <10 F, 10-19 F, 20-29 F, 30-39 F, 40-49 F, 50-59 F, 60-69 F, 70-79 F, 80-89 F, and ≥ 90 F. For example, the variable T_{iym}^j represents the number of days in month m of year y in county c in which the temperature fell into temperature bin j . I use a similar convention for the precipitation variables P_{iym}^k , with five bins: 0 mm, 1-4 mm, 5-14 mm, 15-29 mm, and ≥ 30 mm. Because of the possibility that changes in crime rates due to weather shocks may

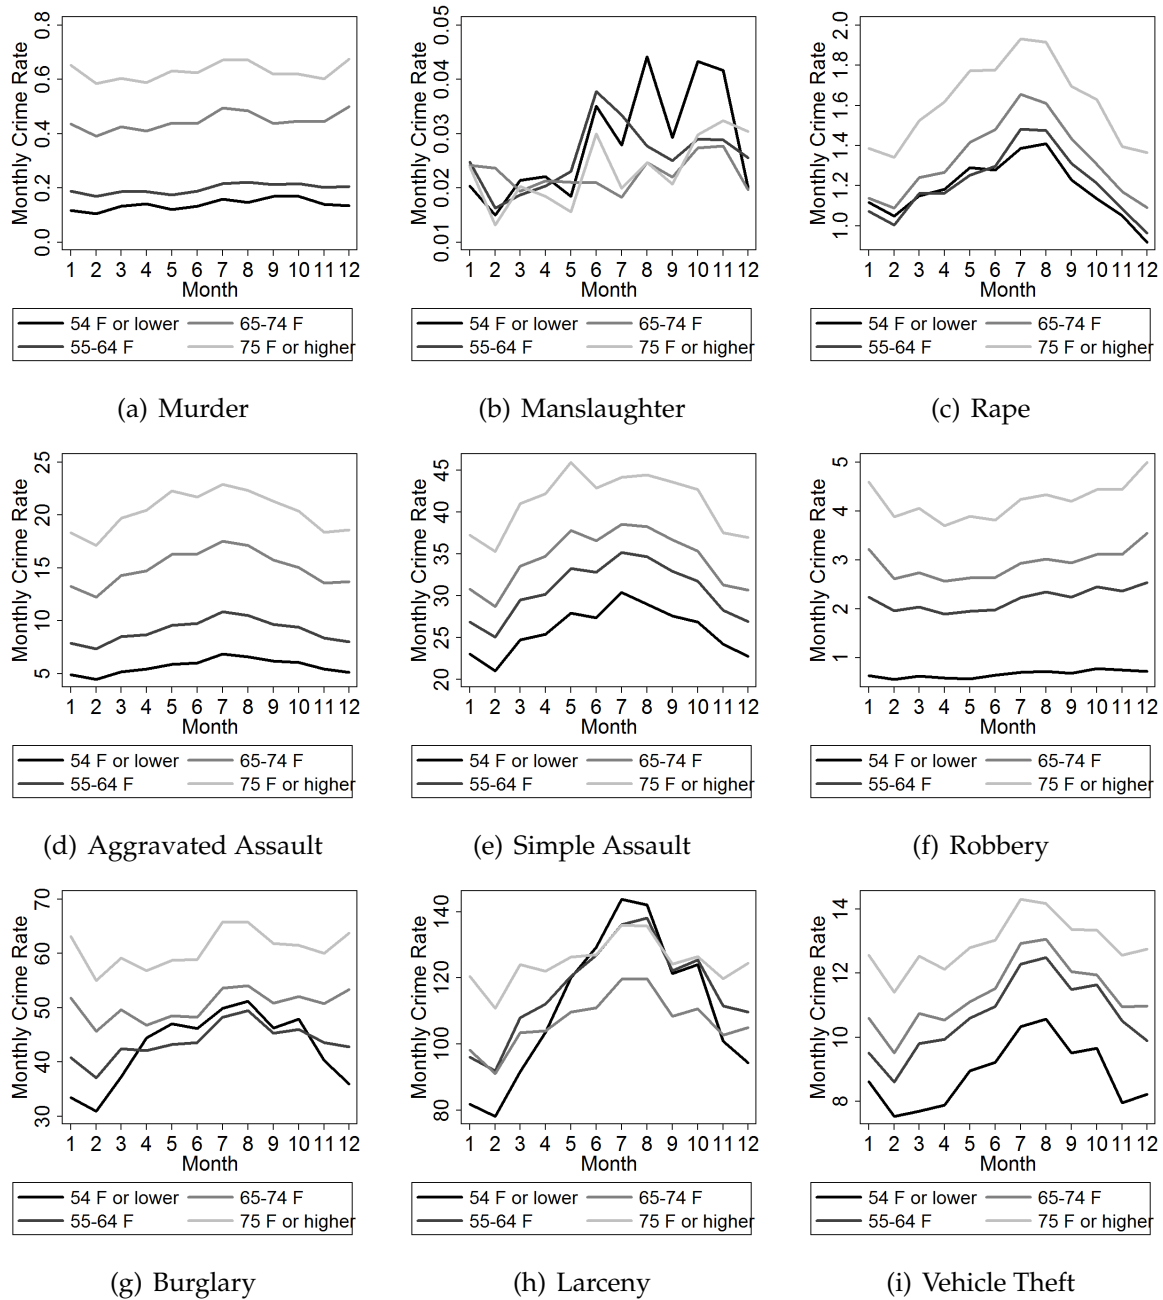


Figure 2.4: Seasonal Crime Rate Trends, by Climate Zone

Note: Each panel shows the mean crime rate across counties within each climate zone, by month. The crime rate variables represent the monthly number of crimes per 100,000 persons.

exhibit negative serial correlation (Jacob, Lefgren, and Moretti, 2007), I also include a one month lag of each temperature and precipitation bin variable. Furthermore, because weather patterns in a particular month are highly correlated between adjacent geographic areas, I cluster all standard errors at the year-by-month level. I also weight each county-by-month-by-year observation by the county population in that year.

Equation (2.1) includes several features designed to address issues that have been problematic in previous analysis of the effect of weather on criminal behavior. First, by using a semi-parametric specification for weather, I avoid imposing structural assumptions on the relationship between weather and crime. Previous analyses have used as independent variables mean weekly temperature and precipitation (Jacob, Lefgren, and Moretti, 2007) or mean yearly temperature and temperature squared (Rotton and Cohn, 2003). These specifications assume that weather has a linear or quadratic effect on crime—which, as the results from this chapter show, may fail to capture important features of the relationship.

Second, Equation (2.1) includes an extraordinarily comprehensive set of fixed effects. In addition to including dummy variables for typical monthly patterns in weather and crime with each county, I include dummy variables that capture the average crime rate and weather conditions in each county-by-year set of observations. In other words, my identification strategy is based on only the residual variation in crime and weather remaining between months within a particular county and year, after controlling for average monthly crime levels.

The motivation for this extensive set of fixed effects is related to the quality of the FBI's crime data. As Figure 2.2 illustrates, the UCR crime data exhibit strong interannual trends that appear to be driven at least partially by differences in reporting. Examination of the microdata shows that at the level of individual counties, these trends are exacerbated, with crime rates in many counties jumping wildly from year to year as the set of reporting agencies changes over time. In the two previous national stud-

ies of crime and climate change (Anderson, Bushman, and Groom, 1997; Rotton and Cohn, 2003), the authors addressed this problem by modeling annual changes in aggregate national or state crime rates as an autoregressive process. Because this approach is not an entirely satisfactory method for dealing with measurement error in the dependent variable, I choose an alternative methodology that requires no consistency in reporting between years. Instead, as discussed in the data section, I construct monthly crime rates within each county-by-year by aggregating the total number of reported crimes each month only for agencies that reported twelve complete months of data for that year. Thus, although the set of reporting agencies within each county changes between years, making interannual comparisons invalid except under very strong assumptions, an identical set of agencies report for each month within a particular year. Thus, the identifying assumption for my analysis is that after controlling for county-by-year and county-by-month fixed effects, differences in weather and crime between months within a county represent the true effect of weather on crime.

2.5 Results

This section presents the main results from the analysis.

2.5.1 Weather and Crime Rates

I begin by presenting the regression results from estimating Equation (2.1). Because of the large number of coefficients, the results are easiest to understand using a graphical approach. For example, Figure 2.5 plots the regression coefficients on the temperature and lagged temperature bin variables. In each subfigure, the horizontal axis represents the daily maximum temperature bins, and the vertical axis represents the coefficient, with units of number of crimes per 100,000 persons per month. The figure shows that across all types of crime, higher temperatures cause statistically significant increases in crime rates. As an illustration, compared to a day in the 60-69 degrees F bin, an extra

day in the 30-39 degrees F bin leads to 0.002 fewer murders, 0.08 fewer aggravated assaults, and 1.1 fewer larcenies, per 100,000 persons per month. In comparison, the mean monthly crime rates for these three offenses are .35 cases of murder, 14 cases of aggravated assault, and 114 cases of larceny. Although the estimated coefficients appear small relative to mean crime rates, the coefficients represent the effect of only a single day of weather per month, and in aggregate imply substantial effects. For example, in a spring month with 10 unusually cold days (in the 30-39 degrees F bin), crime rates for these three offenses would be approximately seven to ten percent lower than crime rates in a spring month with 10 unusually warm days (in the 60-69 degrees F bin).

Figure 2.5 also shows the significant non-linearities in the effect of temperatures on crime. These non-linear effects are most apparent for property crimes such as burglary and larceny. For bins below 40 degrees F, increases in temperature have a strong positive effect on the number of burglaries and larcenies reported. However, above 40 degrees F, increases in temperature have little or no effect on these crimes. The degree of nonlinearity varies by offense, with violent crimes tending to have a much more linear relationship through the entire range of temperatures.

In addition to showing the effect of current monthly temperatures on current monthly crime, Figure 2.5 also presents coefficients and confidence intervals for the effect of lagged temperature from the previous month. For most offenses, the coefficients on lagged temperatures are close to zero and not statistically significant. Thus, unlike Jacob, Lefgren, and Moretti (2007), who find a significant and opposite coefficient on lagged weekly temperatures that dampens the effect of weather on weekly crime, I conclude that at the monthly level, there is little evidence that weather has a lagged effect on crime patterns.

Table 2.3 presents complete regression results from estimating Equation (2.1), including the results for the precipitation bin variables. The table shows that the effects of precipitation on crime rates vary by offense. Although precipitation causes statistically

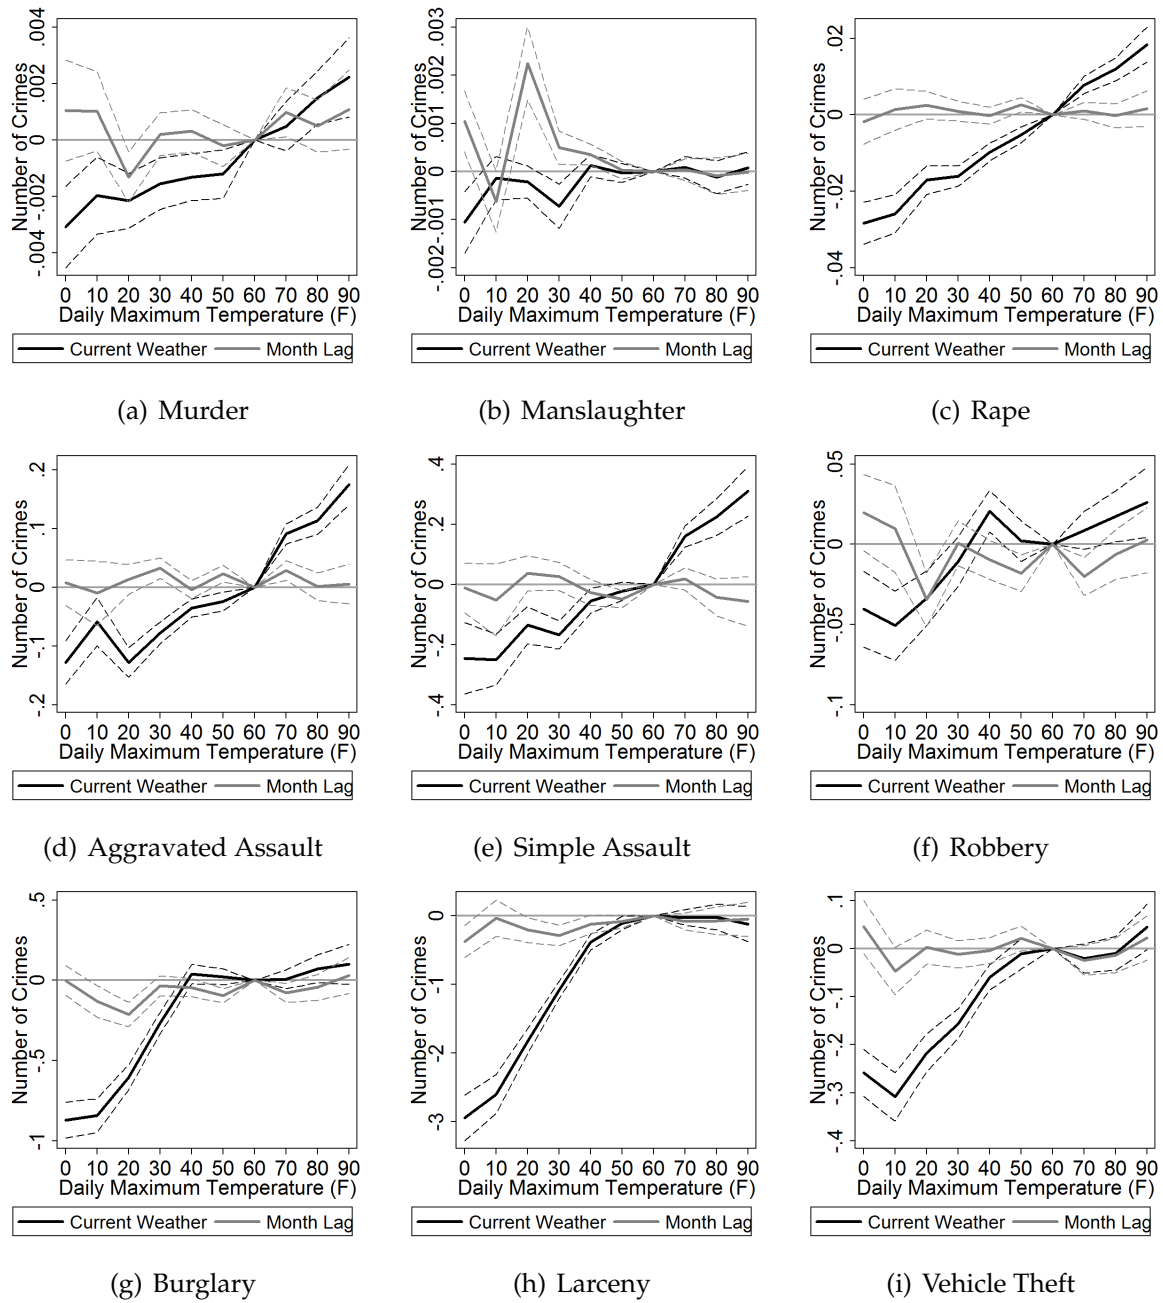


Figure 2.5: The Effect of Daily Maximum Temperature on Monthly Crime

Note: Each figure shows coefficients from a regression of the monthly crime rate per 100,000 persons on a semiparametric set of weather bin variables. The solid black line represent the effect of current weather; the solid gray line represents the lagged effect of the previous month's weather. Dashed lines represent 95 percent confidence intervals for the estimated coefficients. All coefficients are relative to one day in the 60 to 70 degrees F bin.

significant decreases in larceny, the opposite is true for vehicle theft: more vehicles are stolen in months with many rainy days.

One key question about the analytical approach used in this chapter is whether one month is a sufficiently long time period to account for any lagged impacts of weather on crime. Although the mostly insignificant coefficients on the one-month lag of the weather bins suggest that this is the case, I also conduct sensitivity analyses in which I run regressions using data that have been aggregated to quarterly and half-year time periods. Figure 2.12 in the Appendix shows the results of this analysis. Although regressions results based on more aggregate time periods are noisier than the results based on month-long time periods, the estimated coefficients from the three types of regressions are generally similar. For example, although the relationship between temperature and crime rates for aggravated and simple assault appears somewhat weaker based on the quarterly and half-year data, the effect of temperature on burglary and larceny is even stronger in the quarterly and half-year data. Overall, the figure suggests that a one-month aggregation period is sufficient to account for “harvesting” that might occur as a result of negative serial correlation in crime rates.

In addition to the main specifications presented in Figure 2.5 and Table 2.3, I have run a variety of other sensitivity analyses in which I allow the coefficients on the weather bin variables to vary by climate zone, monthly mean temperature, and decade. The results from these analyses are qualitatively similar to the main specification presented here, and are presented in the appendix.

