

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

NON-MARKET ADAPTATIONS TO CLIMATE CHANGE AND RECREATIONAL POLICY
EVALUATION: THREE ESSAYS IN ENVIRONMENTAL AND APPLIED
MICROECONOMICS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

SEAN THOMAS O'CONNOR

Norman, Oklahoma

2021

NON-MARKET ADAPTATIONS TO CLIMATE CHANGE AND RECREATIONAL POLICY
EVALUATION: THREE ESSAYS IN ENVIRONMENTAL AND APPLIED
MICROECONOMICS

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF ECONOMICS

BY THE COMMITTEE CONSISTING OF

Dr. Kevin Kuruc, Chair

Dr. Daniel Hicks, Co-Chair

Dr. Myongjin Kim

Dr. Jonathan McFadden

Dr. Pallab Ghosh

Dr. Deven Carlson

© Copyright by SEAN THOMAS O'CONNOR 2021
All Rights Reserved.

DEDICATION

to:

my wife Cassie and my parents, Josie and Pat.

Acknowledgements

I'd like to express my gratitude to all the members of my committee. First I would like to thank Dr. Kevin Kuruc for his consistent support and advise. This was especially true in this last year, when the stresses of the pandemic and job market were most intense. I also want to thank Dr. Dan Hicks for taking me on as an advisee despite his other commitments.

Dr. Myongjin Kim was the first faculty member to respect and trust me as a researcher, and for that I'll forever be grateful. Dr. Jonathan McFadden challenged me to become the best climate researcher I could be, and any expertise that I have developed in this field is in large part because of his advising. Dr. Pallab Ghosh opened my eyes to work in applied econometrics and was a tremendous instructor to TA for. Dr. Deven Carlson generously stepped in on short notice to serve as the outside member on the committee, which was a huge relief.

I'd also like to thank my sister Jackie, my dog Molly, and TPTS for keeping me honest.

Finally, I'd like to thank my uncle, Dr. John Dalzell, for encouraging me to pursue this degree in the first place.

Contents

1	Geographic Heterogeneity in Climate Change Adaptation: Behavioral Evidence from Participation in Outdoor Activities	1
1.1	The Weather, Climate, and Americans' Daily Activities	5
1.1.1	The Difference Between Weather and Climate	5
1.1.2	Data	7
1.1.3	Empirical Justification	10
1.2	Empirical Strategy	12
1.2.1	Estimation Strategy	13
1.3	Results: How Weather and Climate Explain Leisure Behavior	15
1.3.1	First Stage	16
1.3.2	Second Stage	17
1.3.3	Robustness Checks	18
1.4	Outdoor Leisure in a Changing Climate	19
1.4.1	Forecasting the Effect of Climate Change on Outdoor Time	20
1.4.2	Back of the Envelope Welfare Calculations	21
1.5	Conclusion	22
2	Weather, Climate, and the Police in America	42
2.1	Productivity, Conflict, and the Weather	46
2.2	Empirical Strategy	48
2.2.1	Data Description	48
2.2.2	Estimation Strategy	50
2.3	Results: How Weather Changes Police Productivity and Behavior	54
2.3.1	Change in the Number of Foot and Traffic Patrol Stops	54
2.3.2	Change in the Speed Drivers are Stopped For	55
2.3.3	Change in the Veil of Darkness	56
2.3.4	Robustness Checks and Alternative Specifications	57
2.4	Conclusion	57
3	Retail Marijuana Deregulation and Housing Prices	71
3.1	Background	75
3.1.1	Medical and Recreational Marijuana Legalization	75
3.1.2	The Housing Market Connection	77
3.1.3	Mechanisms	78
3.2	Data	81

3.2.1	Housing Data	81
3.2.2	Marijuana Laws	82
3.2.3	Dispensary Data	84
3.3	Empirical Strategy	84
3.4	Results	87
3.4.1	Housing Prices Following Statewide Marijuana Legalization	87
3.4.2	Spatial Model	90
3.4.3	Robustness Checks	93
3.5	Conclusion	94
3.6	Tables	96
	Appendix A	122
	Appendix B	124
	Appendix C	127

List of Tables

1.3	CARE Stage 2	29
2.1	Proportion of Drivers Pulled Over Who are Black	59
2.4	Veil of Darkness Models	63
3.1	Summary Statistics	96
3.2	Marijuana Legalization Laws	97
3.3	Effect of Marijuana Legalization on House Price per Sq. Foot (OLS)	98
3.4	Effect of Marijuana Legalization on Home Price (Fixed Effects)	99
3.5	Summary Statistics by Spatial Difference-in-Difference Treatment	100
3.6	Spatial Difference-in-Differences	101
A1	Effect of Recreational Marijuana Legalization on House Price per Sq. Foot	128
A2	Effect of Recreational Marijuana Legalization on House Price (Robustness)	129
A3	Heterogeneous Effect of Marijuana Legalization on House Price across Q_τ	130
A4	Heterogeneous Effect of Marijuana Legalization on House Price across Q_τ	131
A5	Spatial Difference-in-Differences: Colorado Subsample	132
A6	Spatial Difference-in-Differences: Washington Subsample	133

List of Figures

1.1	Outdoor Activities More Pleasant	30
1.2	Simulation of the Temperature Distribution Shifting 3 Celsius	31
1.3	Annual Number of Days with Average Temperature in 5 F Bins Under Three Climate Scenarios	32
1.4	The Share of Time Spent Outside by Activity Group	33
1.5	Estimated Change in the Number of Non-Work Minutes Spent Outside in Cities in the Bottom and Top Quintiles of the Temperature Distribution	34
1.6	Daily Share of Non-Work Time Spent Outside	35
1.7	Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time	36
1.8	Stage 1 Temperature Response Function Across Cities	37
1.9	Placebo Exercise: No Pattern for Shopping	38
1.10	Predicted Second Stage Response in a Synthetic Cold and Hot City	39
1.11	Percentage Change in the Amount of Time Spent Outside	40
1.12	Change in the U-Index Under the RCP4.5 Scenario	41
2.1	The Time Drivers of Different Races Get Pulled Over	64
2.2	States, Counties, and Cities in the Sample	65
2.3	Average Annual Temperature by County	66
2.4	Change in the Number of Pullovers Per Week	67
2.5	Change in the Number of Pullovers Per Week – Precipitation	68
2.6	Change in the Speed Pullover	69
2.7	Change in the Number of Pullovers Per Week	70
3.1	Trend of Average House Price Per Sq. Foot (CO, WA)	102
3.2	Trend of Average House Price Per Sq. Foot (OR)	103
3.3	Trend of Average House Price Per Sq. Foot (CA, MA, NV)	104
3.4	Illustration of Spatial Difference in Difference Model in Denver, Colorado	105
3.5	106
3.6	107
.0A1	Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time with 3-Day Temperature Lags	122
.0A2	Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time with No Fixed Effects	123

Abstracts

The first chapter examines the role of non-market behavioral adaptation to climate change in the United States for the case of outdoor leisure with a novel estimation procedure that accounts for both short-run weather and long-run climate adjustments. First, I comprehensively review the temperature sensitivity of all activities in the American Time Use Survey using a flexible non-linear estimation procedure. Predictably most activities are found to be unresponsive to temperature with the exception of those that take place outdoors. Time spent outdoors is studied further using the Climate Adaptation Response Estimation approach, which allows for temperature responses to vary geographically. I find the sensitivity to temperature varies across the country, and this variation is especially pronounced for cold-weather cities in which inhabitants modify their outdoor and physical activities in response to temperatures more than warm-weather cities. Simulating the expected change in outdoor activity time using climate models compiled by the Intergovernmental Panel on Climate Change implies a large increase in outdoor time driven by warmer winters.

In the second chapter, I explore how policing in the United States is influenced by weather and climate. Policing is under renewed scrutiny following a number of high-profile murders at the hands of law enforcement officers. I use data from the Stanford Open Policing Project, which includes over 200 million traffic and pedestrian stops between 2011-2019, to estimate the sensitivity of police productivity to weather. Results from non-linear temperature response functions suggest a ten percent increase in the number of stops police conduct on the hottest days. These results are further supported with estimates showing that the speeds at which drivers are stopped for speeding violations are five percent lower on hot days relative to cool days. Finally, using the “veil of dark-

ness” approach, I find that the gap in the rate of pullovers of black drivers during daylight versus at night increases on hot days. Taken together, these results suggest that police activity is influenced to a large degree by the weather.

The third chapter investigates the consequences of marijuana legalization on housing markets in the United States. Despite federal law, twelve American states and Washington D.C. have legalized recreational marijuana since 2012. Using a national housing data set from the online real estate listing database Zillow.com, we identify the cross and inter-state effects of marijuana legalization on house prices in different points of the price distribution function. We find positive effects upwards of ten percent in the top half of the price distribution following successful legalization ballot initiatives, and between five and fifteen percent across the distribution after the state enacts the ballot initiative and the first legal sales take place. A spatial difference-in-differences model reveals that within Colorado and Washington, prices in neighborhoods with new dispensary openings nearby experience a seven percent price appreciation. To summarize, our results suggest that there are second order benefits to marijuana legalization.

Chapter 1

Geographic Heterogeneity in Climate

Change Adaptation: Behavioral Evidence from Participation in Outdoor Activities

Globally, the past five years mark the five hottest since record keeping began in 1880, with 2020 on track to be the hottest ever.¹ As concerns about warming have increased, the economics literature has made major progress in evaluating the costs of climate change. Early work on the topic studied the effect on agriculture (e.g. Mendelsohn et al. [1994]), but the literature has expanded to include the effect on health and mortality (Deschênes and Greenstone [2011], Heutel et al. [2017], Barreca et al. [2016]), energy consumption (Davis and Gertler [2015], Mansur et al. [2008], Auffhammer [2018]), and worker productivity (Adhvaryu et al. [2019], Colmer [2019]), among others.²

Within the climate economics literature, few causal estimates exist of the role geographic heterogeneity plays in determining outcomes. In the United States, regions are expected to be exposed to different extreme weather events that will increase in frequency and intensity as climate change progresses. Wildfires in the West, hurricanes in the South and East, and droughts in the Midwest are all examples of headline weather events that fit this criteria and are well-understood by the

public. Despite the prominence of these events, the broader and more gradual effects of climate change will also vary across regions. By expanding our understanding of the geographic heterogeneity of climate responses, we get a more complete picture of the welfare effects of climate change. This research studies the role geography and climate play in Americans' participation in outdoor nonmarket activities with data from the American Time Use Survey (ATUS).

To this end, I estimate the sensitivity of outdoor activities to both short-run weather and long-run climate in cities across the United States and project the results forward to predict future welfare changes. My findings formalize intuition about outdoor activities: doing things outdoors is more pleasant than indoors, and because outdoor activities are sensitive to hot and cold temperatures, American cities with low average temperatures will gain days with pleasant outdoor weather due to warmer winters. This stands in contrast to cities with warm average temperatures, which don't gain from warmer winter days but may lose outdoor time due to extreme hot summer days. To complete the analysis, three questions must be asked. First, which household activities are sensitive to temperature? The activity set is narrowed to those that exhibit some response to temperature, leading to the second question: do these temperature sensitivities vary by local climate? I estimate the temperature response function for outdoor nonmarket activities separately for each city in the United States. Finally, having recovered the city-level response to temperature, the final question is: how will activity time allocation in the United States change when the climate distribution shifts? This is done using climate prediction models from the Intergovernmental Panel on Climate Change (IPCC). My research estimates which activities are sensitive to temperature, how these sensitivities vary by local climate, and how we should expect the amount of time dedicated to these activities to change with the climate.

The non-linear model suggests that most activities Americans spend time on are not influenced by the weather. Although respondents do not adjust their time spent doing particular activities, they do adjust where they do them. Consistent with results in Zivin and Neidell [2014], I find that outdoor activities are very sensitive to temperature. This intuitive result underscores the importance of considering a wider range of outcomes when thinking about climate change. Rather than merely

being an extension of past work, these responses likely have large implications for human welfare. Figure 1.1 plots the share of time certain activity groups are unpleasant, based on the U-Index developed by Kahneman and Krueger [2006]. Non-market activities are less unpleasant when they take place outdoors; averaged across all activities, the U-Index is two-thirds lower outdoors than in. The public health and psychology literature find similar positive benefits associated with spending time outdoors.³ If climate change alters the indoor-outdoor margin through changes in the weather, there could be large welfare implications – either warmer summers induce more time inside or warmer winters induce more time outside.

Those living in warm-weather cities may respond differently in ways relevant for understanding climate adaptation. To capture the climate effect, I use a novel estimation procedure developed by Auffhammer [2018] which allows for agents' temperature sensitivity to vary across geography. The method, referred to as Climate Adaptation Response Estimation (CARE), models the relationship between climate and economic outcomes in two stages. In the first stage, CARE estimates a non-linear temperature response function. This non-linear approach has become common in the climate literature because it relaxes parametric assumptions regarding the temperature response function. The key innovation of the first stage is that the function is estimated separately for each geographic unit in the sample. Due to the large number of cities in the ATUS sample and the non-linear functional form, the first stage produces a large number of estimated temperature coefficients. In the second stage, these estimated coefficients are regressed on the long-run temperature averages, which reflect the local climate. As a result, the first stage can be thought of as the response to weather and the second stage how sensitive the weather response is based on the local long-run climate. The estimated temperature response function from the first stage is then projected forward using climate models aggregated by NASA, which allows for simulation of the potential welfare effects of future climate change under two different emissions scenarios.

The results of the climate model illustrate significant geographic heterogeneity in the United States. There are three primary results from the CARE method and its projection: first, the temperature response function from stage one reflects a distinct inverse-U shape. That is, people spend

less time outdoors at both hot and cold temperatures relative to a baseline “nice” day in the 60-70 degree Fahrenheit range. This suggests that individuals avoid extreme temperatures at both the high and low end of the temperature distribution. Then the estimated coefficients from the first stage are regressed on the long-run climate of their respective metropolitan area. My results demonstrate that metropolitan areas with large shares of cool days are generally more temperature sensitive than those with a greater share of warm days. While both cold and warm climate cities decrease the amount of time spent outside when it is very cold and very hot (i.e. the first stage), cold climate cities have more extreme responses.

This result is the first contribution of my research, suggesting that agents in cold climate cities are more sensitive to temperature than their warm climate counterparts. The stage one estimates measured a temperature response function – i.e. how activity time allocation changed as a result of short term weather events. The weather is a particular day’s temperature, precipitation, or humidity. These weather results are informative, but do not reflect how climate may mediate these outcomes. For example, the climate of Minnesota is quite different from the climate of Florida. If Floridians are less sensitive to hot days, these results can teach us about how the response of Minnesota may change when Minnesota warms. Similar research estimating the relationship between outdoor activities and the weather take a non-linear approach but do not provide estimates of the climate effect.

The most similar research to this exercise is Chan and Wichman [2020], which uses a non-linear fixed effects model to estimate how recreational cycling responds to weather fluctuations. The authors find cyclists are more sensitive to cold temperatures than warm, similar to the results presented here. These results should not be a surprise, as human physiology bounds the range of temperatures that we find desirable (Arens and Bosselmann [1989]; Höpfe [2002]; Stathopoulos et al. [2004]). Importantly, the results of this research demonstrate heterogeneous preferences for temperature across the United States.⁴ In a time use setting, the reallocation of time from outdoors to indoors and from active to passive will have a profound effect on welfare as different regions warm at different rates.

The second contribution is to provide evidence for the relationship between weather and the universe of non-market activities in the American Time Use Survey. The American Time Use Survey provides a nationally representative sample of agents' time use patterns and is not restricted to a particular subset of activities. Empirical exercises of time use surveys date back to the 1970s (Gronau [1976a], Gronau [1976b]), and the theory of time use and allocation to the 1960s (Becker [1965], Johnson [1966]), but my research is the first to provide comprehensive estimates of temperature response functions for all the activities included in the time use survey. I exploit within-city variation in weather to identify the non-linear relationship between temperature and the universe of activities in the American Time Use Survey.

The paper proceeds as follows. Section ?? clarifies the distinction between weather and climate as it pertains to the econometric model, details the data being used, and motivates the research with a descriptive empirical exposition. Section 1.2 outlines the two-step estimation procedure used in the CARE model. Section 3.4 estimates both the short-run temperature response function and the long-run climate adaptation model. Section 1.4 projects the results of the model forward using General Circulation Models (GCMs) from NASA, and Section 3.5 concludes.

1.1 The Weather, Climate, and Americans' Daily Activities

This section discusses of the difference between the weather and climate, which is key to understanding the CARE methodology. Then, to establish a relationship between the weather and activity time allocation, a series of non-linear temperature response functions are estimated using data from the American Time Use Survey. Outdoor activities stand out for their relationship with the weather, leading to the climate analysis in the following section.

1.1.1 The Difference Between Weather and Climate

The climate literature in economics has taken two broad approaches to estimating the effect of climate and economic outcomes. The first, established in Mendelsohn et al. [1994], uses long-run

values of a region's weather outcomes as a proxy for the local climate. The economic outcome of interest is regressed on these long-run values (typically the seasonal or monthly averages of the preceding two-to-three decades) cross-sectionally. The alternative approach, which has increased in popularity in the last decade, uses fluctuations in weather to identify the relationship between the outcome and climate. Identification requires panel data in which the researcher observes regions repeatedly over time. The distinction between these two approaches is that the former explicitly models the climate effect, whereas the latter estimates the weather effect with an assumption that this is informative for climate responses. Understanding the difference is important in interpreting the results of the two stage CARE approach used in this paper.

Early research measuring climate adaptation estimated the effect of weather on different outcomes using cross-sectional data. Mendelsohn et al. [1994], for example, regressed county farm prices in the United States on monthly average temperatures. This approach has the advantage of estimating the true climate effect since farmers are aware of prior climate information at their location, allowing them to optimize production and investment, and therefore are able to maximize returns for their property on the market. The major disadvantage of the cross-sectional approach — and the reason it has fallen out of favor in the climate literature — is that it is susceptible to omitted variable bias. By estimating the temperature response function with panel data and fixed effects, and then recovering the long-run climate effect using cross-sectional data and the estimated coefficients from the first stage, CARE recovers the climate effect while avoiding the pitfalls typically associated with cross-sectional data.

The key idea in the panel data-weather fluctuations approach is that weather observations are drawn from the climate distribution. To demonstrate this idea, Figure 1.2 plots a simulation of the climate distribution under the current climate and a future climate where the mean temperature increases by 3 Celsius (5.4 Fahrenheit). This is a hypothetical scenario, as the true temperature distribution for the United States is not as neatly normal as the distributions depicted due to regional differences. Any individual day's temperature (or precipitation, humidity, etc.) is the realized value of the climate distribution. Figure 1.2 assumes that the two distributions are different due to the

nature of time and expected changes in the climate, but the two could distributions could also reflect the difference in climate between regions. For example, the current climate could instead represent the cold climate of Boston and the future climate could represent the warm climate of Los Angeles.

A common approach in the climate literature estimates an outcome variable on temperature with location and time fixed effects to adjust for unobservable characteristics unique to a place and time. The downside of such models is that fixed effects impose that the estimated temperature response function is the same for each location in the sample (Auffhammer and Schlenker [2014]). As a result, the coefficients reflect only the response to deviation in weather, not climate. The practical implication is that Minnesotans and Floridians are assumed to have the same reaction to a 95 day.⁵ Much of the recent climate literature has taken this approach despite its limitations.