2.5.2 Climate Change and Crime Rates

To assess how climate change is likely to affect crime rates in the United States, I combine the regression estimates from the previous section with data on simulated U.S. weather conditions for the time period from 2010 to 2099. These simulations are based on the IPCC’s A1B scenario, a “middle-of-the-road” climate change scenario

Table 2.3: Maximum Daily Temperature and Monthly Crime

	Murder	Mansltr	Rape	Agg Asslt	Smp Asslt	Robbery	Burglary	Larceny	Veh Theft
Temp: < 10 F	-0.003*** (0.001)	-0.001** (0.000)	-0.028*** (0.003)	-0.128*** (0.019)	-0.245*** (0.061)	-0.040*** (0.012)	-0.872*** (0.057)	-2.945*** (0.169)	-0.258*** (0.025)
Temp: 10-19 F	-0.002** (0.001)	-0.000 (0.000)	-0.026*** (0.003)	-0.059** (0.021)	-0.250*** (0.043)	-0.051*** (0.011)	-0.843*** (0.054)	-2.597*** (0.145)	-0.308*** (0.026)
Temp: 20-29 F	-0.002*** (0.000)	-0.000 (0.000)	-0.017*** (0.002)	-0.128*** (0.013)	-0.136*** (0.031)	-0.034*** (0.009)	-0.604*** (0.040)	-1.828*** (0.094)	-0.217*** (0.020)
Temp: 30-39 F	-0.002*** (0.000)	-0.001** (0.000)	-0.016*** (0.001)	-0.078*** (0.009)	-0.168*** (0.024)	-0.011 (0.008)	-0.267*** (0.034)	-1.081*** (0.064)	-0.157*** (0.016)
Temp: 40-49 F	-0.001** (0.000)	0.000 (0.000)	-0.010*** (0.001)	-0.035*** (0.008)	-0.055** (0.021)	0.021** (0.007)	0.039 (0.031)	-0.385*** (0.059)	-0.060*** (0.013)
Temp: 50-59 F	-0.001** (0.000)	-0.000 (0.000)	-0.005*** (0.001)	-0.024** (0.008)	-0.023 (0.016)	0.002 (0.006)	0.021 (0.025)	-0.108* (0.050)	-0.011 (0.016)
Temp: 70-79 F	0.000 (0.000)	0.000 (0.000)	0.008*** (0.001)	0.091*** (0.009)	0.160*** (0.018)	0.009 (0.006)	0.006 (0.030)	-0.025 (0.058)	-0.020 (0.015)
Temp: 80-89 F	0.002** (0.000)	-0.000 (0.000)	0.012*** (0.002)	0.113*** (0.012)	0.224*** (0.031)	0.017* (0.008)	0.072 (0.044)	-0.019 (0.096)	-0.010 (0.018)
Temp: ≥90 F	0.002** (0.001)	0.000 (0.000)	0.018*** (0.002)	0.175*** (0.018)	0.310*** (0.042)	0.026* (0.011)	0.100 (0.064)	-0.121 (0.129)	0.045 (0.024)
Precip: 1-4 mm	0.001** (0.000)	0.000 (0.000)	0.002** (0.001)	0.011* (0.006)	-0.005 (0.013)	0.002 (0.004)	0.013 (0.020)	0.041 (0.046)	0.052*** (0.009)
Precip: 5-14 mm	0.000 (0.000)	0.000 (0.000)	0.001 (0.001)	-0.020* (0.009)	-0.010 (0.021)	-0.004 (0.007)	-0.040 (0.037)	-0.405*** (0.086)	0.026 (0.019)
Precip: 15-29 mm	0.002* (0.001)	0.001* (0.000)	-0.001 (0.002)	-0.004 (0.013)	0.016 (0.032)	0.009 (0.010)	-0.059 (0.043)	-0.413*** (0.103)	0.136*** (0.027)
Precip: ≥30 mm	-0.001 (0.001)	-0.000 (0.000)	-0.005 (0.003)	-0.070** (0.022)	-0.029 (0.041)	0.046** (0.015)	0.119 (0.067)	-0.920*** (0.163)	0.136** (0.047)
Lag T: < 10 F	0.001 (0.001)	0.001** (0.000)	-0.002 (0.003)	0.008 (0.020)	-0.012 (0.042)	0.020 (0.012)	-0.001 (0.047)	-0.377** (0.119)	0.046 (0.028)
Lag T: 10-19 F	0.001 (0.001)	-0.001 (0.000)	0.001 (0.003)	-0.010 (0.028)	-0.052 (0.061)	0.010 (0.014)	-0.131** (0.050)	-0.037 (0.135)	-0.047 (0.025)
Lag T: 20-29 F	-0.001** (0.000)	0.002*** (0.000)	0.003 (0.002)	0.013 (0.013)	0.037 (0.030)	-0.035*** (0.009)	-0.214*** (0.039)	-0.209* (0.092)	0.003 (0.018)
Lag T: 30-39 F	0.000 (0.000)	0.000** (0.000)	0.001 (0.001)	0.033*** (0.009)	0.026 (0.024)	0.001 (0.007)	-0.034 (0.032)	-0.289*** (0.076)	-0.012 (0.014)
Lag T: 40-49 F	0.000 (0.000)	0.000** (0.000)	-0.000 (0.001)	-0.004 (0.008)	-0.027 (0.022)	-0.010 (0.006)	-0.046 (0.029)	-0.126 (0.068)	-0.004 (0.014)
Lag T: 50-59 F	-0.000 (0.000)	0.000 (0.000)	0.003** (0.001)	0.023** (0.007)	-0.049*** (0.014)	-0.018** (0.006)	-0.097*** (0.023)	-0.086 (0.047)	0.021 (0.013)
Lag T: 70-79 F	0.001* (0.000)	0.000 (0.000)	0.001 (0.001)	0.028*** (0.008)	0.018 (0.019)	-0.020*** (0.006)	-0.079* (0.031)	-0.086 (0.063)	-0.025 (0.016)
Lag T: 80-89 F	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.002)	0.001 (0.012)	-0.043 (0.031)	-0.006 (0.008)	-0.044 (0.041)	-0.074 (0.103)	-0.014 (0.018)
Lag T: ≥90 F	0.001 (0.001)	-0.000 (0.000)	0.002 (0.002)	0.005 (0.017)	-0.057 (0.042)	0.003 (0.010)	0.031 (0.057)	-0.051 (0.127)	0.022 (0.024)
Lag P: 1-4 mm	0.001*** (0.000)	0.000* (0.000)	0.000 (0.001)	0.011 (0.006)	0.019 (0.014)	0.022*** (0.004)	0.051* (0.021)	-0.012 (0.050)	0.028** (0.010)
Lag P: 5-14 mm	-0.000 (0.000)	-0.000* (0.000)	-0.001 (0.002)	0.047*** (0.009)	0.062** (0.022)	-0.003 (0.007)	-0.062 (0.033)	0.017 (0.078)	0.031 (0.016)
Lag P: 15-29 mm	-0.001 (0.001)	-0.000 (0.000)	0.007*** (0.002)	0.024 (0.014)	0.082* (0.032)	0.008 (0.011)	0.021 (0.047)	-0.034 (0.108)	0.048 (0.025)
Lag P: ≥30 mm	0.004** (0.001)	0.000 (0.000)	0.006 (0.003)	0.084*** (0.019)	0.168*** (0.044)	-0.004 (0.017)	-0.023 (0.068)	-0.269 (0.141)	0.055 (0.042)
Observations	1,315,325	848,837	1,315,325	1,237,676	1,237,676	1,315,325	1,315,325	1,315,325	1,315,325
Clusters	539	341	539	506	506	539	539	539	539
R-squared	0.000	0.001	0.002	0.005	0.005	0.004	0.022	0.053	0.011

Note: Each observation represents a unique county-by-year-by-month. The dependent variable in all regressions is the monthly crime rate per 100,000 persons, with each column representing a different type of crime. The independent variables are the number of days per month that daily weather fell into the specified range, with 60-69 F as the omitted temperature bin and 0 mm as the omitted precipitation bin. All regressions control for county-by-year and county-by-month fixed effects. The county-by-year fixed effects are removed by long differencing relative to January of each county-by-year group of twelve months, and the dropping all (zeroed-out) January observations. County-by-month fixed effects are removed by de-meaning. All regressions are clustered by year-by-month, and weighted by county population.

that assumes eventual stabilization of atmospheric CO₂ levels at 720 ppm (IPCC, 2000, 2007). I use predictions from two general circulation models: the U.K. Hadley Centre's HadCM3 climate model, and the U.S. National Center for Atmospheric Research's CCSM3 climate model. The predictions, which are available from an archive maintained by the World Climate Research Programme's Coupled Model Intercomparison Project Phase 3 (CMIP3), have an interpolated resolution of two degrees of latitude by two degrees of longitude (WCRP, 2007; Maurer et al, 2007).

To use these data to estimate how climate change is likely to affect crime rates in each county in my analysis, I follow several steps. First, I use the HadCM3 and CCSM3 projections to calculate average predicted monthly temperature and precipitation for each decade between 2000 and 2099, for each two degree-by-two degree gridpoint. Taking the average monthly values for 2000-2009 as a baseline, I then calculate the absolute change in mean monthly temperature and the proportional change in mean monthly precipitation at each gridpoint for each subsequent decade, relative to 2000-2009. I then assign each U.S. county a predicted change in temperature and precipitation for each future decade and month, based on the changes predicted at the closest HadCM3 and CCSM3 gridpoint.

Next, I use these predicted changes to generate a simulated distribution of days across temperature and precipitation bins for each of the nine decades starting with 2010-2019 and ending with 2090-2099, for each month and county. I begin with the actual record of temperatures for each day, month, and county between 2000 and 2009. For each decade, I then add the predicted absolute change in monthly temperature to each daily temperature, by month and county, yielding a new predicted record of daily temperatures. I generate simulated precipitation data by multiplying the daily precipitation values by the proportional change in predicted precipitation. I then use these counterfactual weather records to calculate the mean number of days that will fall into each temperature and precipitation bin in each county and month, in each future decade. I conduct this procedure separately for the HadCM3 and CCSM3 predictions.

Finally, to predict how the projected change in weather will affect crime rates in each county, month, and decade, I combine the climate projections with the regression coefficients estimated in the previous section. I estimate the change in crime rates ΔC_{idm} in county i , decade d , and month m using the following formula:

$$\Delta C_{idm} = 10 \cdot \left[\sum_{j=1}^{10} \alpha_0^j (\bar{T}_{i,d,m}^j - \bar{T}_{i,2000,m}^j) + \sum_{k=1}^5 \beta_0^k (\bar{P}_{i,d,m}^k - \bar{P}_{i,2000,m}^k) + \sum_{j=1}^{10} \alpha_1^j (\bar{T}_{i,d,m-1}^j - \bar{T}_{i,2000,m-1}^j) + \sum_{k=1}^5 \beta_1^k (\bar{P}_{i,d,m-1}^k - \bar{P}_{i,2000,m-1}^k) \right] \quad (2.2)$$

where $\bar{T}_{i,d,m}^j$ refers to the mean number of days per month in which the simulated temperature in month m in county c in decade fell into temperature bin j . The predicted precipitation variable $\bar{P}_{i,d,m-1}^k$ is defined similarly. The variables $\bar{T}_{i,2000,m}^j$ and $\bar{P}_{i,2000,m-1}^k$ refer to the actual distribution of days across temperature and precipitation bins during the decade from 2000 to 2009. I multiply the entire expression on the right-hand side of the equation by ten to account for the number of years in each decade.¹⁵

Before discussing the results of this analysis, I present information on the changes in weather predicted by the CCSM3 and HadCM3 models. Figure 2.6 shows the distribution of temperature and precipitation across bins for three scenarios: the actual weather patterns observed between 2000 and 2009, the weather patterns predicted for 2090 to 2099 by the CCSM3 model, and the weather patterns predicted for 2090 to 2099 by the HadCM3 model. The figure shows that the baseline (2000-2009) maximum daily temperature distribution is heavily left-skewed. As a result, the increases in temperatures predicted by the CCSM3 and HadCM3 models lead to a sharp increase in the number of days that are predicted to fall into the highest daily maximum temperature bin (≥ 90 degrees F). The number of days in all other bins decreases under both sets of model predictions.

Table 2.4 shows the predicted impacts of climate change on crime in the United States. The first two columns of the table present estimates of the additional number of crimes

¹⁵Note that I also adjust ΔC_{idm} to account for the actual county population.

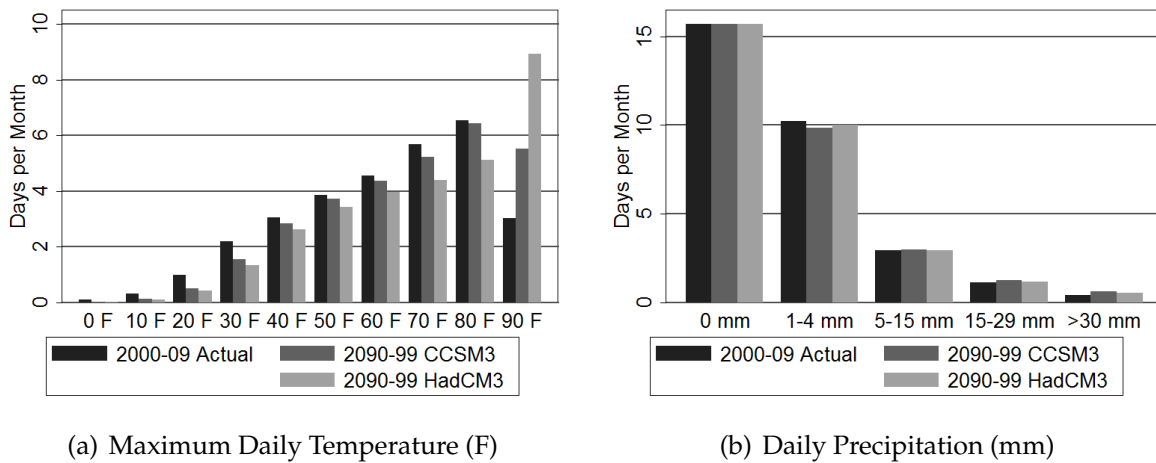


Figure 2.6: Distribution of Daily Weather, by Scenario

Note: Each panel shows the number of days per month that fall into the specified weather bin.

that will occur between 2010 and 2099, compared to the number that would have occurred in the absence of climate change. The table shows that under both climate models, climate change will cause a strikingly large number of crimes during the next century. For example, under the HadCM3 model, there will be an additional 35,000 murders, 216,000 cases of rape, 1.6 million aggravated assaults, 2.4 million simple assaults, 409,000 robberies, 3.1 million burglaries, 3.8 million cases of larceny, and 1.4 million cases of vehicle theft. All of these changes are significant at the five percent threshold. The only category of crime that is expected to decrease is manslaughter, but the expected change is less than 1,000 crimes and is not significantly different from zero. Compared to the baseline number of crimes expected to occur during this 90 year period in the absence of climate change, these figures represent a 3.6% increase in murder, a 2.7% decrease in manslaughter, a 3.8% increase in cases of rape, a 3.1% increase in aggravated assault, a 1.3% increase in simple assault, a 1.6% increase in robbery, a 2.3% increase in burglary, a 0.9% increase in cases of larceny, and a 1.9% increase in cases of vehicle theft.

Because these offenses occur over a 90 year time period and include a variety of types of crimes, it is useful to aggregate them into a social cost metric. I estimate the social costs of future changes in crime using the following valuations per offense: \$5,000,000

Table 2.4: The Predicted Impact of Climate Change on Crime

Crime	Number of Additional Crimes		Social Cost (billions)			
	HadCM3	CCSM3	HadCM3		CCSM3	
			3%	6%	3%	6%
Murder	34,853 (7,890)	25,290 (5,091)	42.6 (9.4)	17.4 (3.7)	30.8 (6.1)	11.9 (2.3)
Manslaughter	-1,066 (2,057)	-1,590 (1,470)	-1.5 (2.5)	-0.6 (1.0)	-2.2 (1.8)	-1.0 (0.7)
Rape	216,258 (29,160)	160,488 (19,034)	2.2 (0.3)	0.9 (0.1)	1.6 (0.2)	0.6 (0.1)
Aggravated Assault	1,582,743 (203,959)	1,154,021 (134,190)	7.5 (1.0)	3.0 (0.4)	5.5 (0.6)	2.1 (0.3)
Simple Assault	2,428,195 (549,913)	1,862,546 (370,464)	2.9 (0.7)	1.2 (0.3)	2.2 (0.5)	0.9 (0.2)
Robbery	409,323 (132,519)	295,152 (90,996)	2.2 (0.7)	0.9 (0.3)	1.6 (0.5)	0.6 (0.2)
Burglary	3,079,435 (712,709)	2,563,501 (481,084)	4.8 (1.1)	2.0 (0.4)	4.1 (0.7)	1.7 (0.3)
Larceny	3,806,456 (1,511,895)	4,180,251 (1,037,660)	3.5 (1.3)	1.5 (0.6)	3.9 (0.9)	1.7 (0.5)
Vehicle Theft	1,427,532 (259,308)	1,099,411 (180,888)	3.6 (0.6)	1.4 (0.3)	3.0 (0.5)	1.2 (0.2)
Total			67.8	27.7	50.5	19.7

Note: The “Number of Additional Crimes” columns represent the number of additional crimes that will occur due to climate change, relative to the number that would occur if temperatures and precipitation stayed at the 2000-2009 averages. The “HadCM3” and “CCSM3” columns show results based on different climate models. The “Social Cost” columns present the present value of the social cost of the additional crimes that will occur due to climate change. Future costs are discounted using two alternative discount rates: 3% and 6%.

for murder and manslaughter, \$41,247 for rape, \$19,537 for aggravated assault, \$4,884 for simple assault, \$21,398 for robbery, \$6,170 for burglary, \$3,523 for larceny, and \$10,534 for motor vehicle theft. The social cost estimates for murder and manslaughter are based on the value of a statistical life (VSL) for workers in U.S. labor markets. Estimates of VSL typically range between \$4 million and \$9 million (Viscusi and Aldy, 2003), and I choose \$5 million as a plausible value. Estimates of the social cost of the remaining offences are drawn from a review article by McCollister, French, and Fang (2010). These valuations represent the tangible costs of crime, including medical expenses, cash losses, property theft or damage, lost earnings because of injury, other victimization-related consequences, criminal justice system costs, and career crime costs.¹⁶ Although McCollister, French, and Fang also report intangible costs of crime (such as pain and suffering), I exclude these estimates because they are based on jury awards that may not accurately reflect individuals' actual willingness to pay to avoid victimization. Exclusion of this category of costs may bias my estimates of the social cost downward.