The CARE method marries the two approaches by estimating both the short-run weather effect and the long-run climate effect, although it is not the first to consider geographic heterogeneity when estimating climate adaptation. Barreca et al. [2016] estimate the relationship between temperature and monthly mortality rates separately for each of the nine Census regions, reporting results for each individually. Similarly, Butler and Huybers [2013] regress maize yields on growing and killing degree days in approximately 1000 counties in the United States. Heutel et al. [2017] take a slightly different approach, interacting a non-linear temperature response function with indicators for whether a U.S. ZIP code falls in the top/middle/bottom climate tercile. The advantage of CARE over these other approaches is that it estimates the temperature response function for each geographic unit using panel data in the first stage and then allows the estimated coefficients to vary cross-sectionally in a second stage.

1.1.2 Data

A number of data sources are used in the analysis, including time use survey data from the U.S. Bureau of Labor Statistics, the Current Population Survey from the United States Census Bureau, and meteorological data from the University of Idaho.

Time Use Data

The primary source of data comes from the Census Bureau and the Bureau of Labor Statistics' American Time Use Survey (ATUS). The ATUS randomly selects one member of a household that has completed eight consecutive months of the Current Population Survey (CPS) to fill out a time use diary. This diary asks that the respondent log the amount of time they spent completing all activities over a 24-hour period. Importantly for this research, the respondents are asked where each of their activities takes place, which helps distinguish between weather sensitive and weather insensitive activities.

Activities are first separated by the categories defined by the BLS. The categories include household activities, education, work and work-related activities, socializing, and others. These groups are quite broad. For example, household activities include meal preparation and gardening. Following Zivin and Neidell [2014], any activity that takes place "outdoors, away from home" or makes references to the exterior of the home, such as "gardening" or "exterior maintenance" is encoded as an outdoor activity. As Zivin and Neidell note, some activities are said to take place "at the home or yard," but since this categorization is ambiguous, it is encoded as not taking place outdoors.

Table 1 presents the summary statistics for each activity group in the ATUS data and for all activities which are coded as occurring outside. Six categories of activities are considered: household labor, market labor, market consumption activities, leisure activities, outdoor activities, and miscellaneous activities. The sub-categories correspond to the 18 major categories defined by the BLS, plus five outdoor activities created for the purpose of this research. Both the unconditional number of minutes spent doing an activity and the number of minutes conditional on participation (i.e. the number of minutes > 0) are included. The three largest activities by time allocation are sleeping, engaging in market labor, and leisure time. Only about a third of respondents participate in outdoor activities but, conditional on participation, respondents average more than ninety minutes outdoors per day. To get a better idea of which activities take place outdoors, Figure 1.4 plots the share of time each of the major activity categories takes place outdoors. Household activities,

including activities such as yard maintenance and pet care, and sports and exercise take place outdoors approximately twenty-five percent of the time, by far the largest out of the major activity groups.

Because the ATUS is administered to individuals who have completed eight rounds of the CPS, the ATUS data can be matched to the CPS data at the individual level. A large number of variables from the CPS are used, including the respondent's age, sex, race, educational attainment, household income, marital status, the number of children, and homeownership status.

Meteorological Data

Meteorological data are obtained from the University of Idaho's gridMET dataset [Abatzoglou, 2013]. gridMET provides daily ground-level meteorological data at a 4-km spatial resolution. The sample used in this research starts at the beginning of 1979 (the earliest year available in gridMET) and ends in 2018. From 1979 to 2004, the only variables derived from gridMET are maximum and minimum daily temperature. As is common in the climate literature, these two variables are averaged to produce daily mean temperature. Starting in 2005 other meteorological measures, including maximum and minimum relative humidity, mean wind speed, and the precipitation amount, are also imported from gridMET.

The most disaggregated geographic level in the ATUS data is the Census' metropolitan statistical area (CBSA). A CBSA TIGER shapefile from the Census is used to merge the gridMET and ATUS data. Since the gridMET is measured continuously in space, the values in each CBSA must be aggregated. For each meteorological variable, the median value in a CBSA is extracted from the continuous gridMET data [Dorman, Rush, Hough, Russel, and Karney, 2020]. The value extraction derives meteorological data for each CBSA-day in the ATUS data.

Miscellaneous Data

In addition to the time use and meteorological data, there are a handful of other data sources used. Elevation data is from the R package "elevatr" [Hollister and Shah, 2020], which provides ac-

cess to Amazon Web Services' Terrain Tiles via an API. End-of-century climate projections come from the NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP). Twenty-one General Circulation Models (GCMs) conducted under the Coupled Model Intercomparison Project Phase 5 (CMIP5) provide daily maximum temperature, minimum temperature, and precipitation. I calculate daily mean temperature for each metropolitan-date between 2080-2099. Creating these daily means allows me to match the expected mean temperature in 2080-2099 to the day that a respondent in the ATUS filled out their diary. NEX-GDDP includes projections for two Representative Concentration Pathways (RCPs): RCP4.5 and RCP8.5. RCP4.5 trajectories model emissions as peaking in 2040 and declining after that. In RCP8.5 models, emissions continue to rise throughout the 21st Century. RCP8.5 models are typically referred to as "business as usual" models. Figure 1.3 shows the number of days per year in temperature bins of 5 under three climate scenarios. The light green color reflects the current climate in the United States, the middle green shade the RCP4.5 scenario, and the darkest green shade the RCP8.5 scenario. There is a clear rightward shift in the average temperature distribution. A large number of cold-weather days in the current climate are lost in both RCP scenarios, while there is a marked increase in the number of days in the top two temperature bins.

1.1.3 Empirical Justification

Before estimating the full model, we first need to determine what (if any) household activities are weather sensitive. This first exercise demonstrates that – unsurprisingly – outdoor and physical activities are most sensitive to the weather.

It is assumed that temperature affects some activities more intensely than others. For example, due to the prevalence of air conditioning in the United States, we should not expect that passive indoor activities such as watching television are influenced by weather conditions outside. The ATUS enumerates a large number of activities over which respondents participate. This allows for an analysis of the temperature sensitivity of different activities. This temperature response function

can be represented by:

$$y_{it} = \sum_{b \in B \setminus \{60-70\}} \beta_b D_{ibt} + \gamma X_{ijt} + \alpha_{jm} + \phi_{my} + \epsilon_{it} \quad (1.1)$$

where y_{it} represents the number of minutes respondent i spent doing activity y at time t and D_{ibt} is an indicator for whether the temperature experienced by respondent i at time t falls into temperature bin b . The model allows for non-linearities in the temperature-outcome relationship by including these indicators. X_{ijt} is a vector of confounding variables which are adjusted for, including other weather variables (precipitation, humidity, and wind speed), demographic characteristics of respondent i from the Current Population Survey (age, race, marital status, number of children, education, household income, home ownership status), ATUS diary information (day of the week the diary was completed for, an indicator for whether the diary day was a holiday) and region j geographic and economic characteristics (elevation, slope, air conditioning prevalence). α_{jm} are region-month fixed effects and ϕ_{my} are month-year fixed effects. The set of fixed effects absorb unobserved region-month and month-year determinants of time allocation.

Table 2 displays the results of the model. The coefficients should be interpreted as the change in the number of minutes for a particular activity as the result of the daily temperature falling into one of eight temperature bins, relative to a baseline day of 60-70. Broadly there are five categories considered: household labor, market labor, consumer activities, leisure, and outdoor activities. The first four categories are exclusive of one another, but outdoor activities is inclusive of the others. For example, the sub-category “exercising outside” is a subset of the “exercising” category in the leisure group.⁶ Looking at the results in Table 2, it is clear that activities which either occur outside or are physically active are those that are most sensitive to temperature. All of the outdoor activities have significant negative coefficients, especially in the bottom half of the temperature distribution. The effects are less pronounced at the top of the distribution however. Physical activities, including exercising and participating in sports similarly show significant negative effects when the weather is the coldest but little when it is hottest. Interestingly, general leisure time increases by a large

amount in the lower bins.

Having established that the number of minutes spent outdoors is particularly susceptible to temperature, Figure 1.5 estimates the outdoor time response function for metropolitan areas in the bottom and top quantile of the temperature distribution separately. This approach is common in the climate literature.⁷ The quantiles here were chosen to demonstrate potential divergences in the temperature response function as a function of local climate. Figure 1.5 shows that the coldest metropolitan areas are consistently more temperature sensitive than the hottest metropolitan areas. It should be noted that although the estimated coefficients in the hottest areas are significant and negative for the second (20-30) and third (30-40) bins, they do not experience any days in the first bin (<20). The coldest areas, on the other hand, do experience days in the ninth bin (>90). In general the hottest cities have a flatter response curve than the coldest cities, supporting the notion that there are heterogeneous temperature preferences in different climate regions.

To further support the notion that different regions of the country have heterogeneous preferences for outdoor activities, Figure 1.6 plots the daily number of minutes spent outdoors each month by state. The graph displays the unconditional survey-weighted average amount of non-work time spent outside. Although the share of time outside is generally quite low, states in Mountain West, North West, and North East increase their shares substantially in the Spring and Summer months relative to the Winter. Surprisingly, the share levels in the South East and South West are quite low and do not vary much throughout the year. The pattern suggests that there is large geographic and temporal heterogeneity in preferences for outdoor time use in the United States. These seasonal and geographic patterns emphasize the need to estimate temperature response functions separately across states, rather than assuming average effects are uniform throughout the country.

1.2 Empirical Strategy

The previous section demonstrated that weather influences the decision to spend time outside. In order to establish the relationship between weather, climate, and leisure time, I use data from

the American Time Use Survey and remote sensing organizations. Following Auffhammer [2018], I detail the Climate Adaptation Response Estimation (CARE) method. Doing so allows me to estimate both the short-run behavioral response to temperature and the long-run response to local climate, allowing for geographic heterogeneity across the United States.

1.2.1 Estimation Strategy

The CARE method estimates both a short-run temperature response function and a long-run climate adaptation effect. The two step procedure is estimated separately for each metropolitan area in the United States, which allows for the current geographic heterogeneity in climate preferences.

In the first stage I measure the temperature response function for each CBSA in the ATUS sample. The model is represented by:

$$y_{it} = \sum_{b \in B \setminus \{60-70\}} \beta_{jb} D_{ibt} + \gamma X_{it} + \alpha_m + \phi_y + \epsilon_{it} \quad (1.2)$$

where y_{it} is the number of minutes spent by individual i doing some activity on day t . The explanatory variables of interest are represented by D_{ibt} , which are dummy variables indicating whether the average temperature experienced by individual i on day t falls in temperature bin b . In addition X_{it} is a vector of observed confounders, including individual characteristics from the CPS and other non-temperature meteorological variables such as the relative humidity and precipitation. α_m and ϕ_y are month and year fixed effects, respectively.

Potential non-linearities in the relationship between temperature and the time spent per day in various activities are captured by the temperature bins b (Schlenker and Roberts [2009]; Deschênes and Greenstone [2011]). As in Auffhammer [2018], the bins are split by decile with exceptions at the top and bottom of the temperature distribution, where they are further divided to create bins for the first, fifth, ninety-fifth, and ninety-ninth percentiles. In total, there are fourteen bins created. Each of the estimated β_{jb} coefficients should be interpreted as the temperature response function relative to the baseline category of the eighth temperature bin. That is, the eighth temperature bin

(from 60 to 70F) serves as a baseline and the deviation measured by the β_{jb} coefficients is not the effect of a day in bin b on y_{it} , but the effect of replacing a day in temperature bin eight with a day in temperature bin b .

One of the primary strengths of this research approach is that the binned temperature coefficients are estimated separately for each CBSA j . The exclusion of location fixed effects — a typical feature in the climate literature — reflects the estimation strategy. By recovering a non-linear temperature response function for each CBSA, CBSA fixed effects are implicitly included. It is also common to interact location fixed effects with time fixed effects to adjust for seasonal characteristics unique to a particular location. Again the nature of the CARE estimation procedure makes the location-time interaction implicit with month and year fixed effects. Still, the exclusion of individual-level fixed effects could create biased estimates if there are unobservable characteristics that influence time-use patterns and are correlated within a geographic location. To account for this, a large number of individual covariates from the CPS are included.

Another key difference between the CARE method and other estimation methods in the climate literature is the inclusion of a second stage. The estimated $\widehat{\beta}_{jb}$ coefficients from the first stage are used as the dependent variable to estimate the sensitivity of CBSA j 's temperature response function to the long-run climate. The second stage model can be represented by:

$$\widehat{\beta}_{jb} = \omega_0 + \omega_1 C_{jb} + \omega_2 Z_j + \eta_{jb} \quad (1.3)$$

where $\widehat{\beta}_{jb}$ are the estimated coefficients from the first stage. C_{jb} is the share of days in CBSA j that occur in temperature bin b from 1979 to 2004. This variable reflects the long-run climate in CBSA j . Z_j is a vector of confounders in a metropolitan area which could influence the long-run response to climate, including income and population density. Adjusting for these variables is important as they could reflect adaptation mechanisms taken through geographic sorting. For example, if cold weather adaptation is more expensive than warm weather adaptation due to heating costs and the costs of heavier winter clothing, then failing to account for income should bias

estimated coefficients downward.

The coefficients of interest are the C_{jb} terms. To clarify interpretation, consider the unconditional number of minutes spent participating in outdoor activities in Table 1. The mean number of minutes outdoors is 32.5 across the sample. In the first stage, if the time response to a day in the 80-90 degree bin relative to a day in the 60-70 degree bin is a decrease of ten minutes, then this would reflect a thirty-three percent decrease in the amount of time spent outdoors on average. Extending the interpretation to the second stage, if the share of days a city experiences above 80 degrees each year increases by ten percent, an estimated coefficient of positive one would be interpreted as an increase in the slope of the first stage coefficient by one minute. More concisely: the hotter a climate is (or becomes), the less sensitive it becomes to hot temperatures. This interpretation can be thought of as the difference under future climates within a city and as the difference between cities with different climates currently.

The purpose of the second stage is to capture geographic differences in the temperature response function that are due to climate. This can be thought of as a cross-sectional approach, in the tradition of Mendelsohn et al. [1994]. The preferences of metropolitan areas for various temperature-sensitive activities reflects not just the temperature on that day, but also the long-run climate as a result of geographic sorting. Deschênes and Greenstone [2011] and Aroonruengsawat and Auffhammer [2011] use a similar approach, but observations of the estimated first stage parameters were at the Census division level, so the results from the second stage were imprecise with only nine observations per temperature bin. The first stage in this research is estimated at the metropolitan level, which provides for many more observations per temperature bin.

1.3 Results: How Weather and Climate Explain Leisure Behavior

Section 3.4 implements the Climate Adaptation Response Estimation procedure. The first stage recovers the temperature response function for each metropolitan area in the sample. The result-

ing parameters are then used to estimate how much of a city’s long-term climate influences its sensitivity to temperature in the second stage. Results suggest that cold-weather areas are more temperature sensitive and therefore would be expected to respond (or adapt) most aggressively to climate change.

1.3.1 *First Stage*

The first stage estimates the temperature sensitivity of household activities separately for each metropolitan area in the American Time Use Survey dataset. The results support heterogeneous temperature preferences across the country.

Figure 1.7 plots the estimated coefficients following Model 1.2, with the number of minutes spent outside per day as the dependent variable. In total there are 31 metropolitan areas in the sample. The orange line represents the median response among the metropolitan areas, with each successive ribbon of blue representing a 20% change in percentile. Since the reference bin is from 60-70, the coefficients can be interpreted as the change in the number of minutes spent outside from the daily average temperature falling in bin b , relative to a day with average temperature between 60-70. The curve has a vague inverse-U shape, but there is significant heterogeneity between the cities in the sample. Estimates in the bottom half of the temperature distribution are more consistently negative than those in the top of the temperature distribution. Response heterogeneity is reflective of the different temperature response functions in Figure 1.5, where the hottest twenty percent of metropolitan areas have much flatter estimated coefficients than those in the coldest twenty percent. Figure 1.8 plots the temperature response curve for each city separately. The cities are arranged to closely resemble their relative geographic position, so New York City is in the top (north) right (east) and San Diego is in the bottom (south) left (west).

To make the comparison more clear, Figure 1.9 represents a placebo exercise which estimates the same equation separately for each metropolitan area, but now the dependent variable is the number of minutes spent shopping per day. Shopping was arbitrarily chosen as it was one of the activities in Table 1 which demonstrated no relationship to temperature. Relative to the results in Figure 1.7, the estimates of β_{jb} are flat and close to zero across the temperature distribution. The

effects, when compared to those in Figure 1.7, are much smaller in magnitude and do not exhibit the same inverse-U shape. These results support the idea that there are certain activities which are uniquely vulnerable to the weather and to climate change, which imply long-run welfare effects since the margin for adaptation is limited.

1.3.2 *Second Stage*

In this section, estimated coefficients from the first stage are used as the dependent variable to recover the effect of local climate on temperature sensitivity. The cross-sectional approach demonstrates the role of the long-run climate distribution on preferences.

Table 3 presents the results. The coefficients of interest are the “Bottom Bin Shares” and the “Top Bin Shares,” which are the pooled shares of days in the bottom five and top three bins over the twenty-five years prior to the start of the ATUS data (1979-2004). The model is estimated at the metropolitan level. A pooled model is estimated because not every metropolitan area experiences days in the most extreme bins, so results are more stable than estimating the climate response on the share of each bin separately. That being said, columns (1) and (2) regress the estimated coefficients for the bottom five bins in stage one on the share of days in those bins, and columns (3) and (4) do the same for the subset of coefficients for the top five bins. Since the slope of the estimated coefficients from stage one can differ dramatically at the extreme bins, columns (2) and (4) include interaction terms for a bin estimate being in the bottom and top three bins.

The coefficients can be interpreted as the change in the slope of the first stage coefficients from increasing the share of days in the bottom/top bins by ten percent. In plain terms this means that the coefficients represent the change in the temperature sensitivity of outdoor time due to the relative coldness/hotness of the local climate. A positive (negative) slope in the first two columns would indicate that a higher share of cold days would induce more (less) time outside when the weather is cold.

The results in columns (1) and (2) suggest that a ten percent increase in the share of cold days makes the slope of the temperature response function *more* negative. For reference, a two minute

decrease in the slope of the first stage temperature response curve accounts for approximately ten percent of the total temperature response. That is, as the share of cold days increases in a city, the less time people spend outdoors on cold relative to pleasant days. Although these results appear counter-intuitive at first glance, they reflect the disparity in adaptation measures already taken with respect to cold weather across the United States. For example, Minneapolis, Minnesota has an eight-mile system of enclosed pedestrian footbridges which connect buildings in the city's downtown, making navigation of the harsh winter more manageable. Avoidance behavior is itself adaptation, and so it should not be a surprise that those most experienced with cold weather are those most able to avoid it. The same experience cannot be said of hot cities with respect to hot weather, however; results for the top bins in columns (3) and (4) are statistically insignificant with large standard errors.