The right-hand side of Table 2.4 shows estimates of the social cost of the climate-related crime that is likely to occur between 2010 and 2099. Including all offenses, the social costs of this crime are between \$20 billion and \$68 billion. Because of the high value of a statistical life, the costs of future murders are by far the largest component of total social cost. As the table demonstrates, the estimates are somewhat sensitive to the choice of climate model and discount rate. For example, based on the HadCM3 model and a three percent discount rate, the present discounted cost of climate-related murder over the next ninety years is \$42 billion. Based on the CCSM3 model and a six percent discount rate, the social cost of murder is only \$12 billion.

One fact that is not apparent from Table 2.4 is that the impacts of climate change on crime are not uniformly distributed across the United States. To investigate distribu-

¹⁶McCollister, French, and Fang (2010) do not report estimates of the social cost of simple assault. For the purposes of this analysis, I value each case of simple assault at 25 percent of the cost of a case of aggravated assault.

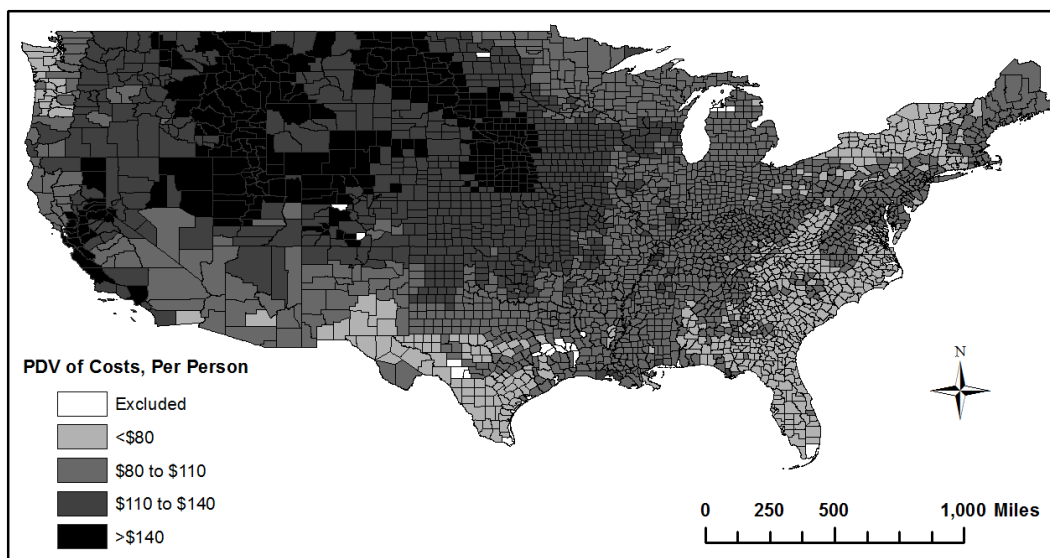


Figure 2.7: Present Discounted Social Cost of Climate-Related Crime, Per Person

Note: The map shows the per capita present discounted value of the social costs of the additional crimes estimated to be caused by climate change between 2010 and 2099. Costs are presented per person, for each county. The costs are based on climate predictions from the HadCM3 model, and are discounted using a discount rate of 6 percent.

tional effects, Figure 2.7 presents—for each U.S. county—the per capita present discounted value of the total social costs of future climate-related crime, by county. In other words, the figure shows the discounted value of the social cost of the number of additional crimes expect to occur in each county over the next 90 years, divided by each county’s current population. The table shows that the per capita cost of climate-related crime is highest in the West, where costs are between \$140 and \$180 per person, and lowest in the South and East, where costs are less than \$80 per person.

2.6 Discussion

The previous sections highlight two main results. First, weather has a strong causal effect on the incidence of criminal activity. For all offenses except manslaughter, higher temperatures lead to higher crime rates. The functional form of the relationship varies across offenses, with some categories, particularly property crimes, showing largest

marginal effects below 40 degrees F. This low-temperature dependency is in some ways surprising. Analyses of the impact of climate change on other economic outcomes, such as agriculture, have highlighted the role of extremely warm temperatures (Schlenker and Roberts, 2009). In contrast, my results suggest that the impact of climate change on property crime may operate largely through changes in the frequency of days with low to moderate temperatures.

Second, climate change will cause a substantial increase in crime in the United States. Relative to the total number of offenses that would occur between 2010 and 2099 in the absence of climate change, my calculations suggest that there will be an additional 35,000 murders, 216,000 cases of rape, 1.6 million aggravated assaults, 2.4 million simple assaults, 409,000 robberies, 3.1 million burglaries, 3.8 million cases of larceny, and 1.4 million cases of vehicle theft. The present discounted value of the social costs of these climate-related crimes is between 20 and 68 billion dollars.¹⁷

In interpreting these results, it is important to keep in mind that climate change will affect humans in a variety of ways (Tol, 2009; Deschenes and Greenstone, 2007, 2011; Hsiang, Meng, and Cane, 2011), and that a comprehensive cost-benefit analysis of climate change should consider all dimensions of costs and benefits. For example, given U.S. residents' high willingness to pay to live in areas with moderate climates (Cragg and Kahn, 1996), it is possible that the social costs of increased crime will be offset, at least in some regions, by the social benefits of more pleasant weather.

It is also worth emphasizing that the estimates presented here do not take into account longer-term adaptation mechanisms. If climate change does cause a permanent

¹⁷To put these dollar values in context, Deschenes and Greenstone (2011) estimate that climate-related changes in mortality and energy consumption will cause welfare losses of \$892 billion over the next century, based on a 3% discount rate and the HadCM3 model's predictions for the A1F1 scenario. Differencing out my estimate of the mortality-related costs of crime (murder and manslaughter together have a cost of approximately \$41 billion) implies that crime-related costs (\$68 billion) are likely to be about eight percent as large as the energy consumption and non-crime-related mortality costs of climate change in the United States (\$851). Of course, this comparison ignores any differences between the A1F1 and A1B emissions scenarios (the A1F1 scenario assumes higher emissions and more warming than the A1B scenario used in this paper).

increase in the frequency of crime, people in affected areas will have the opportunity to modify their behavior to avoid being victimized. Furthermore, it is likely that law enforcement agencies will respond with increased policing activity. The potential for such actions suggests that the estimates in this chapter should be viewed as an upper bound on the potential impacts of climate change on crime.

The estimates in this chapter also assume a static baseline of criminal activity, based on average crime rates between 2000 and 2009. Given the challenges of accurately predicting long-term trends in crime rates (Levitt, 2004), such an assumption is a reasonable analytical strategy. However, if for reasons unrelated to climate change, crime rates were to increase or decrease substantially over the coming decades, then the estimates from this chapter could significantly over- or underestimate climate's effects on future crime.

As a final caveat, I emphasize that this chapter's estimates of the social cost of climate-related crime should be considered to be highly uncertain. Although I monetize the social costs of additional crimes using point estimates drawn from the VSL and crime literatures (Viscusi and Aldy, 2003; McCollister, French, and Fang, 2010), I make no attempt to characterize the range of uncertainties associated with these valuations. Furthermore, consistent with previous literature on the role of discounting in economic analysis of climate change (Weitzman, 2007), I find that the present value of the social costs of additional crime depends heavily on the choice of a discount rate. Thus, the costs presented here are best interpreted as "back-of-the-envelope" estimates, rather than as precise statements of the exact cost of climate-related crime.

2.7 Conclusion

In this chapter, I document a robust statistical relationship between historical weather patterns and criminal activity, and use this relationship to predict how changes in U.S. climate will affect future patterns of criminal behavior. The results suggest that climate

change will have substantial effects on the prevalence of crime in the United States. Although previous assessments of the costs and benefits of climate change have primarily focused on other economic endpoints, the magnitude of the estimated impacts from this chapter suggests that changes in crime are an important component of the broader impacts of climate change.

2.8 Appendix

This appendix presents the results of a variety of sensitivity analyses of the relationship between weather and crime.

One potential concern about the analysis is that the relationship between weather and crime may have changed over time. To address this concern, Figure 2.8 plots the coefficients from separate regressions based on each of the five decades covered by the data: 1960-1969, 1970-1979, 1980-1989, 1990-1999, and 2000-2009. The data show more noise than the main regression results, but the overall pattern of crime increasing with temperature remains similar across decades for almost all crimes.

A second potential question is related to long-term adaptation. In particular, if residents of warmer climates are better adapted to warmer temperatures, then the relationship between weather and crime may vary across geographic regions. To assess whether this is the case, Figure 2.9 shows the results from separate regressions for counties in each of the four climate zones (based on long-term mean annual maximum daily temperature): <55 degrees F, 55 to 64 degrees F, 65 to 74 degrees F, and ≥ 75 degrees F. The figure shows that the effects of moderate and warm temperatures on crime is strikingly similar across climate zones. For very cold temperatures, the coefficients show somewhat more divergence, but this imprecision is primarily due to the fact that there are few days in the dataset in which the warmest climate zones are exposed to very low temperatures.

Another possibility related to adaptation is that people adjust to seasonal conditions, so that crime rates are driven by weather conditions relative to local expectations for that time of year. Under this hypothesis, a 60 degree F day could have very different effects depending on whether it occurred in April or July. As a test of this supposition, Figure 2.10 presents the results of a regression that includes interactions of the weather bin coefficients with three county-month temperature category variables. These categorical

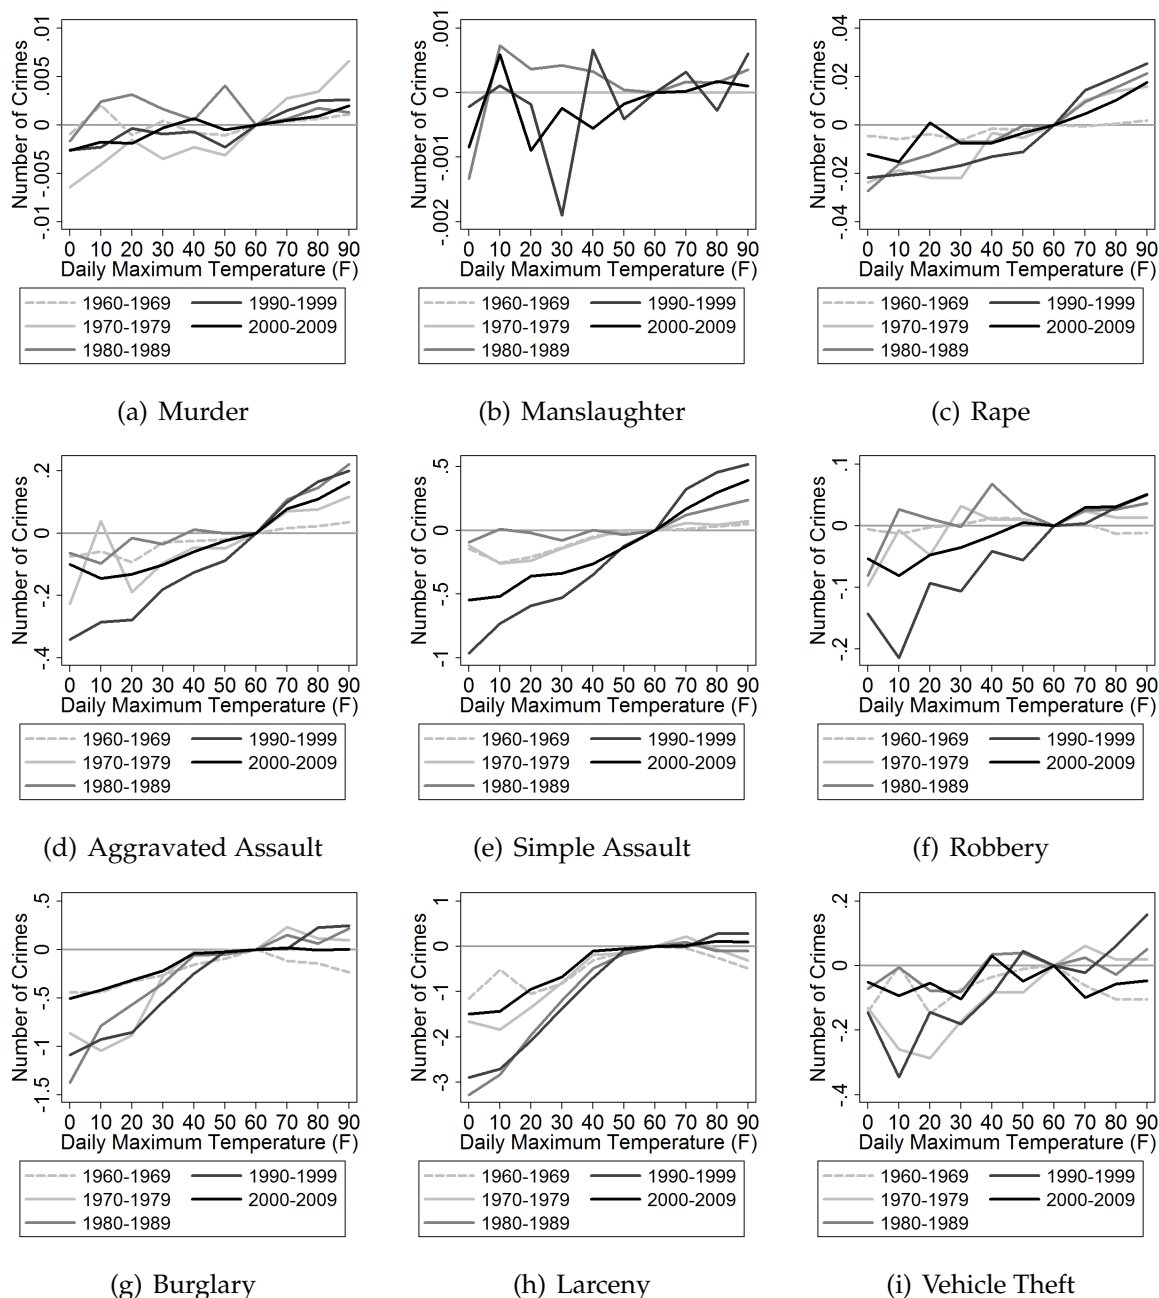


Figure 2.8: Monthly Crime and Daily Temperature, by Decade

Note: Each figure shows coefficients from regressions of the monthly crime rate per 100,000 persons on a semiparametric set of weather bin variables, for separate sets of observations from five different decades. These decades are: 1960-1969, 1970-1979, 1980-1989, 1990-1999, and 2000-2009. All coefficients are relative to one day in the 60 to 70 degrees F bin.