To demonstrate the effect more clearly, Figure 1.10 plots the predicted second stage values for two representative cities: one cold city with the share of cold days equal to that of the tenth percentile in cold day shares, and one hot city with the share of hot days equal to that of the ninetieth percentile in hot day shares. The cold city is predicted to spend much less time outdoors as the share of cold days increases but not change when the share of warm days increases, all else equal. On the other hand, the hot city spends slightly less time outdoors with a higher share of cold days but more time outdoors with a higher share of hot days. The hot city-hot days relationship is less than half that of the cold city-cold days relationship, suggesting that there is more room for adaptation in cold weather than there is in hot weather.

1.3.3 Robustness Checks

If the first stage results are biased I would expect to see a knock-on effect in the second stage, so in this section I run a number of robustness checks to establish that the first stage results are not biased. First, in order to ensure that the shape of the non-linear temperature response function is not being driven by intertemporal substitution, Figure .0A1 presents the same model as in Equation 1.2 with lags for the mean daily temperature in city j over the three days prior to the respondents' diary

days. If respondents were engaging in intertemporal substitution of outdoor activities, we would expect the temperature response function to be flatter across the eight estimated temperature bin than it is in Figure 1.7 where no temperature lags are included. However, what is observed in Figure .0A1 is consistent with previous estimates, exhibiting large negative effects for temperatures below the reference bin and no or little effect in the top temperature bins.

In the main specification of the CARE model, month and month-year fixed effects are used to adjust for seasonal and contemporaneous unobserved confounders. The model is also estimated separately for each metropolitan area in the sample, so location fixed effects are implicit. Although the location “fixed effects” cannot be omitted due to the model specification, the month and month-year fixed effects can be. If the results without time fixed effects maintain their shape (i.e. large negative responses below the 60-70 reference bin and small flat responses above the reference bin), a conclusion can be drawn: the sample period of 2005-2018 does not contain enough temporal variation to alter the results. This suggests that — apart from the seasonality implied by the temperature bins — the time horizon is not long enough to capture any change in the climate. The first stage estimates the short-run temperature response function, not the climate effect, so omitting time fixed effects should have no impact on results. Figure .0A2 plots the first stage model with no time fixed effects. As expected, the temperature response curve maintains its shape.

1.4 Outdoor Leisure in a Changing Climate

The practical implications of the previous estimates depends on the climate changing over the course of the coming decades. Shifting the climate distribution to the right will mean more warm days and fewer cold days. Using a suite of General Circulation Models (GCMs) from the Intergovernmental Panel on Climate Change, this section explores how a warming world could change non-market behavior in the United States and its welfare implications.

1.4.1 Forecasting the Effect of Climate Change on Outdoor Time

Estimates from Section 3.4 are used to project the effect of end-of-century warming on the amount of time spent outdoors. There are two competing effects happening. First, there will be more extremely hot days which – absent of adaptation – should decrease outdoor time. On the other hand, an increase in average temperatures could induce more time outside in the winter as the cost and prevalence of cold weather decreases. The change in the time outside can be represented by:

$$\Delta y_j = \frac{y_{jt+inf}}{y_{jt}} = \frac{\sum_j \widehat{\beta}_{jb} \times \Delta \text{Mean Temp}_j}{\sum_j \text{Mean Time Outdoors}_j} \quad (1.4)$$

where Δy_j is the percentage change in the amount of time spent outside in city j between the current period and the RCP4.5 and RCP8.5 GCMs. $\widehat{\beta}_{jb}$ are the estimated coefficients from Stage 1 of the CARE method from Equation 1.2 which are multiplied by the change in the number of days in bin b for each city j under the two GCM scenarios. This estimate of the change in the amount of time spent outdoors is divided by the average number of minutes spent outdoors in city j during the ATUS sample period (2005-2018). The resulting quotient is the percentage change in the number of minutes spent outside each year for city j .

Daily downscaled projections of the future climate come from NASA’s Earth Exchange Global Daily Downscaled Projections (NEX-GDDP) dataset. The data are at a $2.5\text{km} \times 2.5\text{km}$ grid. NASA publishes 21 climate models from international research groups, each of which provide estimates for the RCP4.5 and RCP 8.5 scenarios. The models are weighted following the methods used in the Fourth National Climate Assessment (Sanderson and Wehner [2017]). The RCP4.5 scenario assumes that greenhouse gas emissions (CO_2 , CH_4 , N_2O , etc.) plateau in the early 2040s and decline in the subsequent years. This is in contrast to the RCP8.5 scenario, which assumes a “business as usual” approach to greenhouse gasses. Emissions continue to rise throughout the century under the RCP8.5 scenario. As a result, the RCP8.5 scenario projects upwards of four degrees Celsius of warming, whereas RCP4.5 results in between two and three degrees of warming.

I will focus on the RCP4.5 scenario for two reasons. First, the 2014 IPCC Fifth Assessment

Report (AR5) which established the RCP system, presented the RCP8.5 scenario as an unlikely worst-case scenario with “with low income, high population and high energy demand due to only modest improvements in energy intensity” (Riahi et al. [2011]). The “business as usual” moniker then is somewhat misleading, since greenhouse gas-abatement technologies such as solar energy become increasingly cost effective. It is unlikely that the current energy status quo persists to the end of the century. Second, uncertainty of natural processes which create deviations in climate models are especially pronounced in RCP8.5 models, where warming feedback loops compound uncertainty. Cloud dynamics, for example, have proven especially difficult to project (Meehl et al. [2020]). For the purpose of this research then, the primary results will be those of the RCP4.5 scenario.

Figure 1.11 presents the results of the simulation. The cities featured are those included in the first stage of the CARE method in Section 1.2.1. Each city has two bars: one for the RCP4.5 scenario and one for the RCP8.5. The change in the amount of time outside is displayed in green if it is increasing and in orange if it is decreasing. Broadly, cities in the top (northern) half of the graph experience increases in their outdoor time, whereas cities in the bottom (southern) half see decreases. The results are reflective of the conclusions from Stage 2 of the CARE method which suggested that cold weather cities spent more time indoors when the weather was cold. In general, it appears that there are margins for increasing outdoor time in northern cities due to warmer winters, but that margin doesn’t exist in southern cities where winters are already warm enough to enable outdoor activities. Instead, these southern cities will likely lose outdoor time as their hot summers become more extreme.

1.4.2 Back of the Envelope Welfare Calculations

Finally, to provide a rough approximation of the welfare consequences of climate change as a result of shifting time use patterns in the United States, I return to the U-Index measure developed by Kahneman and Krueger [2006]. Figure 1.12 plots the change in the U-Index due to the change in the present climate to the predicted end of century climate in cities with at least 1000 ATUS diary

day observations. The change is calculated by multiplying the outdoor-indoor U-Index ratio by the expected change in time outdoors for each city, as presented in Figure 1.11. An increase in the U-Index ratio implies that the share of time spent doing activities outdoors — and therefore time spent doing activities that are more pleasant — increases. A five percent increase in the U-Index ratio means a five percent increase in the amount of time spent doing outdoor activities, which in general are more pleasant than indoor activities. The results of this simple exercise reflect the results of the CARE model: many of the cities that stand to benefit the most are cold-weather cities which will see a decrease in the number of cold days, allowing for more time outdoors. In contrast, the cities most negatively affected are largely warm-weather cities with mild winters already but hot summers, leading to less outdoor time in the aggregate.

1.5 Conclusion

This paper illustrates the importance of accounting for non-market adaptations to climate change. There are two implications of the work. First, the role of behavioral change is considered as a margin for adaptation. Frequently the climate change literature in economics focuses on market activities and investment as the means of adapting, but time use surveys provide a non-market vector of change. Additionally, the surveys are both familiar to economists and will continue to update as the climate changes. Second, there are behavioral preferences for climate that vary within the United States. The breakdown of time spent indoor versus outdoors depends to a great extent on where in the country a person lives. There are clear non-linearities in the temperature response function which are most pronounced in the coldest regions of the country. These areas are expected to warm the most in the coming century, leading to large welfare changes as winters become more mild and summers more extreme. These heterogeneous climate preferences reflect both long-term adaptation and geographic sorting.

One of the benefits of extending the climate literature to non-market activities using time use surveys is that similar surveys are carried out in a number of countries, including in Europe and

East Asia. Further research estimating non-market adaptation in other national contexts could provide new and potentially more generalizable results, especially in the developing world. This research includes strong assumptions about the ability to adapt to temperatures at the top of the climate distribution which have not yet been realized in the United States, so there may be countries or regions with weather observations that are outside the sample used here that might preview what is to come.

Authors' Notes

Opinions expressed in this essay are solely our own responsibility and do not reflect the view of any agency and any errors are ours.

Acknowledgments

Special thanks to Kevin Kuruc, Daniel Hicks, Jonathan McFadden, and Pallab Ghosh at the University of Oklahoma for their support and advice in improving this work. Also thanks to the participants of the OU Economics Department's Friday Seminar Wednesday Brown Bag Seminar Series for their helpful comments. .

Declaration of Conflicting Interests

The author declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author received no financial support for the research, authorship, and/or publication of this article.

Notes

¹<https://www.washingtonpost.com/weather/2020/04/21/earth-warmest-year-likely-2020/>

²Many articles reviewing outcomes exist. For agriculture see Auffhammer and Schlenker [2014], for energy see Auffhammer and Mansur [2014], for conflict see Burke et al. [2015], and for a general overview of climate econometrics, see Hsiang [2016]. Finally, Massetti and Mendelsohn [2018] reviews the adaptation literature.

³For a meta-analysis of public health research, see Bowler et al. [2010]. For a review of the psychology literature, see Pearson and Craig [2014]. And for a meta-analysis of the role greenspace plays in health outcomes, see Twohig-Bennett and Jones [2018].

⁴This conclusion is intuitive. As Senator Bernie Sanders, a Vermont native, put it in an August 2020 interview: “When I’m in Washington, I don’t go outside, and when I’m in Vermont, I don’t go inside. So there you go.”

⁵Time fixed effects impose the same assumption but over time. This is less of a concern, because this style of models typically uses data that spans years, not decades, so the potential for short-term adaptation is limited.

⁶The categorization is inherently subjective, something that many time use research projects have contended with. Aguiar and Hurst [2007] and Ramey and Francis [2009] attempt to create a framework for time use allocation in the context of shifts in Americans’ labor-leisure intensive margin adjustments, but come to different conclusions due to the subjectivity of the categorization process. Krueger and Mueller [2012] creates two methods based on the respondents’ emotional state while carrying out the activities. Since this research does not study time-use trends, the categorization here is strictly for presentation purposes and does not have bearing on the conclusions reached.

⁷Barreca et al. [2016] and Heutel et al. [2017] similarly estimate the temperature-mortality relationship for different time periods and different percentiles of the temperature distribution respectively.

Tables

Table 1: Summary Statistics for Time Use Activities

Number of Minutes per Day

Activity	Unconditional			Conditional on Minutes > 0		
	Mean	Median	Share Zeros	Mean	Median	N. Obs.
Household Labor						
Personal Care	47.1	35	19.8%	58.7	45	141,331
Doing Housework	115.5	65	21.5%	147.2	105	138,209
Caring for HH Member	33.9	0	71.1%	117.1	80	50,986
Caring for non-HH Member	8.2	0	88.2%	69.3	20	20,721
Traveling/Commuting	75.3	60	15.1%	88.6	70	149,569
Market Labor						
Working	157.6	0	62.1%	415.6	460	66,778
Working Outside	0.7	0	99.7%	230.0	120	530
Market Consumption Activities						
Shopping	25.9	0	57.3%	60.6	40	75,159
Using Services	5.0	0	92.3%	65.6	45	13,519
Using HH Services	0.9	0	97.9%	43.3	20	3,709
Sports	19.3	0	80.9%	101.5	60	33,568
Leisure Activities						
Eating/Drinking	68.4	60	4.2%	71.4	60	168,723
Leisure	295.0	255	4.7%	309.7	270	167,802
Using the Telephone	7.7	0	84.0%	47.9	30	28,205
Exercising	17.5	0	81.7%	95.9	60	32,164
Outdoor Activities						
Outdoors	32.5	0	62.5%	86.6	50	66,109
Outdoor Dining	0.3	0	99.3%	39.7	30	1,202
Outdoor Exercise	5.0	0	94.5%	90.4	60	9,686
Outdoor Non-work	31.8	0	62.6%	84.9	50	65,938
Miscellaneous Activities						
Sleeping	529.1	520	0.1%	529.6	520	175,956
Education-Related Activities	16.3	0	94.0%	272.5	240	10,516
Religious Activities	12.9	0	88.0%	107.2	85	21,152
Volunteering Activities	8.9	0	93.5%	136.1	100	11,500

Table 2: Change in the Number of Minutes Spent by Activity
Relative to a Day with Max Temperature from 60-70°

Activity	Temperature Bins								R^2
	≤20°	20-30°	30-40°	40-50°	50-60°	70-80°	80-90°	≥90°	
Household Labor									
Caring for HH Member	-4.3 (2.6)	-1.2 (2)	-1.8 (1.6)	-0.4 (1.1)	1.1 (1)	-0.4 (1.1)	-0.5 (1.8)	0.3 (4.6)	0.235
Doing Housework	-2.3 (4.1)	-4.3 (3.6)	-3.8 (2.6)	-3.8 (2.2)	-1.3 (1.9)	0.2 (1.8)	-0.8 (2.9)	15.1 (8.4)	0.162
Caring for Non-HH Member	-1.4 (1.2)	0.2 (1)	0.2 (0.9)	0.1 (0.6)	0 (0.6)	0.1 (0.7)	1.3 (1)	-1.2 (2.5)	0.072
Personal Care	0.1 (1.7)	0.5 (1.2)	0 (1)	-0.4 (0.9)	-0.8 (0.7)	-0.8 (0.8)	0.2 (1.2)	-3.2 (4.4)	0.091
Traveling/Commuting	-1.2 (2.8)	-3.5 (2.7)	-3.2 (1.8)	-3.5 (1.3)	-0.9 (1.1)	0.8 (1.1)	0 (2)	-9.9 (4.6)	0.071
Market Labor									
Working	-13.4 (9.1)	-2.9 (5.4)	-2.6 (5.2)	-0.1 (4.1)	2.1 (3.5)	0.8 (4)	-3.1 (7.8)	-9.4 (16.7)	0.380
Working Outside	0 (0.6)	0 (0.5)	0.4 (0.4)	0.3 (0.3)	0.5 (0.2)	0 (0.3)	0.2 (0.5)	2.6 (1.9)	0.054
Market Consumption Activities									
Using Government Services	-0.2 (0.2)	0 (0.2)	0.2 (0.2)	0 (0.1)	0 (0.1)	0.1 (0.2)	0 (0.2)	-0.1 (0.3)	0.048
Using HH Services	0 (0.2)	0 (0.2)	0.1 (0.2)	0.1 (0.2)	0 (0.1)	0 (0.2)	0.4 (0.3)	-0.2 (0.6)	0.040
Using Services	-0.5 (0.7)	0 (0.6)	-0.1 (0.5)	0.2 (0.3)	0.5 (0.3)	-0.1 (0.4)	-0.3 (0.6)	-0.1 (2)	0.059
Shopping	0.1 (1.7)	-0.2 (1.2)	0.1 (1)	0.5 (0.9)	0.4 (0.8)	0.1 (0.9)	-0.4 (1.2)	5.6 (5.2)	0.107
Sports	-4.4 (1.7)	-4.2 (1.2)	-4.1 (1)	-2.8 (1)	-1.1 (0.8)	1.3 (1)	3.7 (1.7)	-1.2 (4.9)	0.080
Leisure Activities									
Exercising	-4.3 (1.7)	-3.5 (1.2)	-3.5 (1)	-3 (0.9)	-1.6 (0.8)	1.4 (1)	3.8 (1.5)	-1.3 (4.5)	0.079
Eating/Drinking	-0.3 (1.8)	-1.3 (1.3)	0.4 (1.2)	-0.7 (0.8)	0.1 (0.8)	0 (0.8)	2 (1.1)	5 (4)	0.085
Leisure	21.8 (5.7)	15.1 (4.5)	9.8 (3.8)	6.9 (3.1)	0.6 (2.3)	2 (3.2)	4.3 (4.7)	11.6 (13)	0.195
Using the Telephone	0.7 (0.7)	0.3 (0.6)	0.4 (0.4)	-0.2 (0.3)	-0.4 (0.3)	0 (0.3)	0 (0.5)	2.4 (1.4)	0.071
Outdoor Activities									
Outdoor Dining	0 (0.1)	0 (0.1)	-0.1 (0.1)	0 (0.1)	0 (0.1)	0 (0.1)	0 (0.1)	-0.2 (0.2)	0.041
Outdoor Exercise	-1.7 (0.9)	-1.3 (0.8)	-2 (0.6)	-1.2 (0.6)	-0.8 (0.5)	0 (0.4)	0.3 (0.7)	-1.1 (1.2)	0.069
Outdoor Leisure	-0.5 (0.5)	-0.7 (0.5)	-0.5 (0.4)	-0.6 (0.4)	-0.3 (0.4)	0.5 (0.4)	0.4 (0.6)	0.4 (1)	0.058
Outdoors	-10.2 (2)	-11.1 (2)	-11 (1.7)	-7.5 (1.5)	-2.6 (1.3)	-0.1 (1.3)	-0.2 (2)	-2.2 (4)	0.111

Outdoor Non-work	-10.1 (2.1)	-11.2 (2.1)	-11.5 (1.7)	-7.8 (1.5)	-3.2 (1.2)	-0.1 (1.3)	-0.5 (1.8)	-4.9 (2.9)	0.115
Miscellaneous Activities									
Education	0.4 (2.4)	-1.1 (1.8)	-1.2 (1.8)	-0.4 (1.3)	-1.2 (0.9)	-1.4 (0.8)	-2.5 (1.1)	1.9 (5.9)	0.136
Religious Activities	0.7 (1.3)	-0.2 (1)	0.3 (0.8)	0.4 (0.7)	0.5 (0.7)	0 (0.6)	0 (1)	-1.9 (2.5)	0.134
Sleeping	5.1 (5)	2.9 (3.9)	5.7 (2.7)	4.8 (2.3)	0 (2)	-2.7 (1.7)	-2.9 (2.9)	-9.8 (9.7)	0.167
Volunteering	-1.2 (1.3)	0.1 (1.1)	-0.8 (0.9)	0 (0.7)	-0.2 (0.6)	-0.3 (0.7)	-0.8 (1.1)	-1.9 (2.6)	0.064

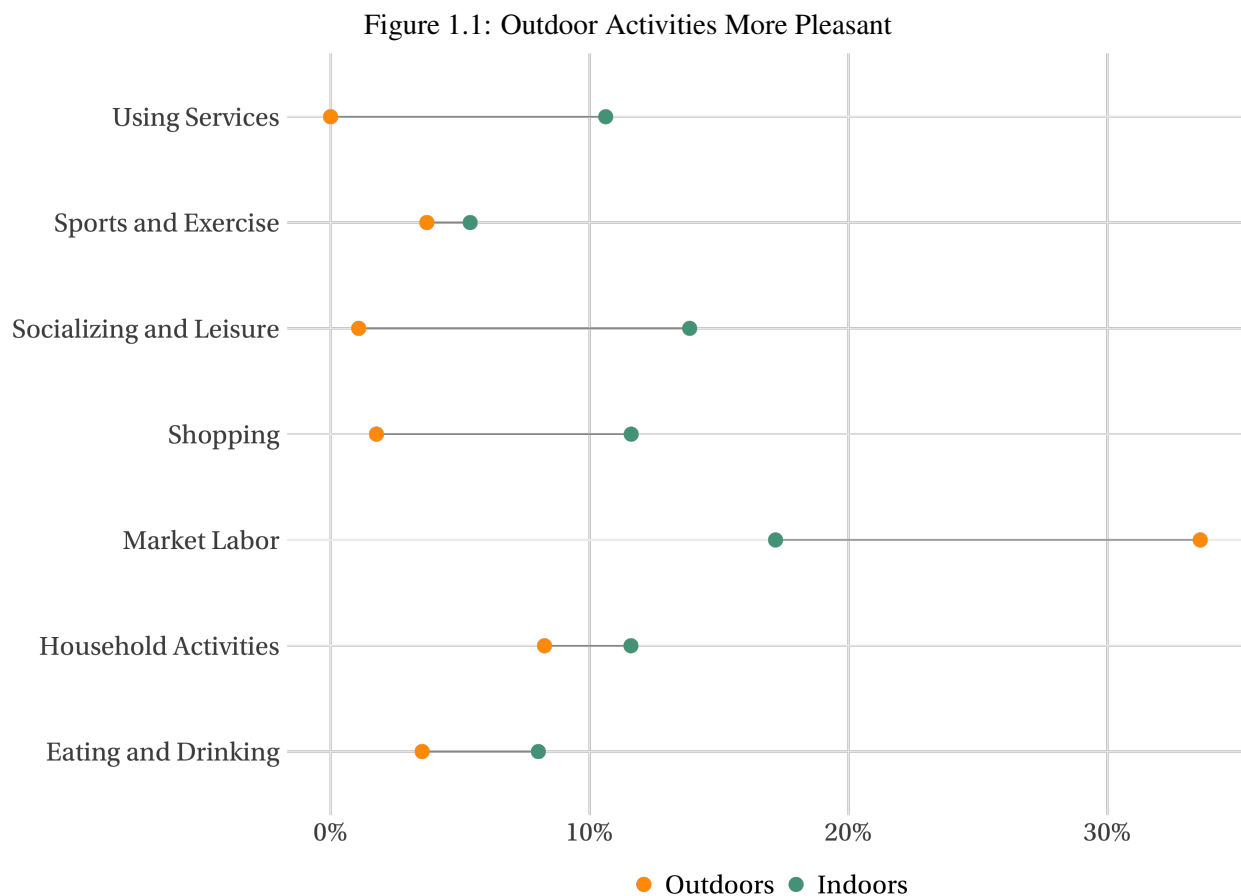
Note: Each activity is estimated separately as a dependent variable with city-month and month-year fixed effects. Standard errors in parentheses are clustered at the state level.

Table 1.3: CARE Stage 2

	Bottom Bins		Top Bins	
	No Interactions	Interactions	No Interactions	Interactions
Bottom Bin Shares	-2.47*** (.913)	-1.94** (.816)		
Top Bin Shares			2.63* (1.35)	3.89** (1.77)
Num.Obs.	160	160	79	79
R2	0.101	0.104	0.092	0.104
Bin FE	Yes	Yes	Yes	Yes
SE	White	White	White	White


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

Figures




Note: The U-Index (Kahneman and Krueger [2006]) calculates the share of time an activity is reported as being unpleasant by survey respondents. If the primary emotion during the activity is negative (e.g. if the respondent reports being more unhappy than happy), then the activity is coded as unpleasant for that respondent. The index is the weighted average of all such responses from the American Time Use Survey's Wellbeing module. The graph demonstrates that activities which take place inside are more unpleasant than the same activities when they take place outside.

Figure 1.2: Simulation of the Temperature Distribution Shifting 3 Celsius




Note: This simulation reflects the change in the temperature distribution of the United States given a three degree Celsius increase in daily average temperatures. The numbers are simplified for demonstration purposes; different areas of the United States are expected to warm at different rates. Plotting the average national temperature distribution ignores the heterogeneous response to climate change across the country.

Figure 1.3: Annual Number of Days with Average Temperature in 5 F Bins Under Three Climate Scenarios



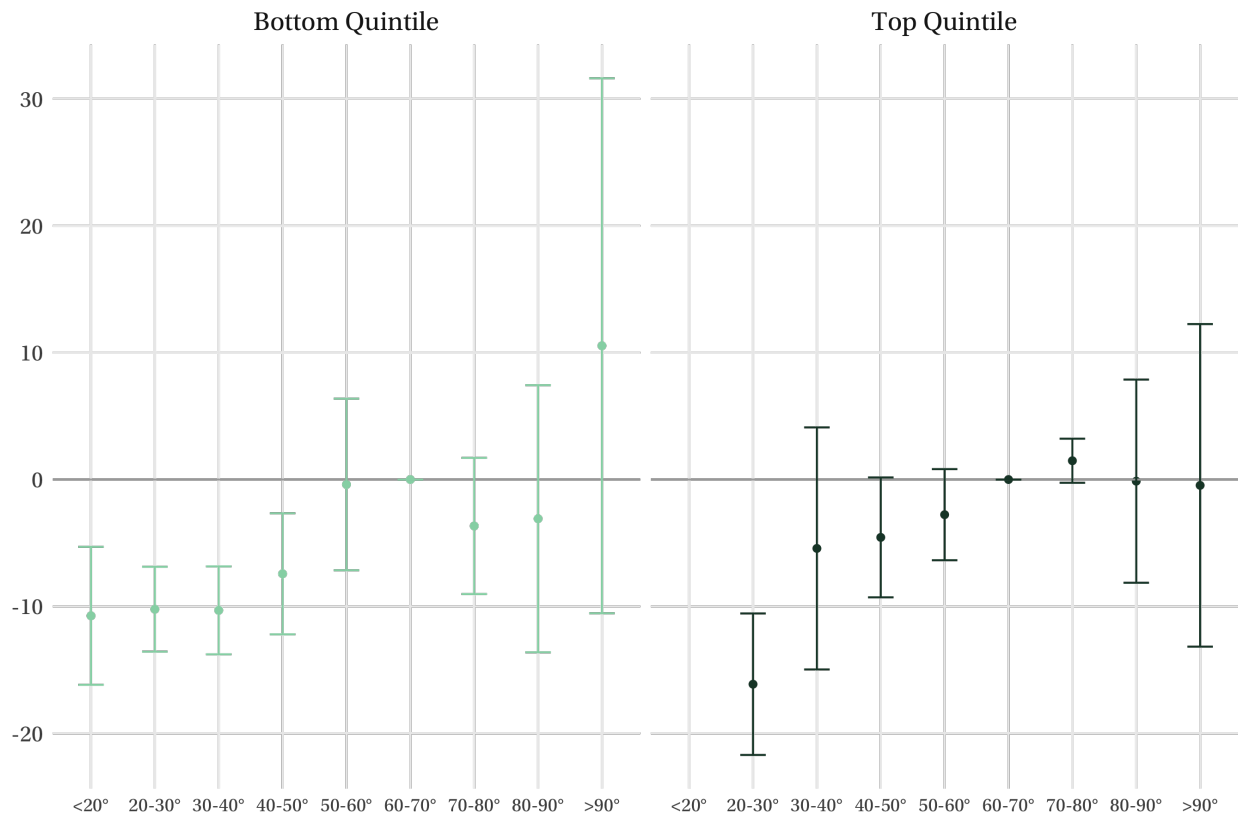
Note: The figure presents the average temperature under three climate scenarios in the sample locations. First in light green is the current climate, defined as the period from 2005 through 2018. The second is the Intergovernmental Panel on Climate Change’s (IPCC) Representative Concentration Pathway (RCP) 4.5 in which greenhouse gas emissions peak between 2040 and 2045 and then begin to decline. This intermediate scenario implies global temperatures rise by 2-3 Celsius. Finally, the darkest shade represents RCP 8.5, under which emissions continue to rise unabated and global temperatures increase by more than 4 Celsius.

Figure 1.4: The Share of Time Spent Outside by Activity Group



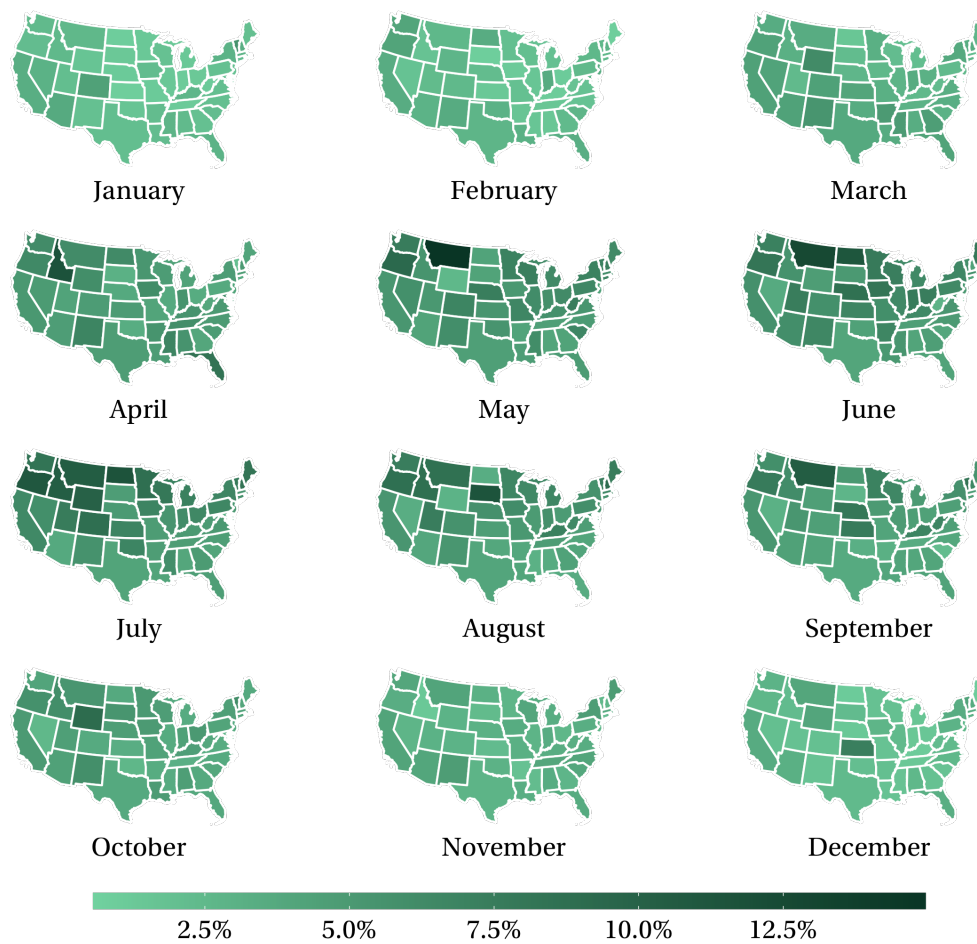
Note: Household activities — including gardening, yard maintenance, and pet care — and participation in sports and recreation take place outside approximately one quarter of the time. Other activity groups have much smaller shares of time spent outside.

Figure 1.5: Estimated Change in the Number of Non-Work Minutes Spent Outside in Cities in the Bottom and Top Quintiles of the Temperature Distribution




Note: The temperature response function is estimated for the bottom and top quintiles of the average temperature distribution. The coldest metropolitan areas appear to be more consistently sensitive to temperature for the coldest bins and are clearly more sensitive than the top quintile at extremely hot temperatures.

Figure 1.6: Daily Share of Non-Work Time Spent Outside



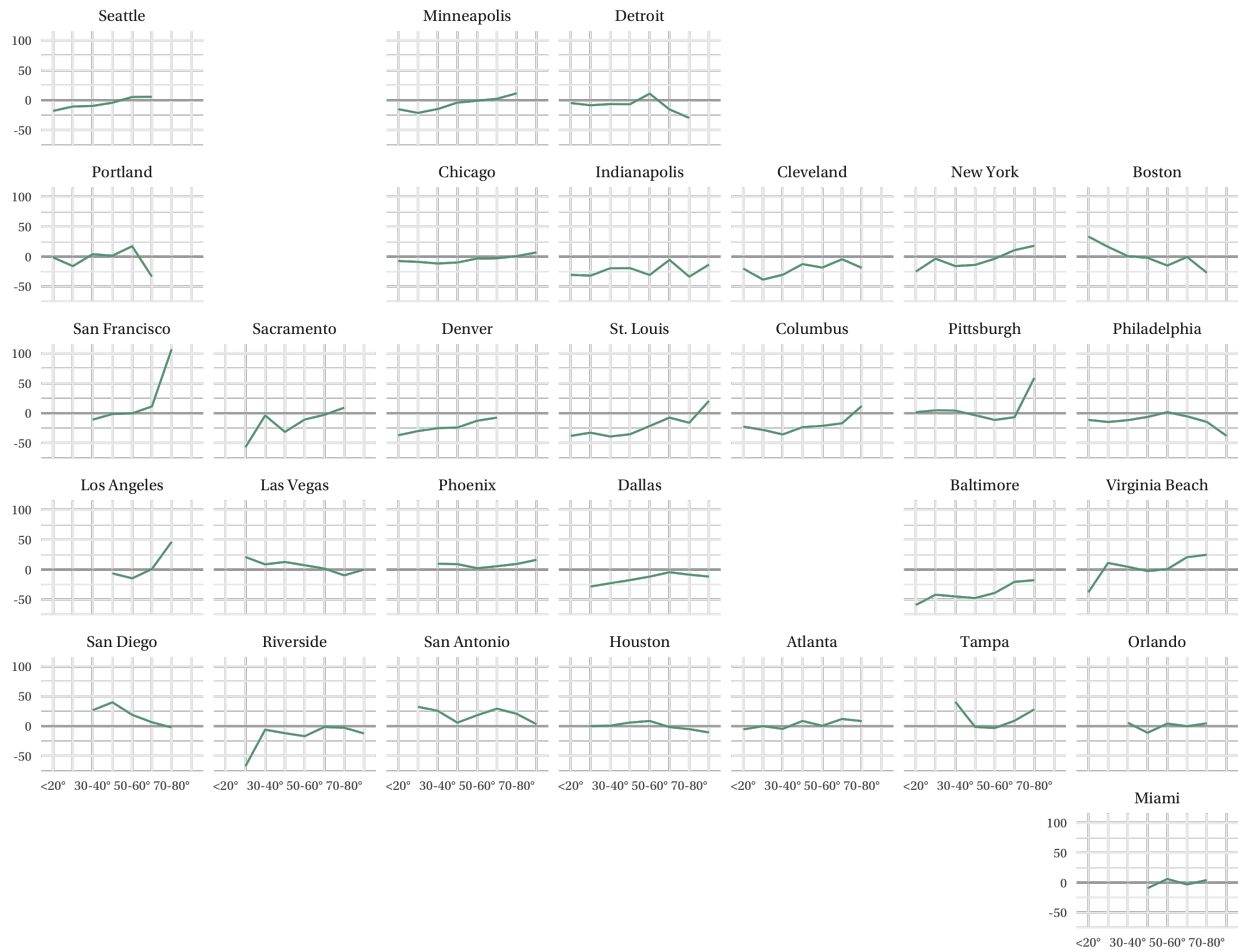
Note: The graph above calculates the daily share of non-work time spent outside by dividing the total number of minutes reported outside by the number of non-work and non-sleeping minutes. Large heterogeneous seasonal patterns across states emphasize the importance of not assuming constant temperature reaction functions.

Figure 1.7: Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time




Note: The solid orange line represents the median response, and each shaded green area represents the surrounding high and low decile. For example, the darkest shaded area is the estimated coefficients that are in the 40th to 60th percentile of the coefficient distribution. The coefficients are interpreted as the change in time outside as a result of being in a particular temperature bin relative to a day in the 60-70bin.

Figure 1.8: Stage 1 Temperature Response Function Across Cities



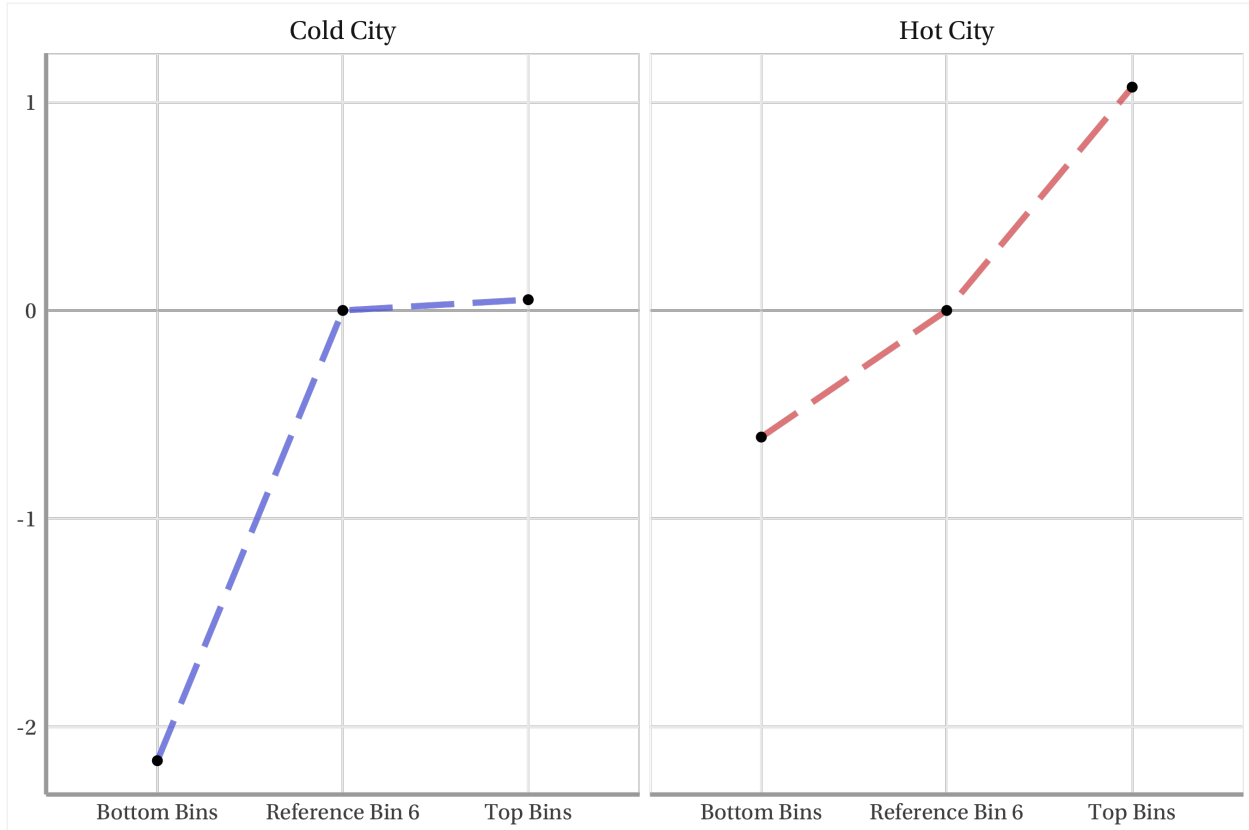
Note: This figure displays the results for each city with at least 1000 ATUS respondents as estimated in Stage 1 of the CARE method. The results in Figure 1.7 are an aggregation of the city-level results presented here. The cities are arranged approximately geographically.