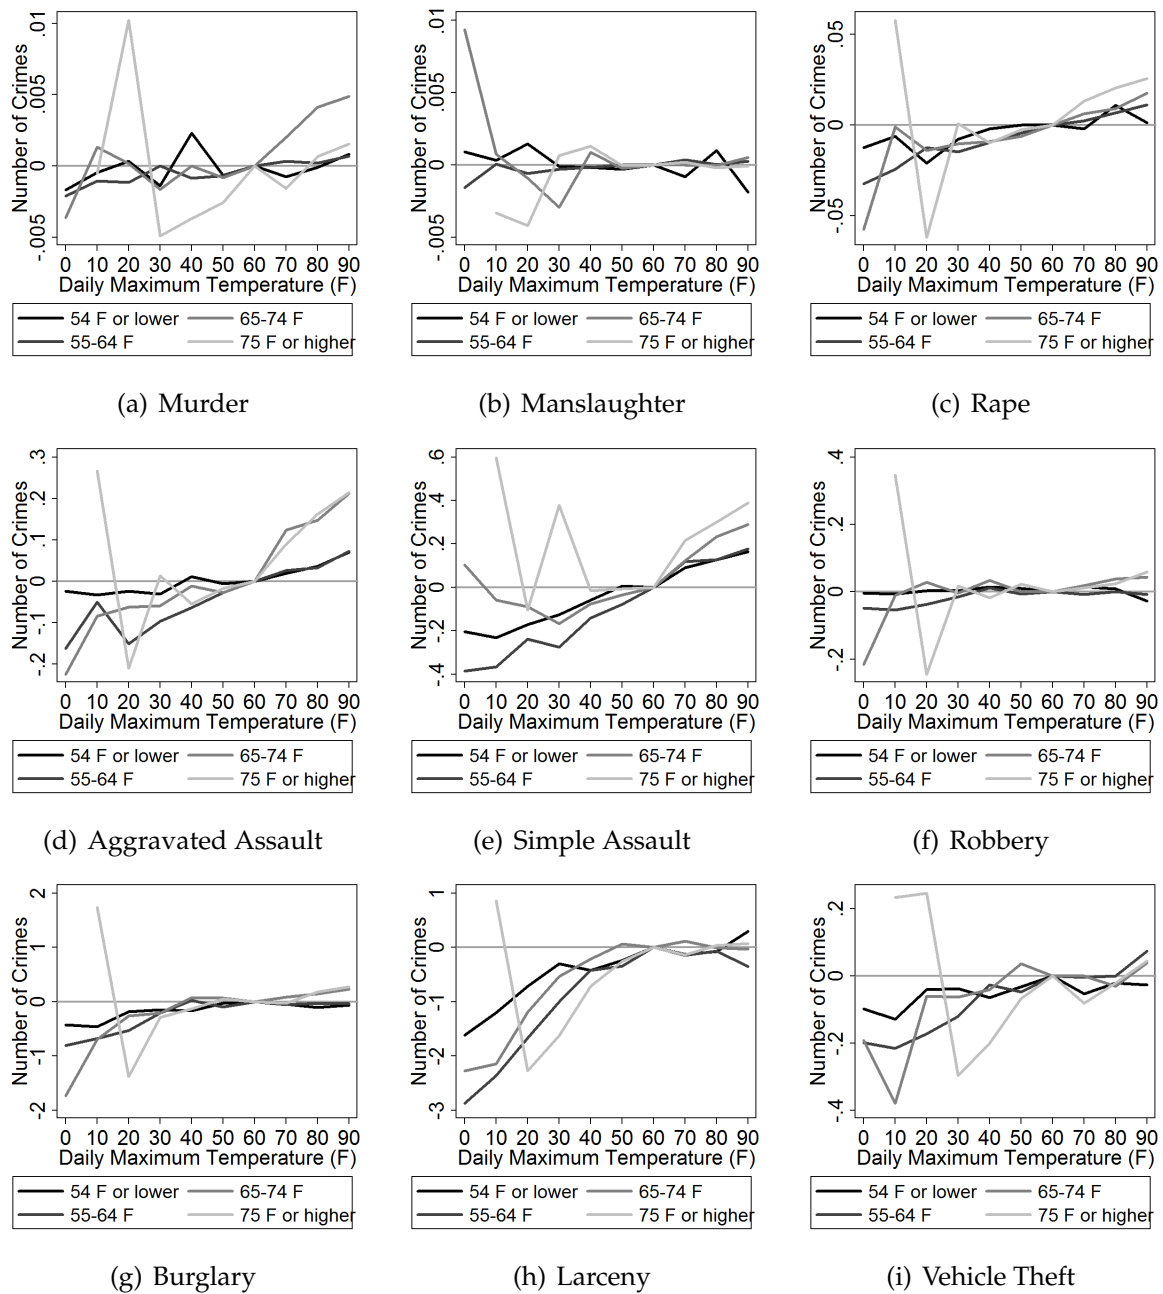


Figure 2.9: Monthly Crime and Daily Temperature, by Climate Zone

Note: Each figure shows coefficients from regressions of the monthly crime rate per 100,000 persons on a semiparametric set of weather bin variables, for counties in each of four climate zones. Each in-sample county is assigned to a climate zone based on whether its long-term mean annual maximum daily temperature falls into one of four ranges: <55 degrees F, 55 to 64 degrees F, 65 to 74 degrees F, and ≥ 75 degrees F. All coefficients are relative to one day in the 60 to 70 degrees F bin.

dummy variables indicate whether the average temperature in each month-by-county, over the period from 1960 to 2009, fell into one of three bins: <45 degrees F, 45 to 69 degrees F, or ≥ 70 F. Although the regressions show a fair amount of noise, particularly for temperatures that are not typical of normal monthly conditions, there are no obvious differences in the effects of temperature on crime that can be attributed to seasonal adaptation.

One additional concern about the analysis is related to heteroskedasticity in the crime rate variables. There is a large degree of variation in absolute crime levels between counties, and plots of time trends for individual counties show that the degree of seasonal variation is roughly proportional to the magnitude of the crime rate. Unfortunately, because the data contain a large number of months in which no crimes were committed (particularly for violent offenses such as murder and manslaughter), using a log transformation would be an inappropriate way to deal with this heteroskedasticity. Instead, as a sensitivity analysis, I estimate separate regressions for counties in each of four crime quartiles. To construct the quartiles, I calculate the mean crime rate for total crimes for each in-sample county, averaging across months and years. I then order these mean crime rates from highest to lowest. Quartile 1 represents counties below the 25th percentile; Quartile 2 represents counties between the 25 and 49th percentiles; Quartile 3 represents counties between the 50 and 74th percentiles; and Quartile 4 represents counties at or above the 75th percentile.

Figure 2.11 shows the results from this analysis. Generally speaking, the coefficients from Quartiles 1, 2, and 3 are of similar magnitude. As expected, the coefficients from regressions using data from Quartile 4 (counties with the highest average crime rate for all crimes) tend to be larger, although the exact degree of difference varies across types of crime.

A final question about the analytical approach used in this chapter is whether one month is a sufficiently long time period to account for any lagged impacts of weather

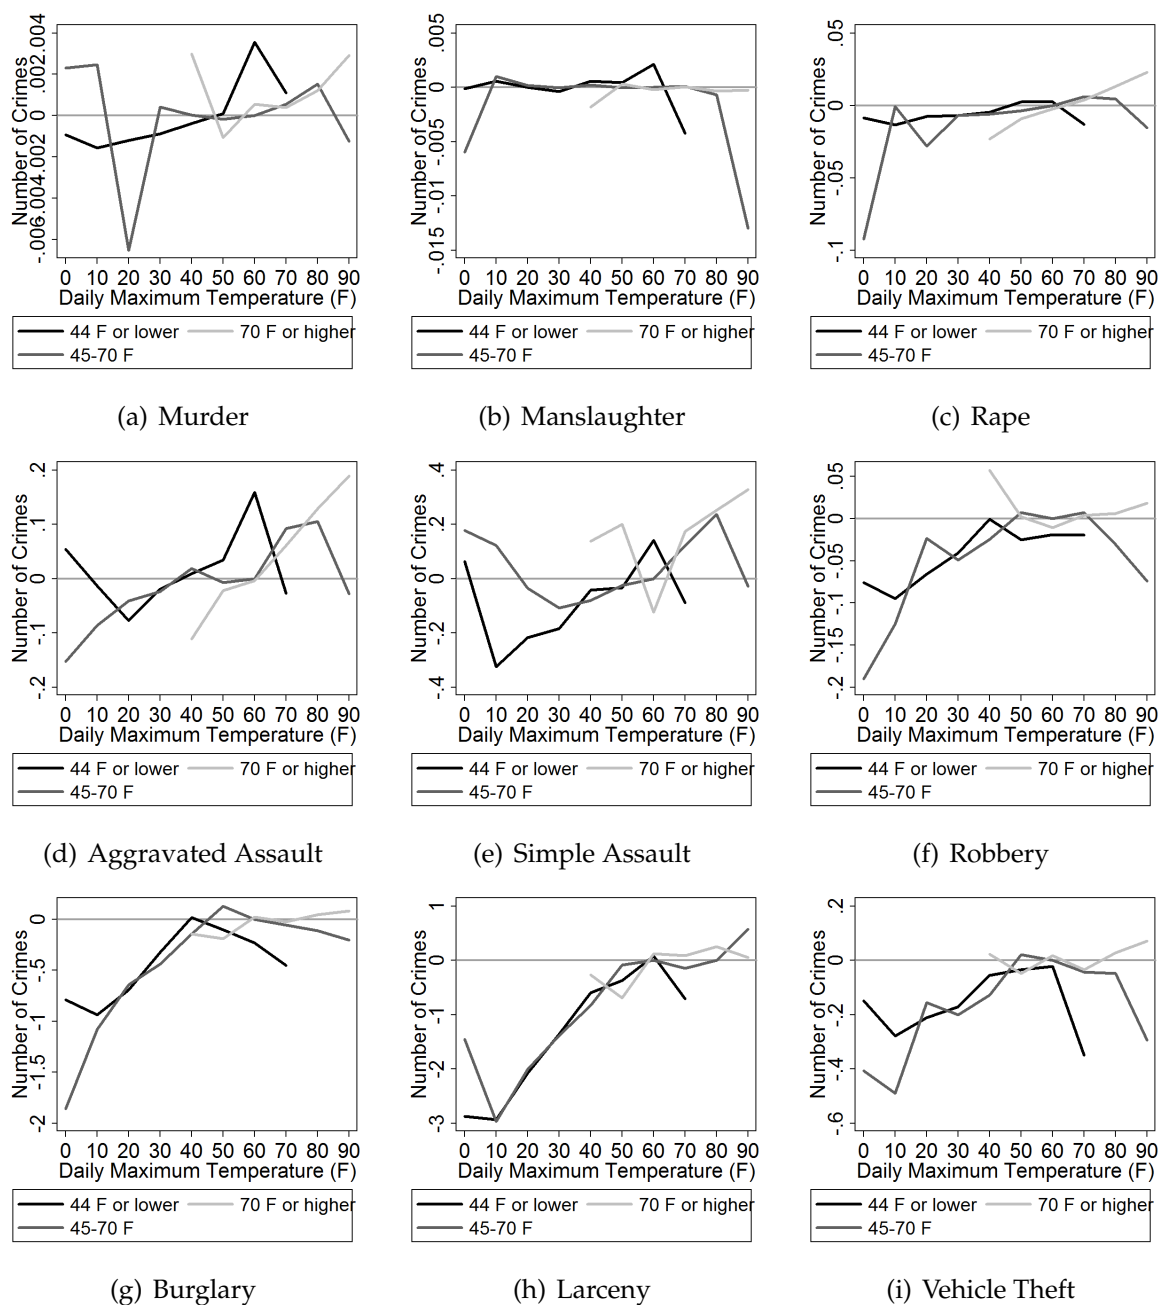


Figure 2.10: Monthly Crime and Daily Temperature, by Mean Monthly Temperature

Note: Each figure shows coefficients from a regression of the monthly crime rate per 100,000 persons on a semiparametric set of weather bin variables, interacted with three county-month temperature category variables. These categorical dummy variables indicate whether the average temperature in each month-by-county, over the period from 1960 to 2009, fell into one of three bins: <45 degrees F, 45 to 69 degrees F, or ≥ 70 F. All coefficients are relative to one day in the 60 to 70 degrees F bin.

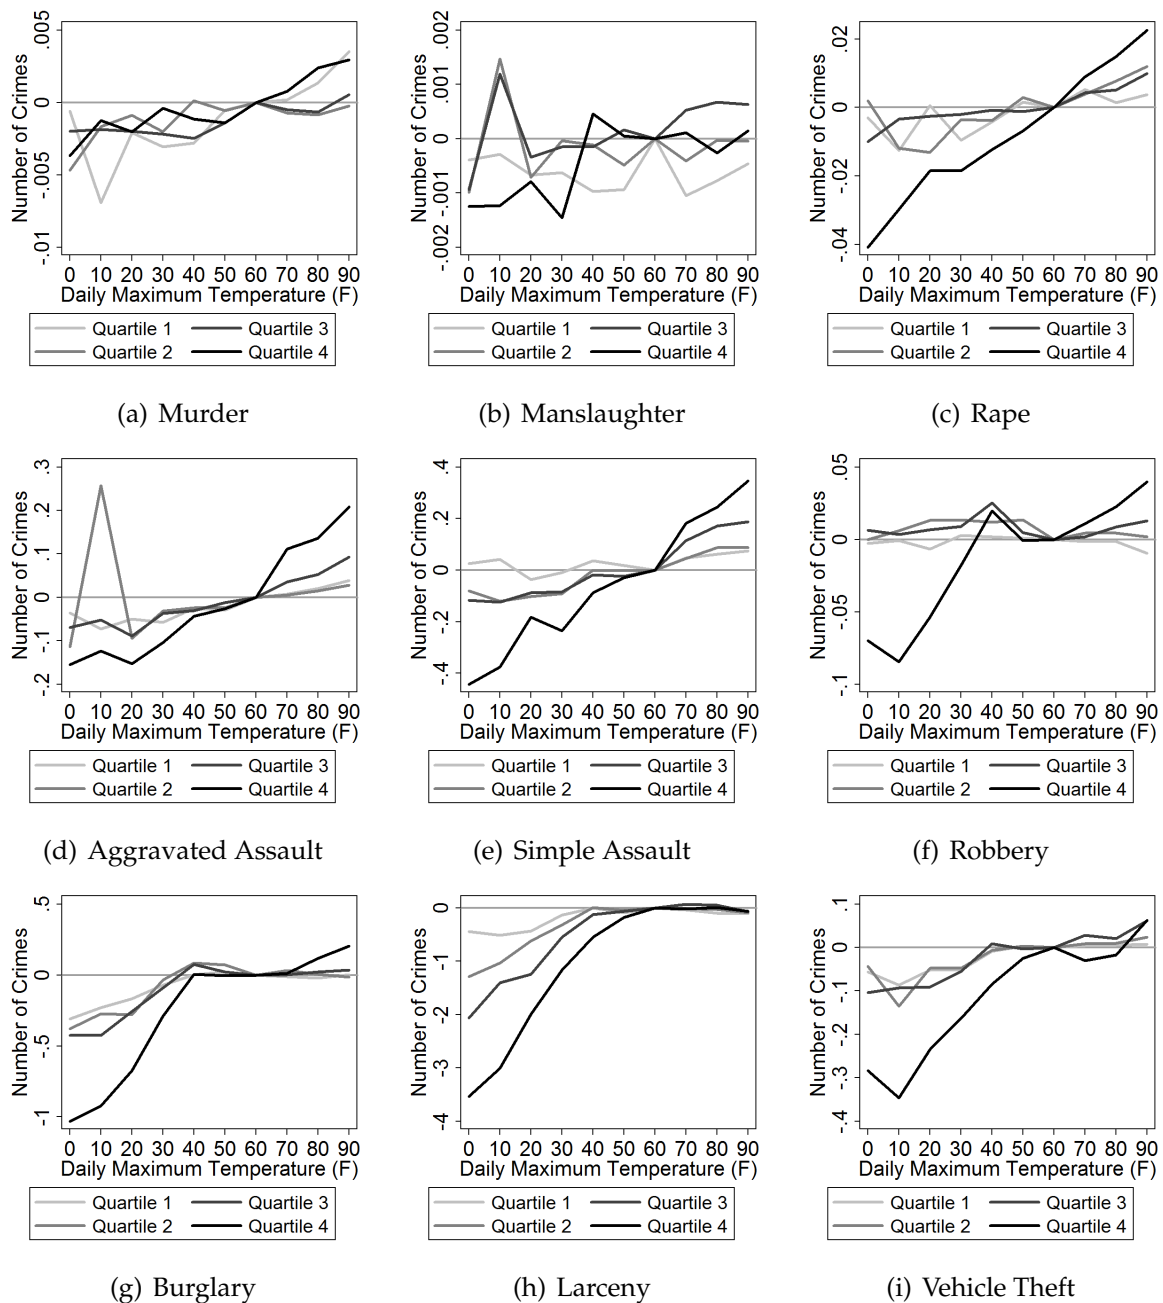


Figure 2.11: Monthly Crime and Daily Temperature, by Crime Rate Quartile

Note: Each figure shows coefficients from regressions of the monthly crime rate per 100,000 persons on a semiparametric set of weather bin variables, for counties in each of four crime quartiles. To construct the quartiles, I calculate the mean crime rate for all crimes for each in-sample county, averaging across months and years. I then order these mean crime rates from highest to lowest. Quartile 1 represents counties below the 25th percentile; Quartile 2 represents counties between the 25 and 49th percentiles; Quartile 3 represents counties between the 50 and 74th percentiles; and Quartile 4 represents counties at or above the 75th percentile. All coefficients are relative to one day in the 60 to 70 degrees F bin.

on crime. Although the insignificant coefficients on a one-month lag of weather suggest this is the case, I also conduct sensitivity analyses in which I run regressions using data that have been aggregated to quarterly and half-year time periods. Figure 2.12 shows the results of this analysis. Although regressions results based on more aggregate time periods are noisier than the results based on month-long time periods, the estimated coefficients from the three types of regressions are generally similar. The relationship between temperature and crime rates for aggravated and simple assault appears somewhat weaker based on the quarterly and half-year data. However, the effect of temperature on burglary and larceny is even stronger in the quarterly and half-year data. Overall, the figure suggests that a one-month aggregation period is sufficient to account for most “harvesting” that occurs as a result of negative serial correlation in crime rates.

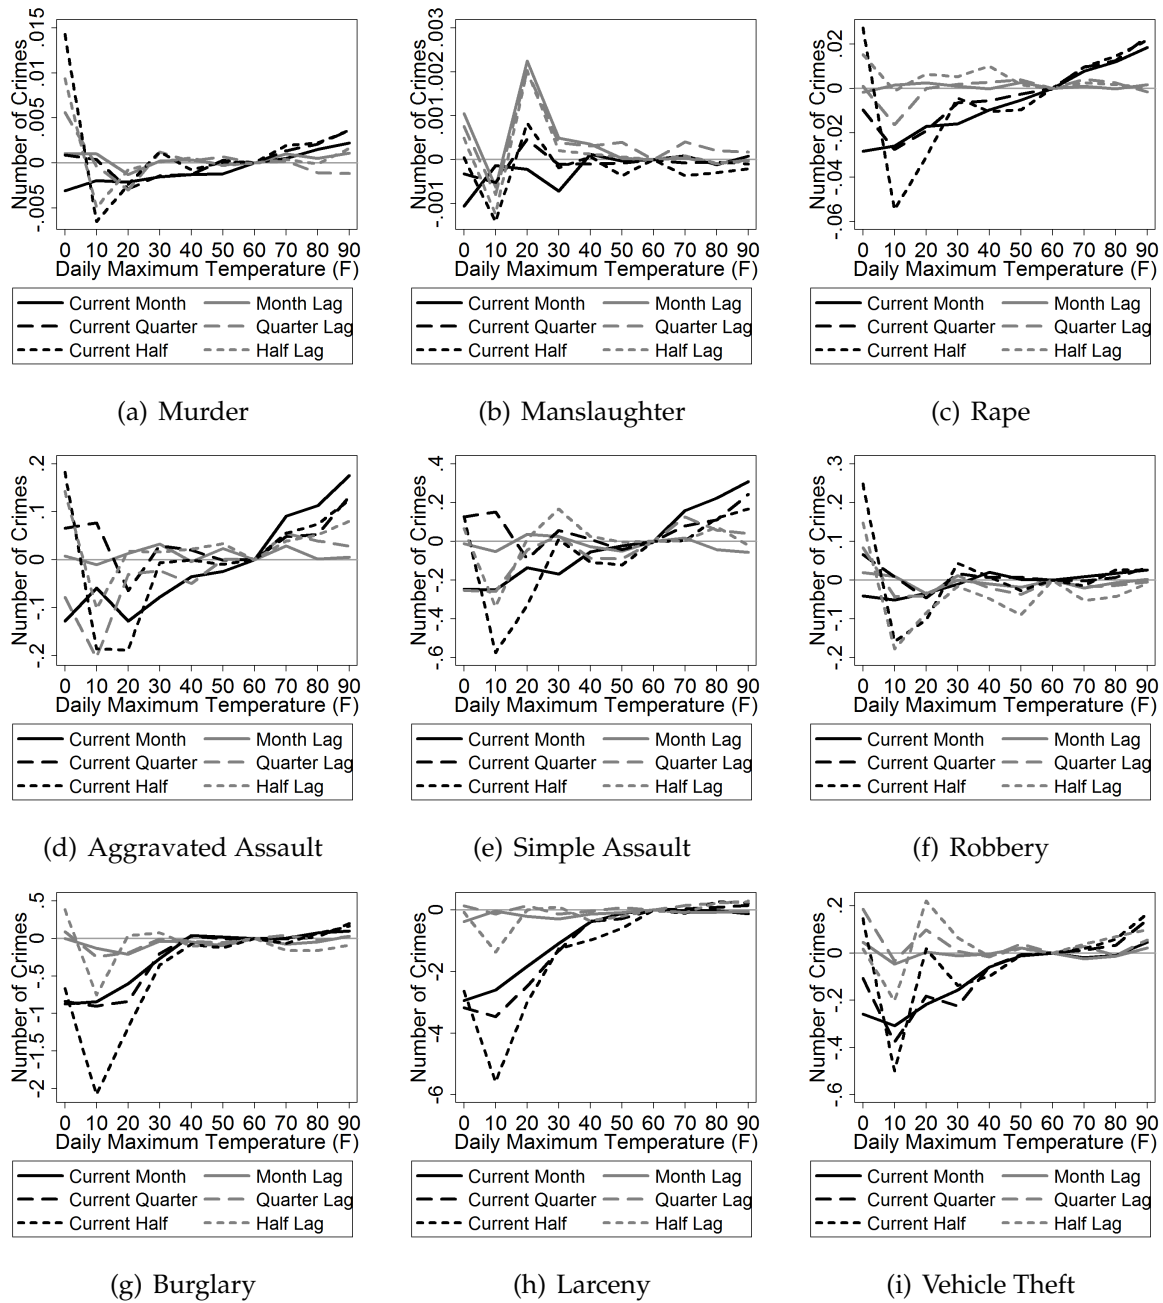


Figure 2.12: Crime and Daily Temperature, by Month, Quarter, and Half-year

Note: Each figure shows coefficients from regressions of the crime rate per 100,000 persons on a semi-parametric set of weather bin variables. Separate regression results are reported for crimes and weather aggregated by month, quarter, and half-year. All coefficients are relative to one day in the 60 to 70 degrees F bin.

Chapter 3:

Bayesian Updating with Potentially-Biased Evidence: A Model of Climate Skepticism

3.1 Introduction

The spectrum of public debate about climate change policy includes a substantial minority of individuals who challenge the scientific basis for global warming. These “climate skeptics” question whether global temperatures are rising and whether anthropogenic greenhouse gas emissions will cause significant changes in temperature.¹⁸ Although their claims are often dismissed by the mainstream scientific community, these skeptics exert considerable influence over public opinion about climate policy. In a 2011 Pew Research Center poll, 28 percent of a nationally representative sample of Americans answered “no” when asked, “From what you’ve read and heard, is there solid evidence that the average temperature on earth has been getting warmer over the past few decades, or not?” (Pew Research Center, 2011). Such beliefs about climate science are difficult to reconcile with the Intergovernmental Panel on Climate Change’s conclusion that “[w]arming of the climate system is unequivocal, as is now evident from observations of increases in global average air and ocean temperatures, widespread melting of snow and ice, and rising global average sea level” (IPCC, 2007).

Why do members of the public continue to hold beliefs that blatantly contradict the “unequivocal” conclusions of scientific experts? In this chapter, I argue that at least part of this discrepancy can be explained by a rational Bayesian learning process with three key features. First, individuals initially disagree about a scientific question that is central to a public policy debate: e.g., whether vaccines cause autism, or whether global temperatures have increased since the Industrial Revolution. Second, new scientific evidence about the question becomes available over time. Third, there is some possibility that the new evidence is systematically biased.

The central issue in such a situation arises from the fact that even though individuals would genuinely like to learn the true answer to the scientific question, their ability to learn is limited by the possibility of systematic bias. Although biased evidence

¹⁸For example, see www.climate-skeptic.com or www.skepticsglobalwarming.com.

does provide some information about the underlying scientific question, it is not quite enough to zero in on the exact answer. To illustrate why this is the case, in this chapter I develop a simple model of an individual's beliefs about whether current global temperatures have increased relative to temperatures in the period before the Industrial Revolution. I model the measured change in temperature as the sum of two random variables: the true change in temperature, and scientific measurement bias. The individual begins with prior beliefs about both the true change in temperature and scientific measurement bias. After receiving new evidence about measured temperature (as reported by scientists), the individual then updates his beliefs.

In contrast to the usual Bayesian result that beliefs converge over time, I find that in this simple model, the posterior distributions for true change in temperature and measurement bias depend on the choice of priors, even when the number of observations of measured change in temperature goes to infinity. From a purely statistical standpoint, this failure of convergence (due to a flat likelihood function over some combinations of beliefs) is well-understood.¹⁹ In the context of my model, however, lack of identifiability generates two predictions that explain several important features of the current debate about climate change.

First, when there is any possibility of systematic measurement bias, no amount of new evidence can reconcile initial differences of opinion about whether global temperatures have increased. Even if two people who initially disagree about temperature share the same initial belief that mean measurement error is zero, as long as they admit some possibility of systematic measurement bias, then they will always interpret new evidence as supporting different conclusions about the true change in temperature. This result holds even in the limiting case with an infinite number of new temperature measurements.

Second, because people who begin with more extreme beliefs about temperature end up with more extreme beliefs about measurement bias, there is a tendency for the

¹⁹See, e.g., Pratt, Raiffa, and Schlaifer (1995), or Feldman (1991).

model to create “climate skeptics”. Thus, the model suggests that skepticism about science is a completely rational response to observing a discrepancy between one’s own prior beliefs and new evidence reported by scientists. Although behavioral biases may certainly play a role in the phenomenon of climate skepticism, this result contrasts sharply with the viewpoint that skeptics are being disingenuous when they argue that the evidence for climate change is overstated.

A natural response to this model is to argue that the main results no longer hold if individuals can also observe separate evidence about the magnitude of the scientific bias. In such a situation, rational Bayesian updating will lead individuals to learn the exact true answer to the scientific question. This is, in some ways, the key policy lesson to be drawn from this chapter: not only is it important for scientists to communicate the conclusions of their research, but it is also important for them to describe how they reached those conclusions. In order to address concerns about potential biases, the scientific community would be well-served by focusing on improving both the quantity *and* transparency of climate research.

This chapter builds on several strands of literature. It is most directly related to a set of climatology papers that develop Bayesian statistical methods for combining results from multiple climate proxy records and general circulation models (Tebaldi et al, 2005; Haslett et al, 2006; Li, Nychka, and Ammann, 2010; Smith et al, 2009). A few of these papers explicitly model measurement bias and the challenge it creates for learning about future climate (Jun, Knutti, and Nychka, 2008; Buser, Kunsch, and Schar, 2010; Buser, 2009). However, unlike my approach, which takes identifiability as a fundamental and unresolvable dilemma, they treat identifiability as a statistical issue to be addressed through ad hoc assumptions such as imposition of zero-sum bias conditions or informative priors (Buser et al, 2009). Furthermore, these papers can best be characterized as “models of climate”, not as “models of people’s beliefs about climate”.

A second set of papers explores optimal policy and the dynamics of learning about

climate change. For example, Weitzman (2009) shows how Bayesian learning in the presence of structural uncertainty about climate sensitivity leads to fat-tailed posterior distributions for climate damages (see also Pindyck, 2009; Newbold and Daigneault, 2009). Other authors focus on Bayesian learning about climate as part of a sequential decision making strategy (Hammitt et al, 1992) or optimal control problem (Kelly and Kolstad, 1999; Nordhaus and Boyer, 2000; Leach, 2007; Karp and Zhang, 2006).

A final set of related papers study whether individuals use a rational Bayesian learning process (Cyert and DeGroot, 1974; Viscusi and O'Conner, 1984; Viscusi and Zeckhauser, 2006; Kelly et al, 2009; Cameron, 2005). These papers identify a variety of behavioral and psychological inconsistencies (Zimper and Ludwig, 2009). For example, people's beliefs about climate change are affected by political beliefs (Borick and Rabe, 2010; Hamilton, 2011), ambiguity aversion (Millner, Dietz, and Heal, 2010), cognitive dissonance (Wagner and Zeckhauser, 2011), and recent local weather fluctuations (Deryugina, 2011; Egan and Mullins, 2010).

Although these behavioral factors undoubtedly influence how individuals learn about climate science, the model in this chapter is based on a completely rational Bayesian framework. The purpose of adopting this rationalist approach is twofold. First, by setting aside behavioral explanations, I am able to focus on understanding the influence of systematic bias on learning about climate change. Although in practice, it is likely that the results presented in this paper explain only part of the empirical discrepancy between public beliefs and scientific evidence about climate change, my model nonetheless emphasizes the potentially important role of this simple but previously neglected rational influence on learning. Second, in order to understand competing hypotheses based on psychological and behavioral explanations, it is important to have a statement of the "correct" baseline rational model. Although my model makes assumptions that may not perfectly describe the actual human brain's learning process, this paper still describes what I believe to be a key qualitative feature of how people learn when evidence is potentially biased.

The remainder of this chapter is organized as follows. Section 3.2 describes the general Bayesian logic underlying my argument. Section 3.3 applies this logic to a simple model of skepticism about climate change. Section 3.4 discusses the findings and concludes.

3.2 General Model

I begin by describing the basic Bayesian logic that leads to the results discussed in the introduction. Consider a simple abstract model, in which an individual would like to learn the distribution of some policy-relevant random variable T , e.g., the increase in global temperature since the Industrial Revolution, or the treatment effect of vaccination on the prevalence of autism. Unfortunately, T is not directly observable. Instead, the individual observes a series of realizations of a signal $Z \equiv Z(T, E)$ that is influenced by both the policy-relevant variable T and by some nuisance variable E that represents systematic measurement error. Like T , the variable E is not directly observable.

Let the joint distribution of T and E be given by $f(T, E; \theta)$, where θ is a vector of parameters with rank two or greater. Furthermore, suppose that at the beginning of the model, the individual's prior beliefs about θ are represented by θ_0 . The individual then observes a series of realizations z_1, z_2, \dots, z_n of the signal Z . Her problem is to use this new information to update her beliefs about the joint distribution of T and E .

Now, after observing n realizations of Z , the individual's posterior beliefs about T and E are given by $\theta_n = \theta | \theta_0, z_1, z_2, \dots, z_n$. It is straightforward to construct counterexamples that show that it is not necessarily true that $\lim_{n \rightarrow \infty} \theta_n = \theta$. Why is this the case? The basic argument is that although $Z \equiv Z(T, E)$ defines a unique mapping from (T, E) to Z , the reverse is not necessarily true. In fact, in most situations, there will be an infinite number of possible mappings from Z to (T, E) . Thus, even though the observations z_1, z_2, \dots, z_n can be used to estimate the precise distribution of the the

random variable Z , these data are insufficient to recover the parameter θ that describes the joint distribution of T and E , or to recover even the marginal distributions of T and E .

In summary, this model implies that that even after observing an infinite number of realizations of the signal, two individuals who start with different prior beliefs about θ will still not agree about the posterior distribution of T . Although the above formulation may seem arbitrary, it captures a central feature of the way that people learn about scientific debates: individuals may be aware of the general results from recent scientific studies, but possess little information with which to evaluate their credibility. This lack of credible information about scientific bias means that people rely heavily on their own prior judgments when updating their beliefs based on new research.

3.3 Application to Climate Skepticism

3.3.1 Model

The previous section presents a general but informal argument that explains why Bayesian updating based on potentially-biased evidence cannot completely resolve prior disagreements. In this section I flesh out the intuition for this result by presenting a more detailed model of how an individual learns about climate change science. I abstract away from the complexities of the climate change debate to a single question: are current global temperatures higher than temperatures during the pre-Industrial Revolution period? The challenge that the individual faces is that it is not possible to observe the true change in global temperature. Instead, the individual observes a set of *measurements* of the change in temperature that are provided by the scientific community. These measurements include two sources of inaccuracy: natural variation in temperature readings, and scientific measurement error.

The individual's task is to use this set of observations about measured change in tem-

perature to update his beliefs about the true change in global temperature. As in the previous section, I assume a rational Bayesian updating process, in which the individual begins with prior distributions for the true change in global temperature and scientific measurement error, and then calculates posterior distributions using a combination of these prior beliefs and the new evidence, following Bayes' Rule.

Unlike other papers on expert opinion (Crawford and Sobel, 1982; Battaglini, 2004; Krishna and Morgan, 2001), I do not explicitly model the incentives of the scientific community. Instead, I assume—perhaps unrealistically—that scientists do their best to provide unbiased estimates of the change in global temperatures. However, by including scientific measurement error in the model, I still allow for the possibility that scientists unintentionally reported biased estimates. Such bias could have a variety of causes, e.g., flawed methodologies, publication bias, or reliance by different researchers on the same underlying sets of data (Jun, Knutti, and Nychka, 2008).

More formally, let Z be a random variable that represents the measured change in global temperature. This variable is defined as the sum of two other random variables: the true change in global temperature T , and scientific measurement error E .

$$Z \equiv T + E \tag{3.1}$$

I model the true change in temperature T as a random variable in order to include uncertainty caused by natural stochasticity in climate. For simplicity, I assume T to be normally distributed:

$$T \sim N(\tau, \phi^2) \tag{3.2}$$

where the mean change in temperature τ is the parameter that scientists and the general public would like to learn, and the variance ϕ^2 can be interpreted as natural sampling variation (e.g., zero-mean differences caused by measuring temperatures in a range of geographic locations).

Similarly, I assume that the systematic measurement error variable E is also normally

distributed:

$$E \sim N(\epsilon, \nu^2) \quad (3.3)$$

where ϵ is the mean measurement bias, and the variance ν^2 captures additional zero-mean noise attributable to the scientific measurement process. For example, a positive value of ϵ might mean that all temperature records based on tree ring cores systematically overestimate the actual change in temperature that has occurred over the last several centuries, and a positive value of ν^2 might indicate that variation in growth patterns across trees introduces additional (zero-mean) measurement error.