Figure 1.9: Placebo Exercise: No Pattern for Shopping




Note: The specification for this figure is the same as in Figure 1.7, but with shopping as the dependent variable. The median response to temperature on shopping is much flatter than the response on time spent outdoors, supporting the conclusions from Section 1.2 that there are some settings that are more temperature sensitive than others.

Figure 1.10: Predicted Second Stage Response in a Synthetic Cold and Hot City




Note: This figure presents the percentage change in the number of minutes spent outside for the in-sample cities under the RCP4.5 and RCP8.5 climate scenarios published by the IPCC. Similar to Figure 1.8, the cities are arranged approximately geographically for visualization purposes. Green indicates an increase in time outside and orange indicates a decrease.

Figure 1.11: Percentage Change in the Amount of Time Spent Outside



Note: This figure presents the percentage change in the number of minutes spent outside for the in-sample cities under the RCP4.5 and RCP8.5 climate scenarios published by the IPCC. Similar to Figure 1.8, the cities are arranged approximately geographically for visualization purposes. Green indicates an increase in time outside and orange indicates a decrease.

Figure 1.12: Change in the U-Index Under the RCP4.5 Scenario



Note: This figure presents the percentage change in the number of minutes spent outside for the in-sample cities under the RCP4.5 and RCP8.5 climate scenarios published by the IPCC. Green indicates an increase in time outside and orange indicates a decrease.

Chapter 2

Weather, Climate, and the Police in America

Three Americans are killed by police every day.⁸ The murder of civilians – like that of George Floyd in 2020 – have propelled social justice movements around the world and forced a reckoning with the use of force by police. The deaths of Floyd, Philando Castile, Eric Garner, and others, along with countless episodes of police brutality more generally, were the result of routine foot and traffic patrol stops. Racial disparities in police interactions with the public are a particular focus of the broader public conversation and research literature (Knox et al. [2020], MacDonald and Fagan [2019], Goncalves and Mello [2020]). Increased focus on police departments comes simultaneously with renewed attention on the potential damages that result from climate change. This research unifies these two emerging narratives, exploring the relationship between police-public interactions in the United States and the climate.

In order to conceptualize the role of weather and climate on police conduct, I argue that there are two channels through which officers are affected. First, the productivity of the police – i.e. how many interactions police have with the public – could be subject to the weather conditions that they experience. Other sectors and industries show productivity being strongly affected by the weather workers experience (Colmer [2019], LoPalo [2020], Kahn and Li [2019]). The second channel is through increased aggression. Similar to the productivity channel, it is well established that people act more aggressively at extreme temperatures (Hennessy and Wiesenthal [1999], Reifman

et al. [1991]). Police aggression is a central theme of the Black Lives Matter movement, and if the weather impacts the police through both their decision making and their behavior, then the future of the police's relationship with the public will be put under more strain as the climate warms, absent reform.

Using data from the Stanford Open Policing Project (Pierson et al. [2020]), I estimate the sensitivity of police conduct to short-run weather and long-run climate among eighty-seven police departments in the United States. I find that the number of encounters police have with the public varies greatly on a week-to-week basis depending on the weather, with both extreme hot and extreme cold temperatures having a pronounced effect. This result reflects the impact of weather and climate on police productivity, as the number of encounters is one of the primary ways cities and police departments evaluate their own performance (Collier [2001], Fielding and Innes [2006]). Having established that police productivity changes with the weather, the next step is to estimate whether the criteria that police require to engage in an encounter with a member of the public becomes more strict with changes in the weather. If the police are simply responding to the behavior of the public, then we would expect the criteria for stops to remain constant. To test this intuition, I estimate the difference between the speed limit and the speed drivers are traveling when they are stopped by the police. Finally, I use the "veil of darkness" approach (Grogger and Ridgeway [2006]) to estimate whether the racial bias exhibited by the police when conducting stops increases in severity on abnormally hot days. These last two exercises – the speed test and the veil of darkness test – demonstrate that weather makes the police more strict and exacerbates their racial biases.

Following work in the climate literature, all the models are estimated non-linearly to allow for the response to temperature vary across the temperature distribution (Schlenker and Roberts [2009], Barreca et al. [2016], Auffhammer [2018]). This allows the response to a day when the temperature is ninety (or twenty) degrees Fahrenheit to be different from the response to a day when it is sixty degrees. In the weekly stops model, police increase the number of stops they make by as much as four percent when the weather is significantly warmer than average, but decrease

the number of stops by the same amount when the weather is significantly colder than average. These results support the notion that police productivity changes when the weather deviates from comfortable temperatures in the 65 to 70 degree Fahrenheit range commonly used as a reference temperature in the climate literature. Despite these pronounced results, the increase (and decrease) in police activity might simply reflect the actions of the public. If, for example, drivers become more aggressive when the weather is hot (Kenrick and MacFarlane [1986]), we could expect the police to respond in kind and conduct more stops.

To test whether the police are responding to public behavior or themselves changing their behavior, the second and third models present evidence that the police become more strict and increase the share of black drivers they stop at hot temperatures. To capture the change in strictness, I estimate a similar non-linear model to that of the weekly stops model, this time at the daily level with the dependent variable representing the difference between the posted speed limit and the speed a driver is traveling when they are stopped. If the police are simply responding to the actions of the public, then there are two possible outcomes for this model. First, there is no change in the dependent variable – i.e. the standards of the police do not change. Second, if the public is speeding faster at extreme temperatures, then the average difference in the dependent variable should increase to reflect the fact that drivers are traveling faster. Despite this expectation, the results of the second model show that police officers actually decrease the speed difference that they stop drivers for. That is, the police become more strict and pull drivers over for smaller infractions at very hot temperatures.

Finally, to test whether the racial biases police exhibit are exacerbated by hot weather, I estimate a veil of darkness model (Grogger and Ridgeway [2006]). the veil of darkness mode tests whether the share of black drivers stopped by police decreases with the onset of night or increases with sunrise. The intuition is that police have a more difficult time deciphering the race of drivers when there is no daylight. Table 2.2 shows a simple cross section, using the variability of darkness throughout the year between 5:30 and 6:00PM to get the proportion of drivers who are black and pulled over by police when it is light and dark out. The share of black drivers stopped decreases

when it gets dark out, suggesting some support for the veil of darkness hypothesis. For a number of reasons though, including the fact that commuting patterns between racial groups may differ (Hamermesh et al. [1996]) this naive comparison is insufficient. To demonstrate the different commuting patterns, Figure 2.1 shows the cumulative number of police stops by race over the average twenty-four hour period. The most salient fact on display is that black drivers' share of stops increases at night relative to white, Hispanic, and Asian and native drivers. This suggests either that the veil of darkness hypothesis is either exactly incorrect, or – more likely – black drivers travel more frequently at night than drivers of other races. Lamberth [2003] finds in a traffic survey that ninety-five percent of drivers' races could be identified during the daylight hours but that nighttime observations were impossible to record without additional lighting, leading to the conclusion that the commuting pattern hypothesis is likely to be correct.

To summarize, the results presented here indicate that when the weather is hotter than usual: the police increase the number of traffic and pedestrian stops they make each week; the police become more strict with speeding violations; and the police do not increase the veil of darkness gap. Taken together, the results of the three models suggest large productivity implications for the police and large welfare implications for the public who are being stopped at higher rates and more arbitrarily when the weather is warm. The racial biases exhibited by the police do not intensify on hot days, but the increase in the number of contacts between the police and the public could exacerbate racial disparities in the use of police force without the temperature having a direct effect on police behavior.

The most similar research to this is Annan-Phan and Ba [2019], which estimates the effect of temperature on civilian deaths at the hands of the police. The authors of that work find that the number of assaults among the public and of police increases in temperature, but the rate of fatal police shootings does not increase. However, they also note that on extremely hot days – i.e. days above 32 degrees Celsius – the rate of deaths from Taser and restraints increases. This research expands on the theme of temperature's effect on police activity and demonstrates that the police become more strict with regard to speeding violations as a result of extreme temperatures. In this

sense, the paper contributes to the climate literature in two regards. First, the effect of temperature on the number of stops the police conduct on a weekly basis adds to the literature on climate change and productivity. The number of stops officers make is one of the primary measures of productivity researchers and departments use when evaluating productivity, so the results presented here add to the growing literature on the climate's effect on worker output. Second, by showing that police officers issue citations for smaller speeding infractions when the weather is hot, this paper contributes to the strain of the climate literature examining the role of heat on behavioral outcomes.

The paper proceeds as follows. Section 2.1 discusses the role weather and climate have on worker productivity and personal behavior, as well as the difference between weather and climate and its implication in econometric models. Section 2.2 discusses the data being used in the analysis and provides details on the non-linear estimation strategy. Section 2.3 estimates the three models – weekly stops, speeding, and veil of darkness – and provides robustness checks and alternative model specifications. Section 2.4 concludes.

2.1 Productivity, Conflict, and the Weather

The climate change literature in economics is rich with research on the effect of climate on both productivity and conflict. Recent advances in computational power and the availability of remote sensing data has precipitated a rush to estimate the effect that weather and climate have on a large number of outcomes as the threat of climate change becomes more acute. To properly establish this research's place in the literature, this section discusses the work that has already been done and reviews the difference between weather and climate, which is one of the central themes in the climate literature.

First, the effect of the weather and climate on productivity. In regions across the globe, heat has been shown to hinder worker productivity in settings ranging from manufacturing (Colmer [2019], Zhang et al. [2018], Adhvaryu et al. [2018]) and survey work (LoPalo [2020]) to outdoor work in

the United States (Zivin and Neidell [2014]). Cognitive tasks become more burdensome (Seppanen et al. [2006]), and even household activities are conducted more sluggishly at high temperatures (Kjellstrom et al. [2016]). Even test scores among school-aged children are subject to heat-induced stress (Park [2017]). It is therefore not a stretch to assume that in a high-stakes occupation such as policing, extreme temperatures may affect officers' productivity and judgement on the job.

Next, the role of weather on conflict. Although the connection between police encounters with the public and conflict usually does not end in violence, it is an ongoing source of public debate and concern in the United States. Therefore, situating this research in the climate-conflict literature is appropriate. At both the interpersonal and international scale, extreme weather has been shown to intensify the propensity for violence across the world. At the interpersonal level, the number of assaults, murders, rapes, and other violent crimes increase with hotter temperatures (Ranson [2014]). Assaults by members of the public on other members of the public and on police officers also increases as the thermostat rises (Annan-Phan and Ba [2019]). At the societal level, the effect of abnormal weather on conflict is not limited to hot temperatures: Iyigun et al. [2017] find that cooling between 1400 and 1900 CE, often referred to as the "Little Ice Age," increased incidents of war and conflict between nations in Europe, North Africa, and the Near East. Finally, Burke et al. [2015] use a hierarchical meta-analysis of the conflict-climate literature, finding that a standard deviation increase in temperature increases interpersonal conflict by 2.4% and inter-group conflict by 11.3%.

Experimental evidence supports the hypothesis that behavioral changes are a key factor in heat-induced increases in conflict. Laboratory settings have demonstrated that extreme temperatures at both ends of the distribution induce aggression and hostile attitudes (Anderson et al. [1996]). Outside the lab in a large number of contexts – from driving (Baron [1976], Kenrick and MacFarlane [1986]) to baseball (Reifman et al. [1991], Larrick et al. [2011]) – people act more aggressively at high temperatures. Combining the effect of heat-induced aggression with the high-stress nature of police work could have large welfare implications for the police and the public, something that this research seeks to explore further.

2.2 Empirical Strategy

2.2.1 Data Description

The primary source of data used in this research is The Stanford Open Policing Project, but there is also meteorological data from the University of Idaho’s gridMET dataset and population-level statistics from the United States Census Bureau.

Police Encounters Data

The Stanford Open Policing Project collects data for encounters between the police and civilians through freedom of information requests to state and local police departments (Pierson et al. [2020]). Due to the federalized nature of policing in the United States, there is no national-level source for interactions between police and the public. The Open Policing Project is the most comprehensive data available at this time. The complete dataset contains over 200 million observations, each of which represent a single police encounter. Data include both pedestrian and vehicle encounters, however the majority of the data is vehicle stops as many police departments do not make pedestrian stops available.

Figure 2.2 displays the sample being used in this analysis. The figure is at the county level, with counties in dark blue representing those in-sample. A number of cities and states available through The Open Policing Project are excluded due to certain variables being unavailable. For example, the Oregon and Nevada State Patrols do not provide information on where stops take place. Location data is necessary to match with weather observations in order to go forward with the analysis. The data are provided by either state police patrols or city police departments. In Figure 2.2, city departments are presented at the county level to more easily highlight their inclusion. A number of major cities and state capitals, such as Chicago, New Orleans, Oklahoma City, and others stand out for being city-level departments within states that do not provide data. There are also cases like those of the major city departments in Texas and North Carolina – e.g. Houston, Austin, Raleigh, etc. – that are in states which have state patrol data available. These cities do not

stand out as well as those cities in states that do not provide data.

To get a complete picture of the cities and states included, Table 2.3 presents the number of observations, sample years, and the analyses – main, speeding, and veil-of-darkness – each location is apart of. The sample spans from 2011 to 2019; the largest single source is the California State Patrol, which includes over 24 million observations. Every city or state department listed appears in the primary model specification, which estimates the change in the number of stops in a county per week as a result of abnormal temperatures. The second analysis, testing whether the speed police pull drivers over for changes on abnormal weather days, is much more restricted as only a handful of departments make the speed limit and driver speed data at the location of a stop available. Despite the restricted sample, there are still over 10 million observations in the speeding model due to the large size of the cities and states included. Finally, the veil of darkness model is more inclusive of the total sample than the speeding model. The veil model requires data on the time of day a stop takes place and the race of the driver, which is provided by most of the departments in the sample.

Meteorological Data

Weather data comes from the University of Idaho’s gridMET dataset (Abatzoglou [2013]). The gridMET data provides measures such as maximum and minimum temperature, maximum and minimum relative humidity, wind speed, and precipitation at the 2.5 \times 2.5 mile spatial resolution for the continental United States since 1979. gridMET combines gridded climate data from The PRISM Climate Group with reanalysis data to, creating a precise and high-resolution dataset that is comparable with weather stations but with more consistent temporal and spatial coverage.

For this analysis, daily weather observations from gridMET are used to find the average temperature, relative humidity, wind speed, and precipitation in the county which a police encounter occurs. For example, a stop in Boston on June 1, 2015 is merged with the average meteorological values in Suffolk County, Massachusetts on that date. A majority of the police agencies in the Open Policing dataset do not make the exact location of a stop – e.g. the coordinates or closest

street address – available, necessitating the use of county averages. Each weather variable is extracted as the mean value in the county of an encounter from the continuous gridMET data using the “terra” package in R (Hijmans [2021]).

Figure 2.3 displays the annual mean temperature for each county in the sample. The gridMET meteorological data is matched to the Stanford Open Policing data to create the figure. Counties are divided into five temperature bins to more clearly display the regional temperature patterns present. In the models presented, identification is derived from within-county temperature variability. Therefore, regional geographic temperature heterogeneity plays an important part in determining the difference in police behavior given certain weather conditions. Average temperature on a county-encounter day is the primary explanatory variable of interest, but as part of a series of robustness checks, I also use the other variables available in the gridMET data, including the daily maximum temperature, the wind-chill temperature, and the wet bulb temperature (Stull [2011]).

2.2.2 Estimation Strategy

Following the recent work in the climate change literature, I estimate the three models allowing for non-linearities in the temperature response function. In practice, this means that the response to temperature is flexible and allowed to change based on where in the temperature distribution a particular day falls. Following the work in Auffhammer [2018], I create a series of temperature bins based on the percentiles of the temperature distribution. The primary model, which estimates the change in the number of stops a police department conducts in a given week, is represented by:

$$y_{cwy} = \sum_b \beta_b D_{cbwy}(T_b) + \sum_j \gamma_j D_{cwy}(P_j) + \rho X_{cwy} + \alpha_{cw} + \phi_{my} + \epsilon_{cwy} \quad (2.1)$$

where y_{cwy} is the log number of police stops in county c for week w in year y . D_{cbwy} is the sum of days in week w that county c falls into temperature bin b . The β_b and γ_j terms are non-linear temperature and precipitation response coefficients, respectively. They represent a series of bins to allow for police departments to have differential responses to more or less severe weather. The

β_b terms can be interpreted as the effect of an extra day per week in bin b on the number of stops per week relative to the reference bin, in this case a day between 64 and 69 degrees Fahrenheit. Similarly, the γ_j term is interpreted as the effect of an extra day per week with rain relative to the reference bin of a day with no rain. X_{cwy} are a series of weather covariates and an indicator for whether the police department in county c includes pedestrian stops rather than just traffic stops. The weather covariates include the average relative humidity and average wind speed.

Following work in the climate literature (Schlenker and Roberts [2009], Deschênes and Greenstone [2011], Barreca et al. [2016], etc.) α_{cw} and ϕ_{my} are county-week and month-year fixed effects. α_{cw} captures annual variation occurring in county c during a particular week of the year. For example it's reasonable to expect the week of Memorial Day – typically associated with the beginning of summer in much of the United States – would see more police activity in counties with popular outdoor recreational amenities available. Similarly, the month-year fixed effects capture the effect of events that are shared across all the counties in the sample that occur in a particular month-year. The sample only runs through 2019, but March and April 2020 provide strong intuition for idea: as lock-downs due to COVID-19 were extended across the country, we would expect that the number of police encounters fell as a result of fewer people traveling from home.

More must be said about the structure of the temperature bins. As mentioned, the models presented follow that of Auffhammer [2018] by creating percentile-based bins. This is relatively uncommon in the literature, where readers more frequently encounter five or ten-degree temperature (Fahrenheit) bins. Both approaches allow for agents to respond non-linearly to the experienced temperature, but the Auffhammer approach has the additional advantage of recovering estimates for the most extreme temperature days. Fourteen bins are created: a bin for observations below the first percentile of average temperatures, a bin between the first and fifth percentiles, a bin between the fifth and tenth percentiles, a bin between the tenth and twentieth percentiles, and a bin for each decile until the ninetieth percentile. The bins above the ninetieth percentile reflect the bins below the tenth percentile: one between the ninetieth and ninety-fifth, one between the ninety-fifth and ninety-ninth, and one above the ninety-ninth. The precipitation bins are also non-linear, but do not

take the same percentile approach. Instead, they are divided into seven bins: below five millimeters of precipitation, between five and ten millimeters, between ten and twenty millimeters, then every subsequent bin represents ten millimeter increments until the seventh, which reflects days with more than fifty millimeters of precipitation.

The second and third models – the change in the speed required for drivers to be stopped and the change in the veil of darkness – are slight modifications on the main model. The speed model can be represented as follows:

$$y_{it} = \sum_b \beta_b D_{it}(T_b) + \sum_j \gamma_j D_{it}(P_j) + \rho X_{it} + \alpha_{cw} + \phi_{my} + \epsilon_{it} \quad (2.2)$$

where y_{it} is the difference between the speed a driver was driving when stopped and the posted speed limit on day t . Unlike in the primary model above, the speed model is at the individual encounter level i rather than at an aggregated county-week level. The bins D_{it} then refer to the weather conditions – mean temperature and precipitation – on the day of the stop. Again X_{it} is a vector of covariates containing humidity and wind speed. However, since the model is at the encounter level there are a number of additional variables available: the race, sex, and age of the driver, the time of day the stop occurred, and the age and type of vehicle being driven. Since the stops in the sample are necessarily traffic stops, the indicator for whether the data contained pedestrian stops is not contained in this model. Again α_{cw} and ϕ_{my} are county-week and month-year fixed effects.

Two aspects of the non-linear approach should be noted. First, although the method is more flexible than a simple quadratic approach, an identifying assumption is that the effect within bins is constant. A majority of the bins are approximately five degrees wide, so in practice this assumption means that we would expect the response to a day with mean temperature of 75 to be the same as a day with mean temperature of 79. The Auffhammer percentile-based bins mitigate concerns about this assumption somewhat by creating bins that are slightly less arbitrary than five or ten-degree bins. Second, in Model 2.1 the dependent variable is the log of weekly stops in a county. Wichman [2018] shows that, in the non-linear temperature response models used in the climate literature, we

should be careful interpreting the coefficients as a percentage change in the outcome of interest if the estimated response is sufficiently large. To account for this, Section 2.3.4 includes a robustness check allowing for the dependent variable to adjust according to the suggestions made in Wichman [2018].

The veil of darkness model can be represented by the following:

$$\Pr(\text{black}_i|t, c, D) = \text{logit}^{-1}(\alpha D_i + \beta D_i \times \text{temp}_i + \delta D_i \times \text{temp}_i^2 + \gamma \text{ns}_6t + \rho X_i + \psi_c + \phi_{my} + \epsilon_{it}) \quad (2.3)$$

where $\Pr(\text{black}_i|t, c, p, D)$ is the probability that the driver in stop i is black, given the date and time t , the county c , and the darkness indicator D . D equals one if stop i occurs after dusk and zero if the stop is before sunset. The D indicator is interacted with both the linear and quadratic forms of temperature. The coefficients of interest are those representing the interaction terms, β and δ . ns_6t represents a natural spline of time, reflecting the thirty-minute time period leading up to or following sunset or dusk. It could be the case that the time variable should vary non-linearly as certain periods within the thirty-minute window have different visibility for the police on patrol. For example, the glare ten minutes before sunset might be more severe than thirty minutes before sunset. The spline allows for this sort of variation. X_i is a vector of variables that may influence the officers' decisions to make a stop and/or be associated with black drivers, such as the other non-temperature weather conditions, the gender of the driver, and the type of vehicle being driven. Finally, ψ_c and ϕ_{my} are county and month-year fixed effects, respectively. These are included to soak up any county-specific and month-year specific variation, as described in the previous models.

As demonstrated in Figure 2.1, racial differences in commuting patterns make a simple comparison of light and dark police pullovers insufficient. To account for this, I follow the analysis in Grogger and Ridgeway [2006] and Pierson et al. [2020] by focusing on traffic stops that occurred up to sixty minutes before sunset and up to sixty minutes following dusk, with the thirty minutes between sunset and dusk filtered out to account for the fact that it is neither light nor dark during this period. The sample is also restricted to the sixty days surrounding the commencement or ex-

piration of daylight savings time (Ridgeway [2009]). The reason for restricting the sample to this period around daylight savings changes is that drivers' schedules are plausibly the same during these periods despite the level of darkness varying. That is, it is unlikely that many drivers change their commuting patterns based on whether it is light or dark out at 5:00 PM in, for example, November.

2.3 Results: How Weather Changes Police Productivity and Behavior

The results presented below show significant changes in both the number of stops that police conduct in the course of a week and in the speed that police pull drivers over for. Section 2.3.1 estimates the main police encounters model detailed in Equation 2.1, finding a large increase in the number of stops at hot temperatures and a large decrease at cold temperatures. Section 2.3.2 demonstrates that, rather than police simply responding to changes in the public's behavior, the police themselves become more strict at hot temperatures.

2.3.1 Change in the Number of Foot and Traffic Patrol Stops

Figure 2.4 and Figure 2.5 present the results of Model 2.1 for the non-linear temperature and non-linear precipitation variables of interest. The reference bin – i.e. the bin that the rest of the bin coefficients are interpreted as being relative to – is the ninth bin, from 64 to 69 degrees Fahrenheit. The dependent variable in both figures is the log number of weekly stops that police departments make. Therefore, the results should be read as the percentage change in the number of stops a police department makes in a week if we were to replace one day with temperatures between 64 and 69 degrees with, for example, a day with mean temperatures above 87 degrees.

Both Figures 2.4 and 2.5 show significant non-linearities across the temperature and precipitation distributions. In the figures, the center dot of each bin represents the estimated coefficient and the error bars are clustered standard errors at the state level multiplied by two to roughly approx-

imate a 95% confidence interval. First, the number of stops per week is affected by both low and high temperatures, with the effect being especially pronounced at the extreme ends of the distribution. On the bottom (i.e. colder) half of the distribution, the number of stops decreases by as much as four percent with an extra day per week in the lowest bin relative to a day between 64 and 69 degrees Fahrenheit.

An important thing to note in the weekly stops models presented here represent the change in the response of police to temperature and precipitation, but this response could be a reflection of changing behavior on the part of drivers and pedestrians. It is well established in the literature that people act more aggressively when the temperature is hot (Kenrick and MacFarlane [1986], Reifman et al. [1991], Larrick et al. [2011]), so the coefficients estimated in this section might be explained by the public's behavior rather than the police's.

2.3.2 *Change in the Speed Drivers are Stopped For*

To get a better sense of whether the effect of temperature on police behavior, Model 2.2 estimates the difference in the speed that police stop drivers for. That is, whether the delta between the speed limit and the speed a driver is stopped for changes due to the weather. If the change in the number of stops per week demonstrated in the previous section were due strictly to the change in behavior of the public, then there are one of two possibilities: first, there could be no change in the speed drivers are stopped for. This would be the case if there were more drivers speeding but their average speeding infraction was not more severe at hot temperatures than under normal conditions. Alternatively, if individual drivers drove faster and therefore more aggressively, then we would expect the average speed that police are stopping drivers to increase.

Figure 2.6 plots the results of Model 2.2. Similar to the previous section, the center dot of each bin represents the estimated coefficient and the error bars are clustered standard errors at the state level multiplied by two. The dependent variable is the difference between the speed a driver was traveling and the speed limit when they were stopped by police. For clarity I will refer to this difference as the speeding gap. The coefficients then can be interpreted as the change in the speed

delta as a result of a day falling in one of the thirteen temperature bins rather than being in the reference bin of 64 to 69 degrees Fahrenheit. Again I use the Auffhammer percentile bins to relax the assumption of constant effects within a bin and to better estimate the effect on the extremes of the temperature distribution.

The results in Figure 2.6 show an upside down-U shape, demonstrating that on days with extreme temperatures the police become more strict with speeding relative to days with temperatures between 64-69 degrees Fahrenheit. Since the coefficients represent the change in the speeding gap, the most extreme hot bin's result of approximately -.50 is a half a mile per hour decrease in the speed drivers get pulled over for. The results at the cold end of the temperature distribution are similarly negative, but the magnitude is smaller with the estimated coefficients ranging from -.30 to -.05 miles per hour. The results are consistent with an interpretation of the police not only responding to driver behavior, but changing their own behavior due to the weather. As discussed in the previous section, if the police were simply reacting to the actions of the public, we would not expect to see any change in the speeding gap. However, the model clearly shows this is not the case, and therefore we can conclude that the police are not immune to the stresses and pressures of the weather.

2.3.3 *Change in the Veil of Darkness*

Table 2.4 displays the results of Equation 2.2.2. The dependent variable in each of the four columns is an indicator for whether the driver pulled over by police was black, and the right-hand side variables of interest are the interaction between the "Dark" indicator and the temperature and temperature quadratic terms. Each column includes county and month-year fixed effects and has standard errors clustered at the state level. The only difference between the models is which covariates are included, as indicated by the rows in the bottom half of the table.

The first column includes three weather variables: relative humidity, wind speed, and precipitation. The second includes the gender of the driver, the third includes the age of the driver, and the fourth includes the outcome of the stop (e.g. ticket issued). The results of all four models

are inconsistent and – for the interactions between darkness and temperature – very small when significant. The estimated likelihood values for “Dark” are large, negative, and significant, which is in line with the existing literature (Pierson et al. [2020]), but its interaction with temperature are extremely small and slightly positive.

2.3.4 Robustness Checks and Alternative Specifications

Figure 2.7 plots the results of Model 2.1, but instead of using the Auffhammer [2018] percentile-based temperature bins, ten-degree bins are used. The interpretation of the coefficients is the same: the effect of an additional day per week in a given bin rather than a day in the reference bin. The key difference in Figure 2.7 is that the reference bin is now 60-70 degrees Fahrenheit, rather than 64-69 degrees as in Figure 2.4. As expected, the results remain consistent regardless of how the bins are divided. Police are less active on cold days but substantially more active on hot days, relative to the reference bin. Days above 90 degrees increase the number of stops per week by over five percent, which days below twenty degrees decrease the number of stops per week by three percent.

2.4 Conclusion

This research demonstrates the importance of expanding our understanding of the effect of climate change on behavioral outcomes. The police are a central part of modern American society, and the increased scrutiny that their actions draw, are subject to the same stresses induced by extreme weather as the rest of the public. Their productivity – measured here by the number of traffic and pedestrian encounters they make in a week – increases as the weather gets hotter. However, since this productivity effect may simply be in response to the actions of the public, I also demonstrate that the police become more strict with regard to speeding violations on high temperature days. Non-linearities in the response to temperature also demonstrate that shifts in the temperature distribution itself will have unforeseen consequences, such as when warmer winters

induce more active behavior on the part of the police and public.

Other measures of police productivity and conduct that are not explored in this work may also be affected by the weather. Further research into changes in police violence and the penalties that they issue as a result of the weather would enhance the public's understanding of what to expect with respect to intensifying climate change and potentially point to areas where police reform can be targeted, one of the central tenets of recent social movements such as Black Lives Matter. It is also unclear whether these results are generalizable to other countries, where police practices vary substantially from those in the United States. Estimates of geographic heterogeneity within and between countries could provide further insight into the role of weather on stressful occupations as well as the effect of climate change more generally.

Tables

Table 2.1: Proportion of Drivers Pulled Over Who are Black
Between 5:30 and 6:00PM when it's light and dark out

Proportion Black	
Dark	17.2%
Light	18.9%

City and State Departments Included in the Sample

Police Agency	State	N. Obs	Sample Years		Outcome Models		
			Min. Year	Max. Year	Main	Speeding	Veil
Little Rock	AR	13,641	2017	2017	Yes	No	Yes
Gilbert	AZ	336,607	2011	2018	Yes	No	No
Mesa	AZ	96,621	2014	2017	Yes	No	Yes
State Patrol	AZ	3,481,294	2011	2017	Yes	No	Yes
Anaheim	CA	87,876	2012	2017	Yes	No	No
Bakersfield	CA	153,691	2011	2018	Yes	No	Yes
Long Beach	CA	208,331	2011	2017	Yes	No	Yes
Los Angeles	CA	4,966,403	2011	2018	Yes	No	Yes
Oakland	CA	133,405	2013	2017	Yes	No	Yes
San Bernardino	CA	90,523	2011	2017	Yes	No	Yes
San Diego	CA	382,844	2014	2017	Yes	No	Yes
San Francisco	CA	474,367	2011	2016	Yes	No	Yes
San Jose	CA	152,833	2013	2018	Yes	No	Yes
Santa Ana	CA	46,268	2014	2018	Yes	No	Yes
State Patrol	CA	24,199,710	2011	2016	Yes	No	No
Stockton	CA	41,629	2012	2016	Yes	No	No
Aurora	CO	172,929	2012	2016	Yes	No	Yes
Denver	CO	1,870,609	2011	2018	Yes	No	No
State Patrol	CO	2,642,566	2011	2017	Yes	No	Yes
Hartford	CT	18,435	2013	2016	Yes	No	Yes
State Patrol	CT	1,175,339	2013	2015	Yes	No	Yes
State Patrol	FL	6,622,051	2011	2018	Yes	No	Yes
Tampa	FL	1,301,854	2011	2018	Yes	No	No
State Patrol	GA	1,906,772	2012	2016	Yes	No	Yes
State Patrol	IA	1,391,911	2011	2016	Yes	No	Yes
Idaho Falls	ID	66,071	2011	2016	Yes	No	Yes
Chicago	IL	846,456	2012	2016	Yes	No	Yes
Fort Wayne	IN	152,265	2011	2017	Yes	No	Yes

Louisville	KY	110,959	2015	2018	Yes	No	Yes
Owensboro	KY	6,921	2015	2017	Yes	No	Yes
New Orleans	LA	450,127	2011	2018	Yes	No	Yes
State Patrol	MA	1,883,756	2011	2015	Yes	No	No
Baltimore	MD	854,759	2011	2017	Yes	No	No
State Patrol	MI	782,459	2011	2016	Yes	Yes	Yes
Saint Paul	MN	223,547	2011	2016	Yes	No	Yes
State Patrol	MS	758,412	2013	2016	Yes	Yes	No
State Patrol	MT	682,388	2011	2016	Yes	No	Yes
Charlotte	NC	625,570	2011	2015	Yes	No	Yes
Durham	NC	136,901	2011	2015	Yes	No	Yes
Fayetteville	NC	222,529	2011	2015	Yes	No	Yes
Greensboro	NC	212,847	2011	2015	Yes	No	Yes
Raleigh	NC	337,390	2011	2015	Yes	No	Yes
State Patrol	NC	7,766,734	2011	2015	Yes	No	Yes
Winston-Salem	NC	176,901	2011	2015	Yes	No	Yes
Grand Forks	ND	46,848	2011	2016	Yes	No	Yes
State Patrol	ND	268,925	2011	2015	Yes	No	Yes
State Patrol	NE	3,034,695	2011	2016	Yes	No	No
State Patrol	NH	259,822	2014	2015	Yes	No	Yes
Camden	NJ	195,100	2013	2020	Yes	No	Yes
State Patrol	NJ	2,801,557	2011	2016	Yes	No	Yes
Henderson	NV	111,089	2011	2019	Yes	No	Yes
Albany	NY	47,600	2011	2017	Yes	No	Yes
State Patrol	NY	6,841,977	2011	2017	Yes	Yes	Yes
Cincinnati	OH	207,987	2011	2018	Yes	No	Yes
Columbus	OH	128,157	2012	2016	Yes	No	Yes
State Patrol	OH	6,267,341	2011	2017	Yes	No	Yes
Oklahoma City	OK	726,433	2011	2017	Yes	No	Yes
Tulsa	OK	311,102	2011	2016	Yes	Yes	Yes
Philadelphia	PA	1,865,096	2014	2018	Yes	No	Yes

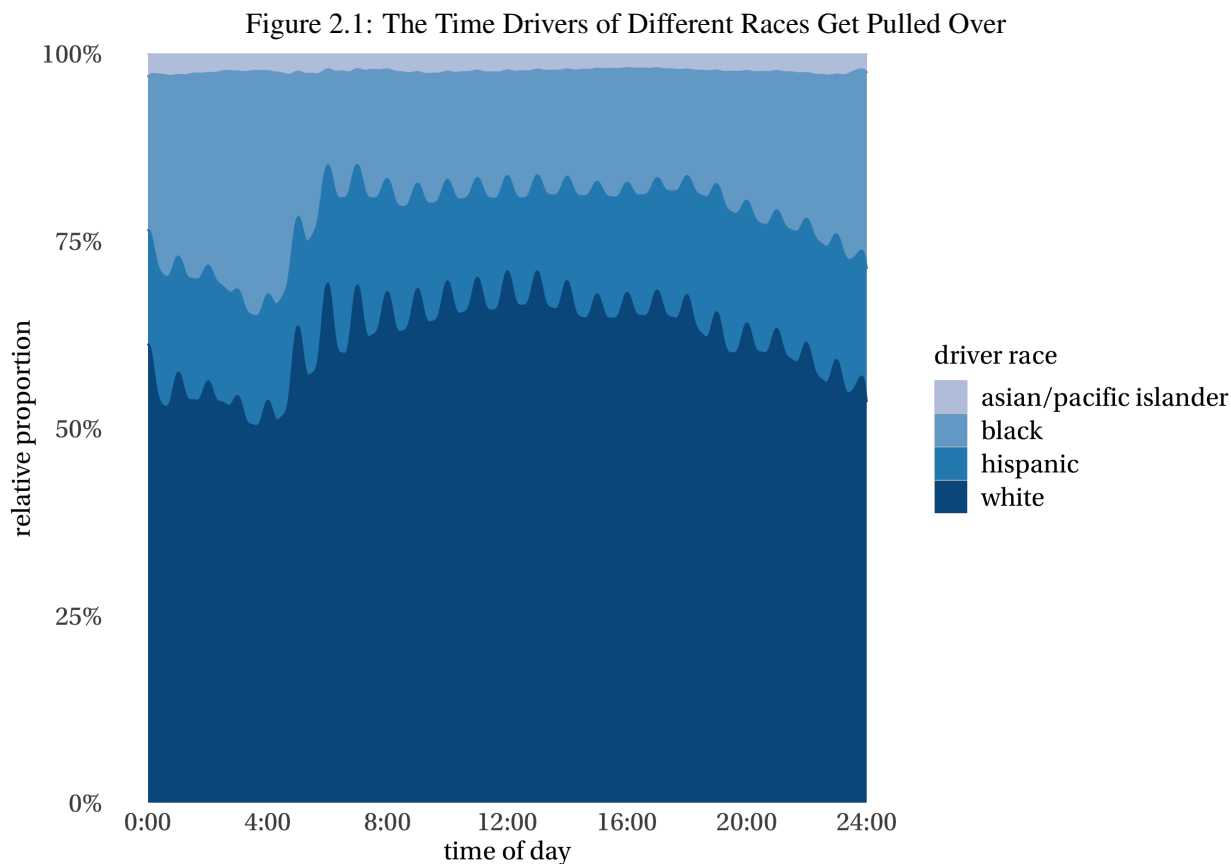
State Patrol	SC	4,382,384	2011	2016	Yes	No	No
State Patrol	SD	435,895	2012	2016	Yes	No	No
Nashville	TN	2,781,729	2011	2019	Yes	No	Yes
State Patrol	TN	2,012,268	2011	2016	Yes	No	Yes
Arlington	TX	112,526	2016	2016	Yes	No	Yes
Austin	TX	259,403	2011	2016	Yes	No	No
Garland	TX	159,840	2012	2019	Yes	Yes	Yes
Houston	TX	1,375,920	2014	2018	Yes	Yes	No
Lubbock	TX	411,790	2011	2018	Yes	No	No
Plano	TX	249,043	2012	2015	Yes	Yes	Yes
San Antonio	TX	1,040,428	2012	2018	Yes	No	Yes
State Patrol	TX	14,812,214	2011	2017	Yes	No	Yes
State Patrol	VA	2,398,402	2011	2016	Yes	No	No
Burlington	VT	31,510	2012	2017	Yes	No	Yes
State Patrol	VT	258,806	2011	2015	Yes	No	Yes
Seattle	WA	121,393	2011	2015	Yes	No	Yes
State Patrol	WA	8,966,807	2011	2018	Yes	No	Yes
Tacoma	WA	141,591	2011	2017	Yes	No	No
Madison	WI	208,310	2011	2017	Yes	Yes	Yes
State Patrol	WI	1,052,838	2011	2016	Yes	No	Yes
State Patrol	WY	172,948	2011	2012	Yes	No	Yes
Total		132,765,275	2011	2019	132,765,275	10,687,063	88,989,983