Let $f(z|\tau, \epsilon)$ denote the probability density function for Z , where the variance parameters ϕ and ν are suppressed for notational convenience. Using the standard result that the sum of two normal random variables is itself normally distributed, equation (3.1) implies that the distribution for Z is:

$$Z \sim N(\tau + \epsilon, \phi^2 + \nu^2) \quad (3.4)$$

Note that this formulation allows for real bias, in the sense that when ϵ is non-zero, observations of Z will be consistently higher or lower than T .

The timing of the model is as follows. Each period s , scientists draw a single realization z_s from Z , the random variable representing measured change in temperature. As indicated by Equation 3.1, this realization is itself based on the sum of (unobserved) realizations t_s and e_s of the random variables T and E . Scientists then report this observation z_s to the general public. Each member of the public then uses this new information to update his or her beliefs about the true change in temperature and systematic measurement bias.

As is typical in models of Bayesian updating, I assume that individuals do not know the parameters τ and ϵ that describe mean true change in temperature and mean systematic measurement bias (although I do assume that the variance parameters ϕ^2 and ν^2 are known).²⁰ Instead, in period 0, before observing any data, each individual uses

²⁰Because ϕ^2 and ν^2 both represent zero-mean noise, these variance parameters have no effect on the

her prior knowledge (based, for example, on theoretical calculations, or anything else) to generate probability density functions that represents her beliefs about the possible range of values that these two parameters might take:

$$g_0(\tau) \sim N(\tau_0, \sigma_0^2) \quad (3.5)$$

$$h_0(\epsilon) \sim N(\epsilon_0, \omega_0^2) \quad (3.6)$$

For example, $g_s(\tau)$ describes an individual's prior beliefs about the parameter τ , and $h_s(\epsilon)$ describes an individual's prior beliefs about the parameter ϵ , based on all information that has been accumulated as of time s . After observing the new temperature measurement z_{s+1} , the individual uses a Bayesian updating process to generate new posterior parameter values τ_{s+1} , σ_{s+1}^2 , ϵ_{s+1} , and ω_{s+1}^2 that better represent her beliefs about τ and ϵ , based on the new information contained in z_{s+1} .

3.3.2 Results

In this section I present the main analytical results. I begin by deriving the posterior distributions for the true change in temperature and for measurement error, after the individual observes n observations $\{z_1, z_2, \dots, z_n\}$ of the measured change in temperature. These observations can be envisioned as distinct pieces of evidence provided by scientists.

Theorem 3. *After observing n observations $\{z_1, z_2, \dots, z_n\}$, the posterior distribution for τ is:*

$$\tau|z_1, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0 \sim N(\tau_n, \sigma_n^2) \quad (3.7)$$

where:

$$\tau_n \equiv \frac{(n\omega_0^2 + \phi^2 + \nu^2)\tau_0 + \sigma_0^2 \sum_{i=1}^n (z_i - \epsilon_0)}{n\omega_0^2 + \phi^2 + \nu^2 + n\sigma_0^2} \quad (3.8)$$

and:

$$\frac{1}{\sigma_n^2} \equiv \frac{n}{n\omega_0^2 + \phi^2 + \nu^2} + \frac{1}{\sigma_0^2} \quad (3.9)$$

overall results as the number of temperature observations becomes large.

Furthermore, the posterior distribution for ϵ is:

$$\epsilon|z_1, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0 \sim N(\epsilon_n, \omega_n^2) \quad (3.10)$$

where:

$$\epsilon_n \equiv \frac{(\phi^2 + \nu^2 + n\sigma_0^2)\epsilon_0 + \omega_0^2 \sum_{i=1}^n (z_i - \tau_0)}{n\omega_0^2 + \phi^2 + \nu^2 + n\sigma_0^2} \quad (3.11)$$

and:

$$\frac{1}{\omega_n^2} \equiv \frac{1}{\omega_0^2} + \frac{n}{\phi^2 + \nu^2 + n\sigma_0^2} \quad (3.12)$$

Proof. See Appendix I. □

Now suppose that $n \rightarrow \infty$, so that the individual observes a (nearly) unlimited number of measurements of the change in temperature. The limiting posterior distributions for τ and ϵ are given by Theorem 4.

Theorem 4. *Let τ^* and ϵ^* denote the true population means for τ and ϵ . As the number of observations n approaches infinity, the limiting posterior distribution for τ is:*

$$\lim_{n \rightarrow \infty} \tau|z_1, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0 \sim N(\tilde{\tau}, \tilde{\sigma}^2) \quad (3.13)$$

where:

$$\tilde{\tau} \equiv \frac{\omega_0^2 \tau_0 + \sigma_0^2 (\tau^* + \epsilon^* - \epsilon_0)}{\omega_0^2 + \sigma_0^2} \quad (3.14)$$

and:

$$\frac{1}{\tilde{\sigma}^2} \equiv \frac{1}{\omega_0^2} + \frac{1}{\sigma_0^2} \quad (3.15)$$

Furthermore, the limiting posterior distribution for ϵ is:

$$\lim_{n \rightarrow \infty} \epsilon|z_1, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0 \sim N(\tilde{\epsilon}, \tilde{\omega}^2) \quad (3.16)$$

where:

$$\tilde{\epsilon} \equiv \frac{\sigma_0^2 \epsilon_0 + \omega_0^2 (\tau^* + \epsilon^* - \tau_0)}{\omega_0^2 + \sigma_0^2} \quad (3.17)$$

and:

$$\frac{1}{\tilde{\omega}^2} \equiv \frac{1}{\omega_0^2} + \frac{1}{\sigma_0^2} \quad (3.18)$$

Proof. Equation (3.14) follows from equation (3.8) by applying l'Hopital's rule to the linear terms and applying the law of large numbers and the relationship from equation (3.4) to the sum $\sum_{i=1}^n z_i$. Similarly, equation (3.15) follows from applying l'Hopital's rule to equation (3.9). Finally, convergence in equations (3.14) and (3.15) implies that $\lim_{n \rightarrow \infty} \tau |z_1, \dots, z_n, \tau_0, \epsilon_0$ converges to $N(\tilde{\tau}, \tilde{\sigma}^2)$, as in equation (3.13). The proof for equations (3.16), (3.17), and (3.18) follows by symmetry. \square

As a first step towards understanding these formulas, it is worth noting that neither ϕ^2 nor ν^2 appear in any of the equations. Intuitively, this is because both of these variance parameters represent zero-mean Gaussian noise that does not matter when the sample size becomes large enough (due to the Law of Large Numbers).

It is now informative to conduct comparative statics. Consider first the reference case where $\epsilon_0 = 0$ and $\omega_0^2 = 0$. This is the standard Bayesian model with no measurement error. In this case, $\tilde{\tau} = \tau^*$ and $\tilde{\sigma}^2 = 0$, and the individual's posterior distribution for the true change in temperature approaches the actual population distribution.

Now suppose instead that ϵ_0 takes some non-zero value but $\omega_0^2 \rightarrow 0$. Once again, $\tilde{\tau} \rightarrow \tau^*$ and $\tilde{\sigma}^2 \rightarrow 0$. The intuition for this result is that when the bias is perfectly known, it can simply be subtracted and then ignored.

Next, consider the case where ϵ_0 and ω_0^2 are both non-zero, but $\sigma_0^2 \rightarrow 0$. The result is that $\tilde{\tau} \rightarrow \tau_0$ and $\tilde{\sigma}^2 \rightarrow 0$. Since the individual believes that his prior estimate τ_0 is exactly correct, there is no need to update based on new evidence.

If all of the parameters take non-zero values, then $\tilde{\tau}$ takes a value that is a weighted average of the prior mean τ_0 and the sample mean net of the prior measurement error, where the weights are determined by the precisions of the prior distributions. Because no information is available that can completely resolve the uncertainty about the amount of measurement bias, the variance $\tilde{\sigma}^2$ does not go to zero.

The most interesting case, however, is when ϵ_0 is zero but ω_0^2 takes a positive value.

Surprisingly, even when $\epsilon^* = 0$, the result is that $\tilde{\tau} \rightarrow \frac{\omega_0^2 \tau_0 + \sigma_0^2 \tau^*}{\omega_0^2 + \sigma_0^2}$ and $\tilde{\epsilon} \rightarrow \frac{\omega_0^2 (\tau^* - \tau_0)}{\omega_0^2 + \sigma_0^2}$. This is the key result of this chapter. The intuition is that as long as ω_0^2 is non-zero, there is some possibility of systematic measurement error. Thus, even when $\epsilon_0 = 0$ and $\epsilon^* = 0$, a rational Bayesian updater who observes a discrepancy between his prior belief about the true change in temperature (τ_0) and his asymptotically correct estimate of the mean of the measured change in temperature ($z_n \rightarrow \tau^*$) will attribute some of that difference to measurement error. Thus, an individual whose initial beliefs satisfy $\tau_0 < \tau^*$ and $\epsilon_0 = \epsilon^* = 0$ will always arrive at posterior beliefs that underestimate the true change in temperature and overestimate the amount of systematic measurement error.

3.4 Discussion

The value of the simple model presented in this chapter lies in the fact that despite a parsimonious specification, the model generates two somewhat surprising insights into climate change skepticism (and more broadly, into skepticism about other scientific debates). Both of these results arise from the fact that when there is any possibility of systematic scientific measurement bias, the posterior distributions for the true change in temperature and measurement bias remain dependent on the choice of priors even as the number of measured temperature observations goes to infinity.

First, when there is any possibility of systematic measurement bias, no amount of new evidence can reconcile initial differences of opinion about whether global temperatures have increased. As long as ω_0^2 and σ_0^2 are non-zero, two individuals with different prior beliefs about temperature will arrive at different posterior beliefs about temperature. This is true even if both individuals start with the same prior beliefs about measurement bias—and furthermore, even if both individuals start with the same prior belief that mean measurement bias is zero.

Second, there is a natural tendency for people who begin with more extreme beliefs about temperature to end up with more extreme beliefs about measurement bias.

Again, as long as there is some possibility of systematic measurement bias (in the sense that ω_0^2 is non-zero), this result holds even if those individuals share the same prior belief that the mean of measurement bias is zero. The intuition for this result is that an individual whose prior beliefs about temperature are further from the true population mean always attributes a greater quantity of the discrepancy between his initial beliefs and the temperature measurements to measurement error. This contrasts sharply with the viewpoint that climate skeptics are being disingenuous when they argue that the evidence for climate change is overstated. Instead, it suggests that skepticism about science is a completely rational response to observing a discrepancy between one's own prior beliefs and new evidence reported by scientists.

Although the results presented in this chapter are based on a specific functional form in which systematic bias enters as an additive component of the observed temperature signal, the results can be generalized to other types of models. For example, a “mixture distribution” version of the model, in which individuals update their beliefs based on the likelihood of observing different signals, produces similar qualitative results.²¹ In this alternative model, scientists again make a series of measurements of the change in global temperatures since the beginning of the industrial revolution. However, there is some probability that their scientific methodology is simply wrong, e.g., because some factor such as widespread plant disease causes the historical relationship between temperatures and tree ring growth to be very different from what it is today. I model this uncertainty by assuming that with probability P scientists observe measurements from the true temperature record, and with probability $1 - P$ they observe meaningless noise. These probabilities represent a one-time draw, so that if scientists start out observing the true temperature record, then they will always observe the true temperature record, and vice versa.

In such a model, individuals have prior beliefs about the distribution of signals (i.e.,

²¹Complete results are available from the author upon request. I plan to explore these models in more detail in future research.

measurements of historical temperatures) that would be observed from each of these two types of data generating processes. Their problem is to use the signals to update both their beliefs about which signal they are observing (the true record or meaningless noise) as well as their beliefs about what the actual change in historical temperatures has been. Although I do not present the mathematical details here, the basic qualitative result is similar: an individual's prior beliefs always influence her posterior beliefs, even as the number of temperature measurements goes to infinity.

3.5 Conclusion

What are the policy implications of the somewhat gloomy results presented in this chapter? One interpretation is that policy makers should not expect future scientific research to build a public consensus for action on climate change.²² Although the model does suggest that additional research will bring beliefs closer together, it cannot completely bridge the divide between people who begin with very different initial beliefs about rising temperatures. Thus, policy makers will have to resort to more traditional political tools—such as strong leadership and emotional appeals—to generate public support for appropriate climate actions.

A second, more positive, interpretation is that science has great potential for helping to inform the public's opinions about climate change. However, the model strongly suggests that only reporting the results from additional scientific studies may not be very useful. Instead, reporting these results is most valuable when scientists also make an effort to communicate information about the quality of their research, so that individuals can make informed decisions about scientific bias. The central obstacle to achieving consensus is that people are unsure whether they can trust climate science. If this obstacle were removed—by making data, methodologies, and results more available to the public in a non-technical form that encourages critical thinking and discussion—it

²²Survey evidence suggests that people who distrust scientists are less likely to believe that climate change is real (Kellstedt, Zahran, and Vedlitz, 2008).

seems likely that at least some doubts about scientific bias could be eliminated and that public opinions about climate change might converge.

As discussed in the introduction, a great deal of evidence suggests that people are at best imperfect Bayesians. Thus, a reasonable objection to this chapter is to say that the Bayesian model of belief formation is not descriptively accurate. I acknowledge that this criticism has empirical strength: beliefs about climate change are influenced by political affiliations, cognitive dissonance, recent weather conditions, and other “non-Bayesian” factors (Borick and Rabe, 2010; Wagner and Zeckhauser, 2011; Deryugina, 2011). Yet, unlike more complicated models, I argue that the basic application of Bayesian decision theory presented here produces a simple and practically useful prescription: make science more transparent.

3.6 Appendix

Proof of Theorem 1:

Proof. Consider first the posterior distribution for the temperature parameter τ . After observing the n observations z_1, z_2, \dots, z_n , the new posterior distribution for τ is:

$$g(\tau|z_1, z_2, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0) = \frac{\left(\int_{\epsilon=-\infty}^{\infty} (\prod_{i=1}^n f(z_i|\tau, \epsilon)) h_0(\epsilon) d\epsilon \right) g_0(\tau)}{\int_{\tau=-\infty}^{\infty} \left(\int_{\epsilon=-\infty}^{\infty} (\prod_{i=1}^n f(z_i|\tau, \epsilon)) h_0(\epsilon) d\epsilon \right) g_0(\tau) d\tau} \quad (3.19)$$

This can be rewritten in Bayesian notation by substituting the formula for the normal distribution and then dropping normalizing constants:

$$g(\tau|z_1, z_2, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0) \propto \int_{\epsilon=-\infty}^{\infty} \left(\prod_{i=1}^n \exp \left(\frac{-(z_i - (\tau + \epsilon))^2}{2(\phi^2 + \nu^2)} \right) \right) \exp \left(\frac{-(\epsilon - \epsilon_0)^2}{2\omega_0^2} \right) d\epsilon \cdot \exp \left(\frac{-(\tau - \tau_0)^2}{2\sigma_0^2} \right) \quad (3.20)$$

Applying some algebra shows that:

$$g(\tau|z_1, z_2, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0) \propto \int_{\epsilon=-\infty}^{\infty} \exp \left(\frac{n\omega_0^2\epsilon^2 + (\phi^2 + \nu^2)\epsilon^2 - 2\omega_0^2 \left(\sum_{i=1}^n (z_i - \tau) \right) \epsilon - 2(\phi^2 + \nu^2)\epsilon_0\epsilon + \omega_0^2 \sum_{i=1}^n (z_i - \tau)^2 + (\phi^2 + \nu^2)\epsilon_0^2}{-2(\phi^2 + \nu^2)\omega_0^2} \right) d\epsilon \cdot \exp \left(\frac{-(\tau - \tau_0)^2}{2\sigma_0^2} \right) \quad (3.21)$$

Now, by completing the square in ϵ , moving terms that contain τ outside of the integral, and dropping normalizing constants, the result is:

$$g(\tau|z_1, z_2, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0) \propto \int_{\epsilon=-\infty}^{\infty} \exp \left(\frac{\left(\epsilon - \frac{\omega_0^2 \left(\sum_{i=1}^n (z_i - \tau) \right) + (\phi^2 + \nu^2)\epsilon_0}{n\omega_0^2 + \phi^2 + \nu^2} \right)^2}{-2\omega_0^2(\phi^2 + \nu^2)/(n\omega_0^2 + \phi^2 + \nu^2)} \right) d\epsilon \cdot \exp \left(\frac{\left(\frac{-\omega_0^4 \left(\sum_{i=1}^n (z_i - \tau) \right)^2 - 2\omega_0^2(\phi^2 + \nu^2)\epsilon_0 \sum_{i=1}^n (z_i - \tau) + n\omega_0^4 \sum_{i=1}^n (z_i - \tau)^2 + \omega_0^2(\phi^2 + \nu^2) \sum_{i=1}^n (z_i - \tau)^2}{(n\omega_0^2 + \phi^2 + \nu^2)^2} \right)}{-2\omega_0^2(\phi^2 + \nu^2)/(n\omega_0^2 + \phi^2 + \nu^2)} \right) \cdot \exp \left(\frac{-(\tau - \tau_0)^2}{2\sigma_0^2} \right) \quad (3.22)$$

Since the integrand is the probability density function of a normal distribution with mean $\frac{\omega_0^2 \sum_{i=1}^n (z_i - \tau) + (\phi^2 + \nu^2)\epsilon_0}{n\omega_0^2 + \phi^2 + \nu^2}$ and variance $\frac{\omega_0^2(\phi^2 + \nu^2)}{n\omega_0^2 + \phi^2 + \nu^2}$, the integral evaluates to one. After combining the exponents and then writing out and combining the terms that contain τ , the equation simplifies to:

$$g(\tau|z_1, z_2, \dots, z_n, \tau_0, \sigma_0, \epsilon_0, \omega_0) \sim \exp \left(\frac{- \left(\tau - \frac{(n\omega_0^2 + \phi^2 + \nu^2)\tau_0 + \sigma_0^2 \sum (z_i - \epsilon_0)}{n\omega_0^2 + \phi^2 + \nu^2 + n\sigma_0^2} \right)^2}{2 \left(\frac{(n\omega_0^2 + \phi^2 + \nu^2)\sigma_0^2}{n\omega_0^2 + \phi^2 + \nu^2 + n\sigma_0^2} \right)} \right) \quad (3.23)$$

This posterior distribution can be rewritten as the normal density function described in equations (3.7), (3.8), and (3.9).