Veil of Darkness Models

	(1)	(2)	(3)	(4)
Dark	-0.090 (0.057)	-0.132** (0.057)	-0.179** (0.078)	0.025 (0.078)
Avg. Temp	-0.002 (0.002)	-0.002 (0.002)	-0.002 (0.003)	-0.001 (0.003)
Avg. Temp ²	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Dark \times Temp	0.002 (0.002)	0.003 (0.002)	0.006** (0.003)	0.003 (0.004)
Dark \times Temp ²	0.000 (0.000)	0.000 (0.000)	0.000* (0.000)	0.000 (0.000)
Num.Obs.	2,214,542	1,977,556	1,055,907	724,624
FIPS FE	X	X	X	X
Month-Year FE	X	X	X	X
Weather	X	X	X	X
Driver Sex		X	X	X
Driver Age			X	X
Stop Outcome				X

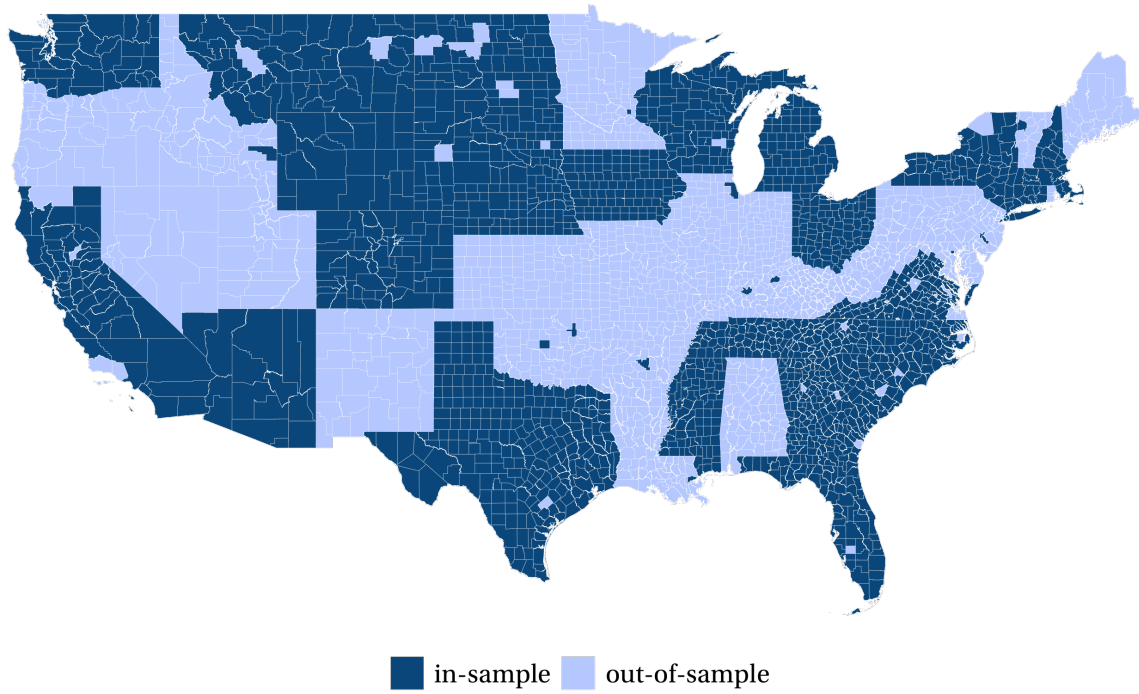
* p < 0.1, ** p < 0.05, *** p < 0.01

Figures



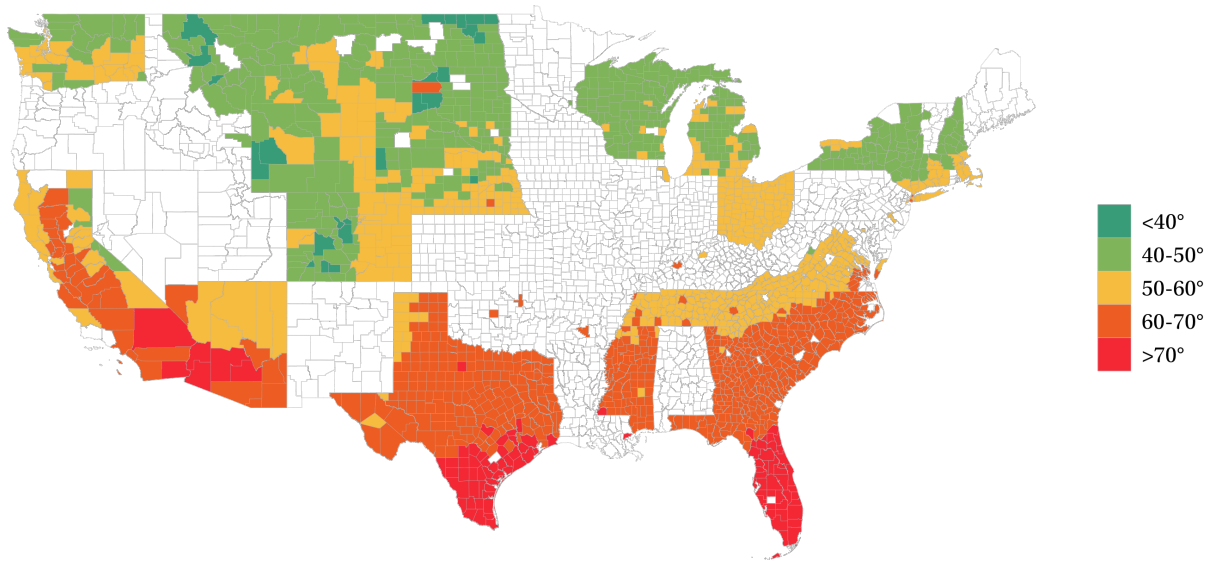
Note: This plot reflects the relative proportion of stops by police for each race in the sample over the course of an average twenty-four hour period. The afternoon shares are largely constant across the four groups, but black drivers are stopped at much higher rates during the overnight hours, suggesting different driving patterns from whites and Hispanics in particular.

Figure 2.2: States, Counties, and Cities in the Sample




Note: This figure plots the state and city police departments in-sample at the county level for clarity. Not all states are fully represented, and some states which provide data from their state patrols also have data from city departments. A prominent example of this is in Texas, where both the state patrol and the departments of major cities such as Houston, Austin, and Dallas all provide data.

Figure 2.3: Average Annual Temperature by County



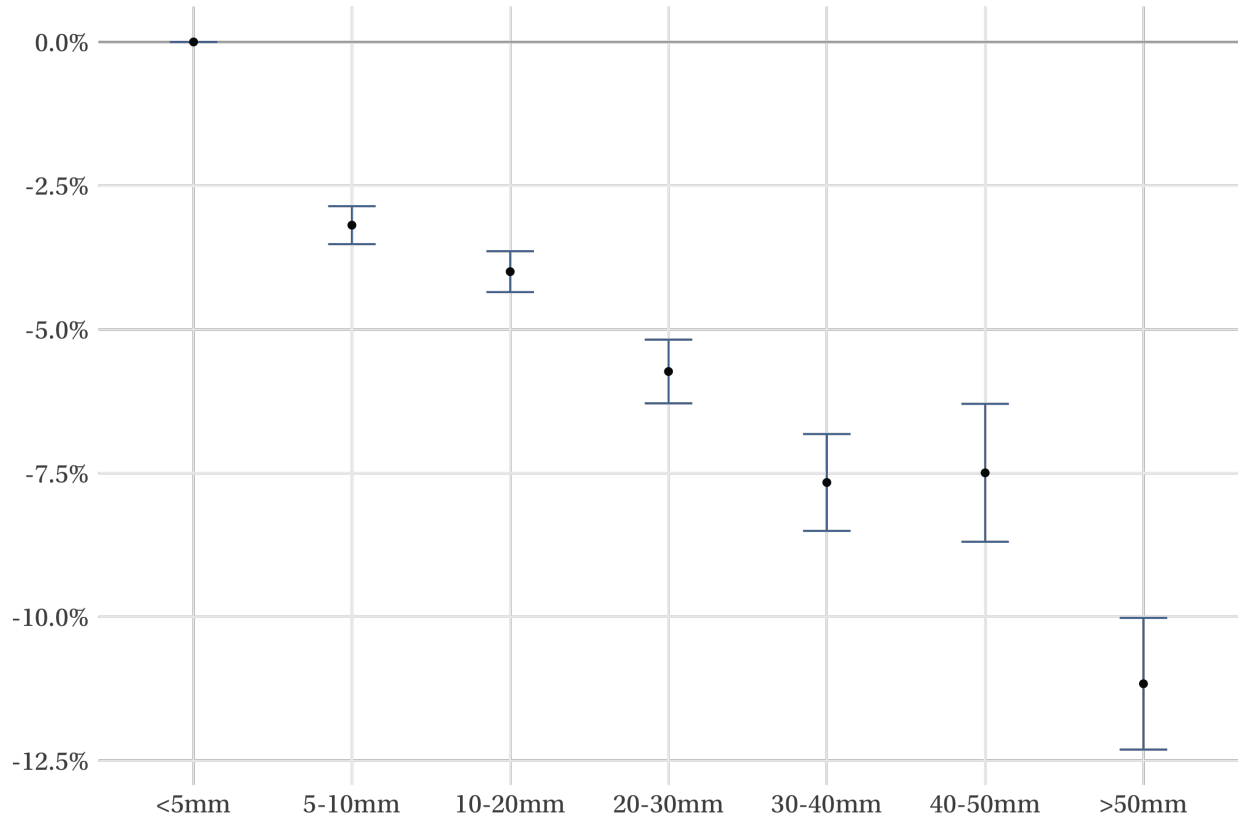
Note: This plot demonstrates the heterogeneity in temperature across the counties in the sample. The counties are grouped into five temperature bins based on their average annual temperature.

Figure 2.4: Change in the Number of Pullovers Per Week




Note: The temperature response function is estimated using a non-linear model to allow for the effect of temperature to vary across the temperature distribution. Cold weather decreases the number of stops police make per week relative to days in the 64-69 degree bin, but hot temperatures increase the number of stops police make.

Figure 2.5: Change in the Number of Pullovers Per Week – Precipitation




Note: This plot displays the estimated coefficients for the non-linear precipitation model. The dependent variable is the log number of stops per week, as it is in the temperature model. The reference bin is zero precipitation. Unsurprisingly, the more precipitation that falls, the fewer stops police make per week.

Figure 2.6: Change in the Speed Pullover



Note: This figure plots the change in the speed required for police to pull drivers over. Again a non-linear approach is taken to allow the response to vary across the temperature distribution. At extreme temperatures, police reduce the speed required to be stopped, suggesting that the weather affects police behavior.

Figure 2.7: Change in the Number of Pullovers Per Week



Note: This figure plots the same temperature response function with the log number of weekly stops as the dependent variable, however the bin structure is modified to include 10-degree bins rather than the percentile-based bins from Auffhammer [2018].

Chapter 3

Retail Marijuana Deregulation and Housing Prices

The rapid legalization of recreational marijuana has created a new industry in the United States. Despite the quick succession of states passing these legalization measures, there is little evidence of how the local economy responds. Immediately upon passage of legalization laws, states increase revenues with marijuana sales taxes and decrease costs by reducing the burden of marijuana-related arrest and incarcerations. Both of these examples create second-order effects on markets which have yet to be considered. This research contributes to the growing marijuana legalization literature by studying the cross-state effect of recreational marijuana legalization (RML) on the housing market.

Other studies has considered the impact of marijuana legalization on residential home prices, many of which are concerned with the effect of marijuana dispensaries. Thomas and Tian [2017], Conklin et al. [2020], Tyndall [2019], and Burkhardt and Flyr [2019] all estimate the housing market response to new dispensary openings in nearby neighborhoods. Among these papers, the evidence is decidedly mixed with negative, positive, null, and positive results respectively. Cheng et al. [2018] use the staggered adoption of city-level marijuana regulations within Colorado to estimate a difference-in-differences model, finding a six percent price increase in the housing market.

Our key contribution to the literature is the estimation of a cross-state model, which is made possible with a rich national level housing data set from the online real estate database Zillow.com. We also provide new evidence of the effect of dispensaries on nearby home values in both Colorado and Washington.

Twelve states and Washington D.C. have passed initiatives legalizing the use of marijuana for recreational purposes since 2012. Additionally, 33 states and Washington D.C. have passed medical marijuana laws since 1996. This quick shift in policy puts the states at odds with the federal government, which still classifies marijuana as a Schedule 1 narcotic on par with cocaine, heroin, and lysergic acid diethylamide (LSD).⁹ The disconnect between the public and the federal government reflects the evolution of the perceived benefits of marijuana. Large majorities of American adults believe that marijuana has medical benefits (Keyhani et al. [2018]), and adolescents have low risk perceptions of the drug (Roditis and Halpern-Felsher [2015]) even though medical professionals are unsure of its efficacy (Kondrad and Reid [2013]; Carlini et al. [2017]; Fitzcharles et al. [2014]; Braun et al. [2018]). Despite the public's beliefs, most states have been reluctant to legalize marijuana for recreational use. Concerns about the potential effect on crime rates and the difficulty in policing impaired driving have been cited as reasons to slow-walk the path to full recreational legalization.

Legalization could increase crime rates, as the drug's effect can make users act more erratically, and easy access to marijuana creates a low-risk trafficking network across state lines. It is well established that crime and the perception of crime negatively impact home prices (Pope [2008]; Buonanno et al. [2013]), so legalization might put downward pressure on the housing markets of states with successful ballot measures. Counter to the crime narrative however, early research suggests that there are no negative effects. Brinkman and Mok-Lamme [2019] find a 19 percent decrease in crime rate in Denver neighborhoods with dispensaries relative to the average crime rate in the sample period. Similarly, Morris et al. [2014] and Huber et al. [2016] find decreases in violent and property crimes following the passage of medical marijuana legalization. There is also evidence that RML increased crime clearance rates by police in Colorado and Washington (Makin

et al. [2019]). Research on traffic incidents suggest similar null or negative results in states with legal recreational (Hansen et al. [2020]) and medical marijuana (Bartos et al. [2018]).

An emerging literature studies the impact of medical marijuana legalization on labor market outcomes. Sabia and Nguyen [2018] find no effect on adult wages, employment or hours worked and a small decrease in wages among young men with access to marijuana dispensaries. Nicholas and Maclean [2019] focus on older adults, finding an increase in the labor supply of those over the age of 51 with the largest effect coming for adults with health conditions which qualify them for legal medical marijuana use. If there are positive labor supply effects, then it is possible that the housing market could be impacted directly through in-migration as individuals from non-legalization states seek to enjoy the perceived benefits. Zambiasi and Stillman [2020] use a synthetic control approach to estimate Colorado's in and out-migration following its passage of RML. Their results suggest that Colorado experienced a large positive inflow of migrants as a result of legalization and no change in out-migration.

Immigration inflows have been shown to increase single family home prices in Switzerland (Degen and Fischer [2017]), but decrease in the United Kingdom as wealthy native homeowners leave the newly immigrant-populated neighborhoods (Sá [2015]). Despite these mixed results, the combination of reduced crime rates (and arrests), migrant inflows, and a new source of sales tax revenue could increase demand for housing in states that pass RML. Some of these states have used the new tax revenue specifically for school funding, which is a mechanism through which home prices might increase. There is a long literature on school resources and student outcomes (Card and Krueger [1998]; Jackson et al. [2016]; Martorell et al. [2016]) and school capital investment's impact on the value of nearby homes (Cellini et al. [2010]; Neilson and Zimmerman [2014]). The combined effect of RML – increased revenue for public goods, decreased crime, little or no change in traffic incidents, and positive labor supply and migration inflow effects – lead naturally to the question of the real estate market. This paper contributes to the literature by estimating the cross and within-state impacts of RML on housing.

First we estimate the cross-state impact using Zillow housing data. The Zillow data is at the

individual property transaction level. The treatment group consists of home transactions in states which have legalized the recreational use of marijuana and the control group consists of home transactions in states which have not legalized it. We find consistent positive effects in the RML case of around 8 percent across a number of specifications which include time and location fixed effects ranging from the county level to the ZIP code level. The estimates are most pronounced when we consider the date that the sale of recreational marijuana is made legal, suggesting that housing demand responds primarily once the drug is being sold, not when the law is victorious at the ballot.

We then extend the cross-state analysis by estimating an unconditional quantile regression (UQR) as in Firpo et al. [2009] with city level fixed effects. Using city level fixed effects controls for unobserved local property taxes which have long been recognized to influence the housing market (Oates [1969]; Anderson [1986]). Doing so provides additional insight into the forces driving our treatment effect. Due to the large heterogeneity in housing markets across the country, the UQR estimates are more robust against extreme value observations than our fixed effects models and provide a more complete understanding of central tendency and dispersion measures. The results of the UQR show positive effects in the top of the distribution following the success of the ballot measure legalizing recreational marijuana, but no effect in the lower half. The greatest impact occurs once it becomes legal to sell marijuana, with large positive effects across the price distribution, especially in the middle three deciles. Heterogeneous responses to a policy shock have not been well-researched in the housing literature, making the findings here one of our major contributions.

Finally, we estimate a spatial model within Colorado and Washington using the Zillow housing data and dispensary location information from the Marijuana Enforcement Division of the Colorado Department of Revenue and the Washington State Department of Health. Our identification strategy follows that of Dronyk-Trosper [2017], who use the staggered construction of municipal buildings such as fire stations to estimate their impact on home prices. In our application, homes which are within two miles of a dispensary at time t and have a second dispensary open within

a half mile of the home at time $t + 1$ increase in value by over 6 percent. The price appreciates the closer to the new dispensary a home is, suggesting that the dispensary itself is a neighborhood amenity which has some positive value among home buyers.

This paper contributes to the existing literature by providing robust evidence that marijuana legalization has beneficial spillover effects at both the state and local levels. Taken together, our three sets of results show that states which pass RML ballot measures benefit relative to other states and that marijuana dispensaries provide a boost to the home values in the immediate vicinity. Marijuana's liberalization provides a novel source of tax revenue which states have used to fund capital expenditures, especially in education and it acts as an amenity via the dispensaries that distribute it. The creation of a new legal market has direct implications for the local economy, as it establishes new dispensary jobs and reduces arrest rates. All of these factors have well-established impacts on housing markets. Indeed our results show that the spillover effects of marijuana legalization on the housing market are both statistically and economically significant.

The paper proceeds as follows. Section 3.1 discusses the history of medical and recreational marijuana legalization in the United States, as well as potential mechanisms through which legalization could impact the housing market. Section 3.2 details three data sources used for estimation and presents summary statistics. Section 3.3 describes the empirical strategy and Section 3.4 presents the impact of marijuana legalization on housing markets.