By symmetry of the problem in τ and ϵ , the proof for equations (3.10), (3.11), and (3.12) follows the same steps as above. \square

References

4.1 References for the Introduction

Adams, Richard, Bruce McCarl, Daniel Dudek, and David Glycer. 1988. "Implications of Global Climate Change for Western Agriculture." *Western Journal of Agricultural Economics* 13(2): 348-356.

Cragg, Michael, and Matthew Kahn. 1997. "New Estimates of Climate Demand: Evidence from Location Choice." *Journal of Urban Economics* 42: 261-284.

Dell, Melissa, Benjamin Jones, and Benjamin Olken. Forthcoming. "Climate Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics*.

Deschenes, Olivier, and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review*. March.

Deschenes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3(4): 152185.

Intergovernment Panel on Climate Change. 2007. "Fourth Assessment Report, Climate Change 2007: Working Group I: The Physical Science Basis." http://www.ipcc.ch/publications_and_data/ar4/wg1/en/spmsspm-direct-observations.html. Accessed March 25, 2012.

Mendelsohn, Robert, William Nordhaus, and Daigee Shaw. 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *The American Economic Review* 84(4): 753-771.

Nordhaus, William. 2007. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature* XLV: 686702.

Rozenzweig, Cynthia and Martin Parry. "Potential Impact of Climate Change on World Food Supply." *Nature* 367(6459): 133-138.

Schlenker, Wolfram, and Michael Roberts. 2009. "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change." *Proceedings of the National Academy of Sciences* 106(37): 15594-15598.

Wagner, Gernot, and Richard Zeckhauser. 2011. "Climate Policy: Hard Problem, Soft Thinking." *Climatic Change* 110(3-4): 507-521.

Weitzman, Martin. 1998. "Why the Far-Distant Future Should be Discounted at Its Lowest Possible Rate." *Journal of Environmental Economics and Management* 36(3): 201-208.

Weitzman, Martin. 2007. "A Review of The Stern Review on the Economics of Climate Change." *Journal of Economic Literature* 45(3): 703-724.

4.2 References for Chapter 1

Abadie, Alberto, and Guido Imbens. 2011. "Bias Corrected Matching Estimators for Average Treatment Effects." *Journal of Business and Economic Statistics* 29(1): 1-11.

Abelson, and P. W., and A. Markandya. 1985. "The interpretation of capitalized hedonic prices in a dynamic environment." *Journal of Environmental Economics and Management* 12(3): 195-206.

Bajari, P. and L. Benkard. 2005. "Demand Estimation With Hetrogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach." *Journal of Political Economy* 113(6): 1239-1276.

Bajari, Patrick, Lanier Benkard, and John Krainer. 2005. "Home Prices and Consumer Welfare." *Journal of Urban Economics* 58: 474-487.

Bajari, Patrick, Jane Cooley, Kyoo il Kim, and Christopher Timmins. 2010. "A Theory-Based Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution." Working paper: <http://www.econ.umn.edu/~bajari/research/airhedonic.pdf>

Bajari, Patrick, and Matthew Kahn. 2005. "Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities." *Journal of Business and Economic Statistics* 23(1) 20-33.

Bajari, Patrick, and Matthew Kahn. 2007. "Estimating Hedonic Models of Consumer Demand with an Application to Urban Sprawl." pp.129-155 in Baranzini, Ramirez, Schaerer and Thalmann (eds.) *Hedonic Methods in Housing Markets: Pricing Environmental Amenities and Segregation*, New York: Springer.

Bayer, Patrick, Nathaniel Keohane, and Christopher Timmins. 2009. "Migration and hedonic valuation: The case of air quality." *Journal of Environmental Economics and Management* 58(1): 1-14.

Bayer, Patrick, Robert McMillan, Alvin Murphy, Christopher Timmins. 2008. "A Dynamic Model of Demand for Houses and Neighborhoods." <http://www.econ.as.nyu.edu/docs/IO/9178/bayer.pdf>

Bell, Frederick, and Vernon Leeworthy. 1990. "Recreational Demand by Tourists for Salwater Beach Days." *Journal of Environmental Economics and Management* 18(3): 189-205.

Berger, Alan, Case Brown, Carolyn Kousky, and Richard Zeckhauser. 2009. "The Five Neglects: Risks Gone Amiss." *Mitigating and Recovering from Natural and Unnatural Disasters*. H. Kunreuther and M. Useem, eds. Philadelphia, PA: Wharton School Publishing.

Bin, Okmyung, Craig Landry, Christopher Ellis, and Hans Vogelsong. 2005.

"Some Consumer Surplus Estimates for North Carolina Beaches." *Marine Resource Economics* 20(2): 145-161.

Bin, Okmyung, Thomas Crawford, Jamie Kruse, and Craig Landry. 2008. "Views-capes and Flood Hazard: Coastal Housing Market Response to Amenities and Risk." *Land Economics* 84(3): 434-448.

Bishop, Kelly, and Alvin Murphy. 2011. "Estimating the Willingness-to-Pay to Avoid Violent Crime: A Dynamic Approach." *American Economic Review, Papers and Proceedings* 101(3): 625-629.

Bishop, Kelly, and Christopher Timmins. 2008. "Recovering the Willingness to Pay to Avoid Ozone: How Panel Data Can be Used to Easily Estimate Hedonic Demand Functions." Working paper, available at http://www.ssc.wisc.edu/econ/workshop/bishop_timmins_1.pdf

Bishop, Kelly, and Christopher Timmins. 2008. "Simple, Consistent Estimation of the Marginal Willingness to Pay Function: Recovering Rosen's Second Stage without Instrumental Variables." Working paper, available at http://econ.duke.edu/~timmins/bishop_timmins.pdf.

Black, Sandra. 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *Quarterly Journal of Economics* 114(2): 577-99.

Blackley, Dixie, and James Follain. 1996. "In Search of Empirical Evidence that Links Rent and User Cost." *Regional Science and Urban Economics* 26:409-431.

Brown, Gardner Jr. and Henry Pollakowski. 1977. "Economic Valuation of Shoreline." *Review of Economics and Statistics* 59(3): 272-278.

Capozza, Dennis and Paul Seguin. 1996. "Expectations, efficiency, and euphoria in the housing market." *Regional Science and Urban Economics* 26: 369-385.

Case, Karl, and Robert Shiller. 1989. "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79: 125-37.

Chay, Kenneth, and Michael Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy* 113(2).

Clark, Todd. 1995. "Rents and Prices of Housing Across Areas of the United States: A Cross-Section Examination of the Present Value Model." *Regional Science and Urban Economics* 25:237-247.

Davis, Lucas. 2004. "The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster." *American Economic Review* 94(5): 1693-1704.

Diewert, W. Erwin, Alice Nakamura, and Leonard Nakamura. 2009. "The housing bubble and a new approach to accounting for housing in a CPI." *Journal of Housing Economics* 18(3): 156-171.

DiPasquale, D., and Wheaton, W., 1994. "Housing Market Dynamics and the Future of Housing Prices." *Journal of Urban Economics*. 35, 1-27.

Dougherty, Ann, and Robert Van Order. 1982. "Inflation, Housing Costs, and the Consumer Price Index." *American Economic Review* 72(1): 154-164.

Edwards, Steven and Frank Gable. 1991. "Estimating the Value of Beach Recreation from Property Values: An Exploration with Comparisons to Nourishment Costs." *Ocean and Shoreline Management* 15(1): 37-55.

Ekeland, Ivar, James J. Heckman, and Lars Nesheim. 2004. "Identification and Estimation of Hedonic Models." *Journal of Political Economy* 112(1): S60-S109.

Florida Department of Environmental Protection (FDEP). 2000. "Shoreline Change Rate Estimates: Duval County." <ftp://ftp.dep.state.fl.us/pub/water/beaches/HSSD/reports/duvsre.pdf>.

Florida Department of Environmental Protection (FDEP). 2001. "Shoreline Change Rate Estimates: Bay County." <ftp://ftp.dep.state.fl.us/pub/water/beaches/HSSD/reports/baysre.pdf>.

Florida Department of Environmental Protection (FDEP). 2008a. "Historic Shoreline Data: MHW Files." <ftp://ftp.dep.state.fl.us/pub/water/beaches/HSSD/MHWfiles/New8390/>. Accessed May 10, 2009.

Florida Department of Environmental Protection (FDEP). 2008b. "Consultant's List of Beaches Controlled by Placements of Large Amounts of Sand." <ftp://ftp.dep.state.fl.us/pub/water/beaches/HSSD/MHWfiles/>. Accessed May 10, 2009.

Florida Department of Environmental Protection (FDEP). 2011a. "Digitized Aerial Photographs of Florida Coastline, 2006-2009." Available at <http://publicfiles.dep.state.fl.us/dwrm/beaches/reg-mon/>. Accessed September 8, 2011.

Florida Department of Environmental Protection (FDEP). 2011b. "Critically Eroded Beaches in Florida." Available at: <http://www.dep.state.fl.us/beaches/publications/pdf/CritEroRpt7-11.pdf>.

Florida Department of Revenue. 2011. "2010 Map Data." Available at <ftp://sdrftp03.dor.state.fl.us/>. Accessed 4/23/2011.

Florida State University. 2009. "Beach Erosion Control Monitoring Database." <http://beach15.beaches.fsu.edu/Marco/design2.asp?RowNum=1>. Accessed July 8, 2009.

Gallin, Joshua. 2004. "The Long-Run Relationship between House Prices and Rents." Finance and Economics Discussion Series 200450, Board of Governors of the Federal Reserve System.

Geolytics. 2003. "Neighborhood Change Database: Tract Data from 1970-2000."

Gopalakrishnan, Sathya, Martin Smith, Jordan Slott, Brad Murray. 2010. "The Value of Disappearing Beaches: A Hedonic Pricing Model with Endogenous

Beach Width." Duke Environmental Economics Working Paper Series, Working Paper EE 10-04, September.

Greenstone, Michael and Justin Gallagher. 2008. "Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program." *Quarterly Journal of Economics* 123(3): 951-1003.

Gyourko, Joseph, and Joseph Tracy. 1991. "The Structure of Local Public Finance and the Quality of Life." *The Journal of Political Economy* 99(4): 774-806.

Hapke, Cheryl, Emily Himmelstoss, Meredith Kratzmann, Jeffrey List, and E. Robert Thieler. 2010. "National Assessment of Shoreline Change: Historical Shoreline Change along the New England and Mid-Atlantic Coasts." USGS Open-File Report 2010-1118. Available at <http://pubs.usgs.gov/of/2010/1118/>.

Intergovernmental Panel on Climate Change (IPCC). 2007. "Chapter 10.6.5: Projections of Global Average Sea Level Change for the 21st Century." IPCC Fourth Assessment Report: Climate Change 2007. Available at http://www.ipcc.ch/publications_and_data/ar4/wg1/en/ch10s10-6-5.html.

Kriesel, Warren, Alan Randal, and Frank Lichtkoppler. 1993. "Estimating the Benefits of Shore Erosion Protection in Ohio's Lake Erie Housing Market." *Water Resources Research* 29(4): 795801.

Kueth, Todd, Kenneth Foster, and Raymond Florax. 2008. "A Spatial Hedonic Model with Time-Varying Parameters: A New Method Using Flexible Least Squares." Working paper. <http://ageconsearch.umn.edu/bitstream/6306/2/467673.pdf>

Kuminoff, Nicolai V. and Abdul Salam Jarrah. 2010. "A New Approach to Computing Hedonic Equilibria and Investigating the Properties of Locational Sorting Models." *Journal of Urban Economics* 67(3): 322-335.

Kuminoff, Nicolai, Christopher F. Parmeter and Jaren C. Pope, 2010. "Which Hedonic Models Can We Trust to Recover the Marginal Willingness to Pay for Environmental Amenities?" *Journal of Environmental Economics and Management*, 60(3), 145-160, lead article.

Kuminoff, Nicolai, and Jaren Pope. 2010a. "Hedonic Equilibria, Land Value Capitalization, and the Willingness to Pay for Public Goods." Working paper, available at <http://www.public.asu.edu/~nkuminof/KP10.pdf>

Kuminoff, Nicolai V. and Jaren C. Pope, 2010b. "A Note on Using Quasi-Experiments to Identify Hedonic Demand Parameters." Working paper.

Landry, Craig, Andrew Keeler, and Warren Kriesel. 2003. "An Economic Evaluation of Beach Erosion Management Alternatives." *Marine Resource Economics* 18(2): 105-127

Lee, David, and Thomas Lemieux. 2009. "Regression Discontinuity Designs in Economics." NBER Working Paper 14723. Available at <http://www.nber.org/papers/w14723>

Mankiw, N. Gregory, and David Weil. 1989. "The baby boom, the baby bust, and the housing market." *Regional Science and Urban Economics* 19(2): 235-258.

McCluskey, Jill, and Gordon Rausser. 2001. "Estimation of Perceived Risk and Its Effect on Property Values." *Land Economics* 77(1): 42.

Meese, Richard, and Nancy Wallace. 1994. "Testing the present value relation for housing: should I leave my house in San Francisco." *Journal of Urban Economics* 35(3): 245-266.

Meese, Richard, and Nancy Wallace. 2003. "House Price Dynamics and Market Fundamentals: The Parisian Housing Market." *Urban Studies* 40(5-6): 1027-45.

Mendelsohn, R, D Hellerstein, M Huguenin, R Unsworth, and R Brazee. 1992. "Measuring Hazardous Waste Damages with Panel Models." *Journal of Environmental Economics and Management*. 22:259-271.

Milon, J. Walter, Jonathan Gressel, and David Mulkey. 1984. "Hedonic Amenity Valuation and Functional Form Specification." *Land Economics* 60(4):378-387.

NationalAtlas.gov. 2011. "Profile of the People and Land of the United States." http://nationalatlas.gov/articles/mapping/a_general.html

National Research Council (NRC). 1995. *Beach Nourishment and Protection*. Washington DC: National Academy Press.

NOAA Coastal Services Center. 2009. "Beach Nourishment: A Guide for Local Government Officials." <http://www.csc.noaa.gov/beachnourishment/html/human/socio/geodist.htm>

Pakes, Ariel. 2003. "A Reconsideration of Hedonic Price Indices with an Application to PC's." *The American Economic Review* 93(5): 1578-1614.

Palmquist, Raymond. 1982. "Measuring Environmental Effects on Property Values Without Hedonic Regressions." *Journal of Urban Economics* 11(3): 333-347.

Parsons, George, and Michael Powell. 2001. "Measuring the Cost of Beach Retreat." *Coastal Management* 29(2): 91-103.

Parsons, George, and Yangru Wu. 1991. "The Opportunity Cost of Coastal Land-Use Controls: An Empirical Analysis." *Land Economics* 67(3): 308-317.

Plantinga, Andrew, Ruben Lubowski, and Robert Stavins. 2002. "The Effects of Potential Land Development on Agricultural Land Prices." *Journal of Urban Economics* 52: 561-581

Pompe, Jeffrey, and James Rinehart. 1995. "Beach Quality and the Enhancement of Recreational Property Values." *Journal of Leisure Research* 27(2):143-154.

- Poor, P. Joan, Kevin Boyle, Laura Taylor, and Roy Bouchard. 2001. "Objective versus Subjective Measures of Water Clarity in Hedonic Property Value Models." *Land Economics* 77(4): 482-493.
- Smith, Martin, Jordan Slott, Dylan McNamara, and A. Brad Murray. 2009. "Beach Nourishment as a Dynamic Capital Accumulation Problem." *Journal of Environmental Economics and Management* 58(1): 58-71.
- Riddel, Mary. 2001. "A Dynamic Approach to Estimating Hedonic Prices for Environmental Goods: An Application to Open Space Purchase." *Land Economics* 77(4): 494-512.
- Roback, Jennifer. 1982. "Wages, Rents and the Quality of Life." *Journal of Political Economy* 90(6): 1257-1278.
- Rondinelli, Concetta, and Giovanni Veronese. 2011. "Housing rent dynamics in Italy." *Economic Modelling* 28(1-2): 540-548.
- Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82: 34-55.
- Rosenbaum, P., and D. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70: 4155.
- Rubin, Donald. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66(5): 688-701.
- Titus, James G. 1989. "Chapter 7: Sea Level Rise." *The Potential Effects of Global Climate Change on the United States. Report to Congress.* Washington, D.C.: U.S. Environmental Protection Agency. EPA 230-05-89-052.
- U.S. Army Corps of Engineers. 1996. "Shoreline Protection and Beach Erosion Control Study." IWR Report 96-PS-1.

Van Wellen, E., A. J. Chadwick, and T. Mason. 2000. "A Review and Assessment of Longshore Sediment Transport Equations for Coarse-grained Beaches." *Coastal Engineering* 40(3): 243-275.

Verbrugge, Randal. 2008. "The puzzling divergence of rents and user costs, 1980-2004." *Review of Income and Wealth* 54(4): 671-699.

Western Carolina University. 2009. "Beach Nourishment Database for Florida." <http://www.wcu.edu/1038.asp>. Accessed May 10, 2009.

Woglom, Emily. 2003. *The Dynamics of Shoreline Management: An Approach Using Renewable Resource Economics*. Duke University Masters Thesis. <http://dukespace.lib.duke.edu/dspace/bitstream/10161/275/1/Woglom%20MP%202003.pdf>

Woods Hole Oceanographic Institute Sea Grant. 2000. "Shoreline Change and the Importance of Coastal Erosion." *Focal Points*, April 2000. <http://www.whoi.edu/science/seagrant/education/focalpoints/shoreline.html>

4.3 References for Chapter 2

Agnew, Robert. 2012. "Dire forecast: A Theoretical Model of the Impact of Climate Change on Crime." *Theoretical Criminology* 16(1): 21.

Anderson, Craig. 1989. "Temperature and Aggression: Ubiquitous Effects of Heat on Occurrence of Human Violence." *Psychological Bulletin* 106(1): 74-96.

Anderson, Craig, Brad Bushman, and Ralph Groom. 1997. "Hot Years and Serious and Deadly Assault: Empirical Tests of the Heat Hypothesis." *Journal of Personality and Social Psychology* 73(6): 1213-1223.

Baumeister, Roy, and Todd Heatherton. 1996. "Self Regulation Failure: An Overview." *Psychological Inquiry* 7(1): 1-15.

- Baron, Robert, and Paul Bell. 1976. "Aggression and Heat: The Influence of Ambient Temperature, Negative Affect, and a Cooling Drink on Physical Aggression." *Journal of Personality and Social Psychology* 33(3): 245-255.
- Becker, Gary. 1968. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76: 169-217.
- Brunsdon, Chris, Jonathan Corcoran, Gary Higgs, and Andrew Ware. 2009. "The Influence of Weather on Local Geographical Patterns of Police Calls for Service." *Environment and Planning B: Planning and Design* 36(5): 906-926.
- Bushman, Brad, Morgan Wang, and Craig Anderson. 2005. "Is the Curve Relating Temperature to Aggression Linear or Curvilinear?" *Journal of Personality and Social Psychology* 89(1): 62-66.
- Card, David, and Gordon Dahl. 2011. "Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior." *Quarterly Journal of Economics* 126(4): 1879-1907.
- Cohn, Ellen. 1990. "Weather and Crime." *British Journal of Criminology* 30(1): 51-64.
- Cohn, Ellen, and James Rotton. 2000. "Weather, Seasonal Trends, and Property Crimes in Minneapolis, 1987-1988: A Moderator-Variable Time-Series Analysis of Routine Activities." *Journal of Environmental Psychology* 20(3): 257-272.
- Cragg, Michael, and Matthew Kahn. 1997. "New Estimates of Climate Demand: Evidence from Location Choice." *Journal of Urban Economics* 42: 261-284.
- Deschenes, Olivier, and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review*. March.

Deschenes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3(4): 152185.

Federal Bureau of Investigation (FBI). 2004. "Uniform Crime Reporting Handbook." http://www.fbi.gov/about-us/cjis/ucr/additional-ucr-publications/ucr_handbook.pdf. Accessed March 8, 2012.

Federal Bureau of Investigation (FBI). 2011a. "Uniform Crime Reporting Data, Return A and Supplement to Return A." Compact disc, obtained from the FBI's Multimedia Productions Group.

Federal Bureau of Investigation (FBI). 2011b. "Crime in the United States, 2010: About the Uniform Crime Reporting Program." <http://www.fbi.gov/about-us/cjis/ucr/crime-in-the-u.s/2010/crime-in-the-u.s.-2010/aboutucrmain.pdf>. Accessed March 8, 2012.

Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman. 1996. "Crime and Social Interactions." *Quarterly Journal of Economics* 111(2): 507-48.

Hsiang, Solomon, Kyle Meng, and Mark Cane. 2011. "Civil Conflicts Are Associated with the Global Climate." *Nature* 476: 438-441.

Intergovernmental Panel on Climate Change (IPCC). 2000. "IPCC Special Report: Emissions Scenarios." Available at <http://www.ipcc.ch/pdf/special-reports/spm/sres-en.pdf>.

Intergovernmental Panel on Climate Change (IPCC). 2007. "Working Group I: The Physical Science Basis." Fourth Assessment Report: Climate Change 2007.

Jacob, Brian, Lars Lefgren, and Enrico Moretti. 2007. "The Dynamics of Criminal Behavior: Evidence from Weather Shocks." *Journal of Human Resources* 42(3): 489-527.

Levitt, Steven. 1996. "The Effect of Prison Population Size on Crime Rates: Evidence from Prison Overcrowding Litigation." *Quarterly Journal of Economics* 111(2): 319-351.

Levitt, Steven. 2004. "Understanding Why Crime Fell in the 1990s: Four Factors that Explain the Decline and Six that Do Not." *Journal of Economic Perspectives* 18(1): 163-190.

Maurer, E. P., L. Brekke, T. Pruitt, and P. B. Duffy. 2007. "Fine-resolution climate projections enhance regional climate change impact studies." *Eos, Transactions, American Geophysical Union* 88(47): 504.

Mendelsohn, Robert, William Nordhaus, and Daigee Shaw. 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *The American Economic Review* 84(4): 753-771.

McCollister, Kathryn, Michael French, and Hai Fang. 2010. "The Cost of Crime to Society: New Crime-specific Estimates for Policy and Program Evaluation." *Drug and Alcohol Dependence* 108(1-2): 98-109.

National Climatic Data Center (NCDC) Climate Services Branch. 2011. "Global Historical Climatology Network – Daily." Available at <http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/>. Accessed 4/28/2011.

Rotton, James, and Ellen Cohn. 2003. "Global Warming and U.S. Crime Rates: An Application of Routine Activity Theory." *Environment and Behavior* 35(6): 802-825.

Schlenker, Wolfram, and Michael Roberts. 2009. "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change." *Proceedings of the National Academy of Sciences* 106(37): 15594-15598.

Simister, John. 2002. "Thermal Stress and Violence in India and Pakistan: Investigating a New Explanation of the Kerala Model." *Development Ideas and*

Practices Working Paper DIP-02-01.

Simister, John, and Cary Cooper. 2005. "Thermal Stress in the USA: Effects on Violence and on Employee Behavior." *Stress and Health* 21: 315.

Simister, John, and Evert Van de Vliert. 2005. "Is There More Violence in Very Hot Weather? Tests over Time in Pakistan and Across Countries Worldwide." *Pakistan Journal of Meteorology* 2(4): 55-70.

Tol, Richard. 2009. "The Economic Effects of Climate Change." *Journal of Economic Perspectives* 23(2): 29-51.

U.S. Census Bureau. 1978. "Consolidated File: County and City Data Book." ICPSR: Washington, DC. Available at <http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/07736>. Accessed May 10, 2011.

U.S. Census Bureau. 2004. "Intercensal County Estimates by Age, Sex, Race: 1970-1979, 1980-1989, 1990-1999." Available at <http://www.census.gov/popest/archives/pre-1980/co-asr-7079.html>, <http://www.census.gov/popest/archives/1980s/PE-02.html>, and <http://www.census.gov/popest/datasets.html#cntyinter>. Accessed May 7, 2011.

U.S. Census Bureau. 2011. "Annual Estimates of the Resident Population by Age, Sex, Race, and Hispanic Origin for Counties: April 1, 2000 to July 1, 2009." Available at <http://www.census.gov/popest/counties/asrh/CC-EST2009-alldata.html>. Accessed May 7, 2011.

Viscusi, W. Kip, and Joseph Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World." *Journal of Risk and Uncertainty* 27(1): 5-76.

Weitzman, Martin. 2007. "A Review of The Stern Review on the Economics of Climate Change." *Journal of Economic Literature* 45(3): 703-724.

World Climate Research Programme (WCRP). 2007. "Coupled Model Inter-comparison Project Phase 3 (CMIP3) Multi-model Dataset." Available from the "Bias Corrected and Downscaled WCRP CMIP3 Climate Projections" archive at http://gdo-dcp.ucllnl.org/downscaled_cmip3_projections/. Accessed March 10, 2012. ²³

4.4 References for Chapter 3

Battaglini, Marco. 2004. "Policy Advice with Imperfectly Informed Experts." *Advances in Theoretical Economics* 4(1): 1-32.

Berkson, Joseph. 1950. "Are There Two Regressions?" *Journal of the American Statistical Association* 45: 164-180.

Borick, Christopher and Rabe, Barry. 2010. "A Reason to Believe: Examining the Factors that Determine Individual Views on Global Warming." *Social Science Quarterly* 91: 777-800.

Buser, Christoph. 2009. "Bayesian Statistical Methods for the Analysis of Multi-model Climate Predictions." PhD dissertation, Eidgenössische Technische Hochschule, Zurich. <http://stat.ethz.ch/~buser/CV/buser2009.pdf>

Buser, C., H. Kunsch, D. Luthi, M. Wild, and C. Schar. 2009. "Bayesian Multi-model Projection of Climate: Bias Assumptions and Interannual Variability." *Climate Dynamics* 33: 849-868

Buser, C., H. Kunsch, and C. Schar. 2010. "Bayesian Multi-model Projections of Climate: Generalization and Application to ENSEMBLES results." *Climate Research* 44: 227-241.

²³I acknowledge the modeling groups, the Program for Climate Model Diagnosis and Intercomparison, the WCRP's Working Group on Coupled Modeling, and the U.S. Department of Energy Office of Science for their roles in making available and supporting the WCRP CMIP3 multi-model dataset.

- Cameron, Trudy. 2005. "Updating Subjective Risks in the Presence of Conflicting Information: An Application to Climate Change." *Journal of Risk and Uncertainty* 30(1): 63-97.
- Casanova, Mara, Pilar Iglesias, Heleno Bolfarine, Victor Salinas, and Alexis Pena. 2010. "Semiparametric Bayesian Measurement Error Modeling." *Journal of Multivariate Analysis* 101(3): 512-524.
- Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50(6): 1431-1451.
- Cyert, Richard, and Morris Degroot. 1974. "Rational Expectations and Bayesian Equilibrium." *Journal of Political Economy* 82: 521-535.
- Dellaportas, Petros, and David Stephens. 1995. "Bayesian Analysis of Errors-in-variables Regression Models." *Biometrics* 51(3): 1085-1095.
- Deryugina, Tatyana. 2011. "How do people update? The effects of local weather fluctuations on beliefs about global warming." Working paper.
- Egan, P., and Mullin, M. 2010. "How Citizens Integrate Information without Ideological Cues: Local Weather and Americans' Beliefs about Global Warming." NYU working paper.
- Feldman, M. 1991. "On the Generic Non-Convergence of Bayesian Actions and Beliefs." *Economic Theory* 1: 301-321.
- Hamilton, Lawrence. 2011. "Climate Change Partisanship, Understanding, and Public Opinion." Carsey Institute Issue Brief 26, Spring 2011.
- Hammitt, J.K., R.J. Lempert, and M.E. Schlesinger. 1992. "A Sequential-Decision Strategy for Abating Climate Change." *Nature* 357: 315-318.

Haslett, J., M. Whitley, S. Bhattacharya, M. Salter-Townshend, S. Wilson, J. Allen, B. Huntley, F. Mitchell. 2006. "Bayesian palaeoclimate reconstruction." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 169(3): 395-438.

Intergovernment Panel on Climate Change. 2007. "Fourth Assessment Report, Climate Change 2007: Working Group I: The Physical Science Basis." http://www.ipcc.ch/publications_and_data/ar4/wg1/en/spmssp-direct-observations.html. Accessed March 25, 2012.

Jun, M., R. Knutti, and D. Nychka. 2008. "Spatial Analysis to Quantify Numerical Model Bias and Dependence: How Many Climate Models Are There?" *Journal of the American Statistical Association* 103:934-947.

Leach, Andrew. 2007. "The Climate Change Learning Curve." *Journal of Economic Dynamics and Control* 31(5): 1728-1752.

Karp, Larry, and Jiangfeng Zhang. 2006. "Regulation with Anticipated Learning about Environmental Damages." *Journal of Environmental Economics and Management* 51: 259-279.

Kellstedt, Paul, Sammy Zahran, and Arnold Vedlitz. 2008. "Personal Efficacy, the Information Environment, and Attitudes Toward Global Warming and Climate Change in the United States." *Risk Analysis* 28(1):113-126.

Kelly, David, and Charles Kolstad. 1999. "Bayesian Learning, Growth, and Pollution." *Journal of Economic Dynamics and Control* 23: 491-518.

Kelly, David, David Letson, Forest Nelson, David Nolan, and Daniel Solis. 2009. "Evolution of Subjective Hurricane Risk Perceptions: A Bayesian Approach." University of Miami, Department of Economics, Working Papers: 0905..

Krishna, Vijay, and John Morgan. 2001. "A Model of Expertise." *Quarterly Journal of Economics* 116(2): 747-775.

Li, Bo, Douglas W. Nychka, and Caspar M. Ammann. 2010. "The Value of Multi-proxy Reconstruction of Past Climate." *Journal of the American Statistical Association* 105(491): 883-895.

Millner, Antony, Simon Dietz, and Geoffrey Heal. 2010. "Ambiguity and Climate Policy." National Bureau of Economic Research, Working Paper 16050.

Newbold, Stephen, and Adam Daigneault. 2009. "Climate Response Uncertainty and the Benefits of Greenhouse Gas Emissions Reductions." *Environmental and Resource Economics* 44:351-377.

Nordhaus, William, and Joseph Boyer. 2000. *Warming the World*. MIT Press, Cambridge, MA.

Pew Research Center for People and the Press. 2011. "More Moderate Republicans See Evidence of Warming Modest Rise in Number Saying There Is 'Solid Evidence' of Global Warming." <http://www.people-press.org/files/legacy-pdf/12-1-11%20Global%20warming%20release.pdf>. Accessed March 25, 2012.

Pratt, John, Howard Raiffa, and Robert Schlaifer. 1995. *Introduction to Statistical Decision Theory*. The MIT Press: Cambridge, MA.

Rannala, Bruce. 2002. "Identifiability of Parameters in MCMC Bayesian Inference of Phylogeny." *Systematic Biology* 51(5): 754-760.

Richardson, Sylvia, and Walker Gilks. 1993. "A Bayesian Approach to Measurement Error Problems in Epidemiology Using Conditional Independence Models." *American Journal of Epidemiology* 138(6): 430-442.

Smith, Richard, Claudia Tebaldi, Doug Nychka, and Linda Mearns. 2009. "Bayesian Modeling of Uncertainty in Ensembles of Climate Models." *Journal of the American Statistical Association* 104(485): 97-116.

Tebaldi, C., R. Smith, D. Nychka, and L. Mearns. 2005. "Quantifying Uncertainty in Projection of Regional Climate Change: a Bayesian Approach to the Analysis of Multimodel Ensembles." *Journal of Climate* 18: 1524-1540.

Viscusi, W. Kip, and Charles O'Connor. 1984. "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?" *American Economic Review* 74(5): 942-956.

Viscusi, W. Kip, and Richard Zeckhauser. 2006. "National Survey Evidence on Disasters and Relief: Risk Beliefs, Self-Interest, and Compassion." *Journal of Risk and Uncertainty* 33(1-2): 13-36.

Wagner, Gernot, and Richard Zeckhauser. 2011. "Climate Policy: Hard Problem, Soft Thinking." *Climatic Change* 110(3-4): 507-521.

Weitzman, Martin. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics* 91(1): 1-19.

Zimper, Alexander, and Alexander Ludwig. 2009. "On Attitude Polarization under Bayesian Learning with Non-additive Beliefs." *Journal of Risk and Uncertainty* 39(2): 181-212.