3.1 Background

3.1.1 Medical and Recreational Marijuana Legalization

Beginning in 1937, the federal government prohibited the use of marijuana for recreational consumption and sale with The Marijuana Tax Act of 1937 (Pub. L. No. 75-238, 50 Stat. 551). The law went into effect on October 1, 1937 and two days later a Mexican-American man named Moses Baca was arrested by Denver police for marijuana possession, the first such arrest in the country.¹⁰ In 1968 Richard Nixon won the U.S. presidency on a platform of law and order, quickly

establishing drug abuse as “public enemy number one in the United States.” The Controlled Substance Act (Pub. L. 91-513, 84 Stat. 1236) of 1970 created tiers of illegal drugs indicating the severity of negative health effects and the level of addictiveness. Marijuana is included in the Schedule 1 tier, indicating that its severity is at the highest possible level alongside addictive narcotics such as heroin. In 1973 the federal government established the Drug Enforcement Agency, which was the primary entity responsible for policing drug use in the country.

Some states introduced marijuana decriminalization proposals in response to the federal government’s aggressive stance on marijuana, but that effort ultimately fell out of favor and the intensity of the War on Drugs escalated in the 1980s and early 90s (Pacula et al. [2003]). In 1996 California became the first state to legalize recreational marijuana, marking the beginning of the end of punitive escalation that began with the Marijuana Tax Act in 1937 and was amplified through the 70s, 80s, and 90s. Once California passed the Compassionate Use Act in 1996, the floodgates were opened and in the ensuing years states across the country legalized marijuana for medicinal purposes. Table 3.2 shows this progress. As of May 2020, 33 states and Washington DC have or are in the process of legalizing medical marijuana consumption.

Despite the progress in MML over the last 20 years, it has been a much slower path to full recreational marijuana legalization. Colorado and Washington were the first two states to approve RML on the ballot in 2012, 16 years after California passed its MML law and after 18 other states had done the same. In the years since, Colorado and Washington have been joined by Alaska, California, Maine, Massachusetts, Michigan, Nevada, Oregon, Vermont, and Washington D.C. Some states have had significant lags between their legalization measures passing a vote and the practical implementation of the law. Massachusetts, for example, voted in favor of RML in November 2016 but it was not until November 2018 that dispensaries selling marijuana opened. It is widely expected that this march of progress will continue in the 2020 election cycle and beyond. This paper contributes another data point to the debate over marijuana legalization, demonstrating that those early adopter states have experienced significant appreciations in home values since legalization has been implemented.

3.1.2 *The Housing Market Connection*

Marijuana legalization comes with a number of trade-offs that make its connection to the housing market ambiguous. The expected direction of legalization's effect depends on a number of forces pushing in opposite directions. Increased public capital expenditures and in-migration would increase demand for housing in the short run and, assuming housing supply is fixed in the short run, raise prices. On the other hand, out-migration, negative health impacts, and increases in crime rates could deflate home values.

To establish the direction of the effect on home prices following marijuana legalization, Figures 1, 2, and 3 show the trend in the national housing market since 2000, divided by when each state adopted RML. There are three cohorts of states. Figure 1 includes Colorado and Washington, the first two states to legalize recreational marijuana in 2012. Figure 2 includes Oregon, which legalized in 2014, and Figure 3 includes California, Massachusetts, and Nevada, all of which legalized recreational marijuana in 2016. The four other states and Washington D.C. which have legalized recreational marijuana are not included because they are outside the sample for reasons discussed in Section 3.1. Solid lines are treatment states across the three figures, and dotted lines reflect states which did not legalize recreational marijuana. To verify that this divergence is a feature of marijuana legalization and not a few wealthy states outpacing the national trend, we divide non-RML states into three groups based on average house price per square foot levels. The six treatment states would all fall into the High average price per square foot grouping with the exception of Nevada, which would be classified in the Middle group if it were not a treatment state. By comparing the trend in those states to other wealthy and middle income states, we can get a better idea of the impact legalization has had on the housing market.

Figures 1, 2, and 3 demonstrate that all three control groups show similar housing market trends since 2000. The RML states meanwhile consistently diverge from the control trends upon their respective cohorts' legalization dates. Across the three graphs, the price trend was similar across RML and non-RML states until 2012. Colorado and Washington display a clear divergence in their housing markets following legalization at the end of 2012. A similar divergence can be

seen in Figure 2 when Oregon voted in favor of RML in 2014. At the end of the time trend, the 2016 legalization cohort also see distinct jumps in the housing markets relative to the non-RML states.

The housing markets of RML states have recovered faster and stronger than those of non-RML states. The effect in Figures 1, 2, and 3 are all despite the period spanning the Great Recession. Volatility in the housing market can be seen clearly in each figure; the market begins accelerating in 2002, peaks in 2006, and reaches its nadir in 2011. The difference in recovery between RML and non-RML states can be seen most dramatically in the first cohort of Colorado and Washington. This could reflect slack in housing as the market over-corrected during the recession, but there can be no doubt that those two states recovered at a faster rate than their economic peers. It appears that the implementation of RML raised house prices despite the burden of the housing market recovery.

3.1.3 Mechanisms

Having established that states which enacted RML laws received a positive boost during the recovery period following the Great Recession, we now turn our attention to the mechanisms responsible. We consider two possible avenues, which we will refer to as the the “economic development” effect and the “amenity” effect. The economic development effect considers long-run changes to the community which legalization induces; increased tax revenue and spending on public goods that results is an example. The amenity effect captures the role dispensaries have on nearby home values. This reflects the local brick-and-mortar changes that occur due to RML. Our cross-state models estimate the economic development effect and our spatial model estimates the amenity effect.

First consider the economic development effect. The illegal marijuana market prior to legalization is necessarily un-taxed. In the political debate over legalization, supporters often advocate for a mandate that marijuana sales taxes fund public goods investment, including infrastructure improvements and education funding. For example the disposition of Colorado marijuana tax revenue is first distributed to the Public School Capital Construction Assistance Fund, and any revenue

over \$40 million is transferred to the Public School Fund.¹¹ There is a long literature on school resources and student outcomes (Card and Krueger [1998]; Jackson et al. [2016]). The physical condition of school capital and government investment as a vehicle for student achievement is also of interest in the existing literature (Martorell et al. [2016]). There is further evidence that school capital investment increases the value of local homes. Cellini et al. [2010] use a regression discontinuity design method, exploiting local referenda on bond issuances for capital expenditures to identify the causal effect of referenda passage on the local housing market. Their results suggest a sizable and immediate positive impact on local home values. Neilson and Zimmerman [2014] study the staggered implementation of a school construction project in New Haven, Connecticut, finding that home prices increase in the local neighborhood by approximately 10%. We contribute to this literature by examining whether the passage of recreational marijuana legalization laws – and therefore new sources of tax revenue – affect local home prices.

Another potential mechanism of the economic development effect is migration. By legalizing the use of marijuana, Colorado and other RML states become an attractive option for residents of other states who value the ability to consume marijuana without fear of legal repercussions. Zambiasi and Stillman [2020] find large migration inflows following Colorado's passage of RML, supporting this hypothesis. For individuals who migrate to a state with legal recreational marijuana, the cost of moving is less than the consumption cost. Those who use marijuana for medicinal purposes could fall into this category, as easy access to legal marijuana decreases the cost of obtaining and consuming an ameliorative drug.

Assuming that housing supply does not increase in response to the success of RML, in-migration of these individuals could affect local housing markets. The effect of inter-country migration on housing markets is ambiguous in the existing literature (Degen and Fischer [2017]; Sá [2015]). However, there is substantial evidence that the number of people migrating within the United States is shrinking and local labor markets conditions and home equity have explain much of the decision to migrate (Henley [1998]; Foote [2016]; Zabel [2012]; Koşar et al. [2019]). Despite this downturn in internal migration, young educated households frequently move to areas with high

quality business environments (Chen and Rosenthal [2008]). Recreational marijuana legalization liberalizes the criminal code, but it also creates a new industry in the states that enact it. Business creation increases employment opportunities and growth (Baptista and Preto [2011]; Andersson and Noseleit [2011]), which in turn puts upward pressure on housing markets (Liu et al. [2016]; Reichert [1990]). Benefits (and potential costs) of industry job creation and demand for marijuana from non-locals could be capitalized into housing values (Cheng et al. [2018]).

We estimate the effect of marijuana legalization at different points of the process (i.e. at the time of the vote to legalize, when the law goes into effect, and when the first dispensaries open), which provides insight into the magnitude of the economic development effect. Since the two-way fixed effects and UQR models define treatment as all homes in a state, the coefficients should reflect the broad treatment inside each state. Homes without nearby dispensaries therefore are likely not experiencing the positive shock through an amenity effect, but through secondary mechanisms such as increased school funding and capital investment. We estimate the UQR model to capture the sensitivity of the price distribution to the economic development effect. The hedonic price function frequently estimated in the housing literature can be highly non-linear. For this reason, the UQR model is our preferred model specification and the primary contribution of this research's estimates of RML on the economic development effect in housing.

The amenity effect will be captured in our Spatial Difference-in-Differences model (see Section 4). By restricting our sample to just homes near dispensaries in Colorado and Washington, we recover the dispensaries' effect on the nearby housing market. This approach is in line with previous research, as prices exhibit localized variation based on a number of amenity factors, including public school quality (Bogart [2000]; Cheshire and Sheppard [2004]), public transit options (Bajic [1983]; Dewees [1976]), water quality (Epp and Al-Ani [1979]; Young and Teti [1984]; Leggett and Bockstael [2000]), rail lines (Bowes and Ihlanfeldt [2001]; Gibbons and Machin [2005]; McMillen and McDonald [2004]), and crime (Hellman and Naroff [1979]). Home prices vary significantly as households are heterogeneous in their amenity preferences and income (Gibbons and Machin [2008]). If dispensaries are an amenity – either positive or negative – then we should

be able to recover an effect with the Spatial Difference-in-Differences model. Indeed other research has estimated the dispensary-housing market connection (Conklin et al. [2020]; Tyndall [2019]; Burkhardt and Flyr [2019]), but either did not use a spatial model as part of their identification strategy or are limited to particular cities which might raise external validity concerns. Recovering the amenity effect of dispensaries in Colorado and Washington using a novel estimation method is the second major contribution of this research.

3.2 Data

This research relies on three primary sources of data. First is a national housing data set from the online real estate database company Zillow (Zillow [2017]). The second is a hand-compiled data set identifying each states' laws regarding the liberalization of marijuana use. Finally, we have yearly data on the construction of marijuana dispensaries in Colorado and Washington.

3.2.1 *Housing Data*

Zillow is a popular tool used by the public to search for properties available for sale in the United States. The company provides a centralized source of property transactions through its Zillow Transaction and Assessment Dataset (ZTRAX).¹² This dataset compiles multiple listing services (MLS) from all fifty states, Washington D.C., and other U.S. territories to provide a comprehensive resource for real estate transactions.

The information includes not only details of a given housing market transaction, such as the sales price and date, but also information about the house itself. The ZTRAX repository provides access to a large number of home characteristics, such as the number of rooms, square foot area of the property, and any structures on it. Table 3.1 shows the summary statistics for all homes in our sample, as well as annual state-level economic variables, such as GDP. The differences among both the home characteristic and local economic variables suggest that local fixed effects will be an important factor in our model specifications.

We consider all homes in each state, conditional on the data being representative of a state's housing market. This is not the case for every state, as some do not have MLS public reporting requirements across all counties. For example, North Dakota has only one county which consistently reports transactions to the state's MLS, so we exclude it from our sample. Additionally, since this research is interested in the spillover effect of marijuana legalization of the housing market, we only consider homes which Zillow documents as residential properties. The richness of the data means that some states report business, government, and other non-residential properties. We exclude these observations.

The data is also filtered for observations that are likely non-market transactions. All included observations are categorized as a deed transfer, which signifies the exchange of a property's title from one party to another. Despite this, there are observations where a non-market transfer occurs between, for example, family members in the case of inheritance. These types of observations are often indicated as such, but in order to further exclude cases where reporting standards differ, we also filter for transactions which have a listed sales price below \$10,000 and above \$10,000,000. Doing so substantially reduces the sample size, but it is unlikely that homes below that price are actual market transactions given the price distribution. Additionally, states that have fewer than 100,000 transaction across the sample period are excluded in order to reinforce that a state's housing market sample is properly represented. We provide a more comprehensive examination of our data cleaning process for the Zillow data in Appendix A.

3.2.2 *Marijuana Laws*

In addition to the housing and dispensary data, we used the legalization dates as determined by each state to identify our treatment conditions. As mentioned in the introduction, there are three possible legal states that marijuana can be classified as: legal to use recreationally, legal to use medicinally, and illegal. We used successful laws and ballot measures to indicate the relative legality of marijuana in each state. The information in this data is presented in Table 3.2. The second column reflects the date that a given state votes for and passes recreational legalization.

The third column is the “effective date” for recreational legalization when either the result of a popular vote is approved or a law goes into effect. This is the date when it is no longer illegal to possess or grow marijuana for recreational purposes.

It is not until the date in Column (4) that there is a way to legally purchase recreational marijuana. An important distinction to note is the difference between the “Dispensary Date” and “First Dispensary” columns. In some cases, the ballot question outlines a specific date on which dispensaries are allowed to open. This is not always the case, however, as some states leave the decision when to open dispensaries up to local municipalities. This distinction is why Dispensary Date and First Dispensary are considered two separate treatments. Some states, such as California and Colorado, specify the Dispensary Date in their ballot questions, and as a result have dispensaries open on that date. In that case, the Dispensary Date and First Dispensary column dates are identical. Other states such as Massachusetts and Maine have large time gaps between the two dates due to local governing bodies having discretion over dispensary permit approvals. The preferred treatment and what is presented in our primary models is the Dispensary Date. We provide separate estimates for both variables, and consider the First Dispensary treatment as a robustness check.

We use a similar logic for cases of medical marijuana legalization. This process is significantly more complicated, however, as the regulations enacted by each state vary widely. A state may vote via a ballot measure or through the state legislature to legalize the use of marijuana for medicinal purposes, but the process following that approval has many additional steps. Similar to the recreational case, the law becomes effective as soon as it is passed, but the possession of marijuana is not necessarily legal due to the method through which the state distributes licenses. California, which was one of the first states to enact medical marijuana legalization, distributed medical license cards similar to a driver’s license for those eligible for marijuana possession. Additionally, there are complications with prescriptions that vary by state which add a layer of complexity to identifying the timing of our effective date. It is also not always clear whether dispensaries that can sell medical marijuana to users with a valid prescription have opened, or if there is some other distribution mechanism that the state has adopted. As a result, we use a similar logic to the recre-

ational case and consider the effective medical marijuana legalization date to be the date that a ballot measure is ratified or a state legislative measure is signed by the governor.

3.2.3 *Dispensary Data*

For our spatial analysis we use data from the Marijuana Enforcement Division of the Colorado Department of Revenue and the Washington State Liquor and Cannabis Board, which detail every dispensary location in the two states since their legalization of recreational marijuana. These data include the spatial coordinates of a given dispensary and the year it opened. Our estimation focuses on the opening of new dispensaries, so the data begins in 2014 when the first strictly recreational dispensaries opened in Colorado and Washington. It is worth noting however that there existed dispensaries in both states prior to recreational legalization due to the previous passage of medical legalization. Those dispensaries are taken as given and exist at the start of the data. The spatial identification strategy depends on the opening of new dispensaries, so whether a dispensary was an already-existing medical dispensary should have no bearing on the validity of the estimation. We combine the dispensary data with the Zillow housing data to estimate the effect of new dispensaries opening on the housing market in the immediate vicinity. This represents the within-state amenity effect of legalization.

3.3 Empirical Strategy

Our empirical strategy involves three primary specifications. First is a linear model, which we test with varying fixed effect levels to establish a baseline relationship between marijuana legalization (both MML and RML) and home prices. We estimate the following: $\log(\text{Price}_{ijst}) = \alpha_1 \text{Recreational Vote}_{st} + \alpha_2 \text{Recreational Possession}_{st} + \alpha_3 \text{Dispensary Date}_{st} + \alpha_4 \text{Medical}_{st} + \beta X'_{ijst} + \delta_j + \rho_q + \epsilon_{ijst}$

Since the Zillow housing data is at the transaction level, our primary dependent variable $\text{Price}_{i st}$ is the price of home i in county/city/ZIP j and state s at time t . In this simple model the variables

of interest are $\text{Recreational Vote}_{st}$, $\text{Recreational Possession}_{st}$, $\text{Dispensary Date}_{st}$, and Medical_{st} , which are all binary variables indicating whether state s has adopted RML (for Recreational Vote, Recreational Possession, and Dispensary Dates) or MML (for Medical) at time t . $\text{Recreational Vote}_{st} = 1$ if the state has approved RML by ballot vote or a legislative statute by the transaction date, $\text{Recreational Possession}_{st} = 1$ if the RML law has gone into effect and it is legal to possess marijuana, $\text{Dispensary Date}_{st} = 1$ if dispensaries can apply for permits to sell recreational marijuana, and $\text{Medical}_{st} = 1$ if MML has been approved by state voters or legislators. In addition to these indicators, X'_{ijst} is a vector of housing characteristics and local economic measures including the number of bedrooms, bathrooms, the age of the home, state GDP, state population, and state land area. Finally we include location and time fixed effects, δ_j and ρ_q , respectively. We use year-quarter fixed effects for ρ_q , but the legalization dummies are defined by the exact date of RML voting, possession, and dispensary openings. This makes our models traditional hedonic estimations.

The second model employed is an unconditional quantile regression (UQR), as specified by Firpo et al. [2009] (FFL). Table 3.1 demonstrates the large amount of variation across the data, especially with regard to our outcome variable of choice, home price. The observed prices and house characteristics exhibit significant heterogeneity, which makes a UQR an attractive estimation strategy. As we demonstrated in Figures 1, 2, and 3, response to the housing recovery varied widely between RML states and non-RML states. Extending this idea to the distribution of prices, a UQR model accounts for systematic differences across states that may influence their decision to pass legalization measures. The UQR model is evaluated on the distribution of independent variables marginally. Because of this, the model does not depend on the covariates conditioned on as in a traditional conditional model.

The UQR model evaluates the impact of RML and MML on house prices across the price distribution using a recentered influence function (RIF) (Hampel et al. [2005]). Although the RIF can be applied to any distributional statistic, FFL use it to estimate quantiles along the distribution. The marginal effect of any quantile on the home price can be represented by: $E[\text{RIF}(\text{Price}_{ijst}; q_\tau) | \text{RML}, \text{MML}, X, \delta, \rho]$

$$\alpha_1 \text{Recreational Vote}_{st} + \alpha_2 \text{Recreational Possession}_{st} + \alpha_3 \text{Dispensary Date}_{st} + \alpha_4 \text{Medical}_{st} + \beta X'_{ijst} + \delta_j + \rho_q + \epsilon_{ijst}$$

Model 3.3 is the same equation as in Model 3.3, with the only difference being the estimation of the RIF. q_τ in the RIF reflects each quantile being estimated. In our case we will derive estimates for each decile along the price distribution (i.e. $q_\tau = (0.1, 0.2, \dots, 0.9)$). By estimating each decile, the RIF allows us to interpret the effect of RML across the distribution which may provide additional insight into the mechanisms behind legalization's impact on the housing market.

Like the fixed effects Model 3.3, the UQR estimates the difference in home prices along the distribution across states. It could be the case that there are differences within states that legalized marijuana use as well. To test this we use data from the Marijuana Enforcement Division of the Colorado Department of Revenue, the state agency in Colorado tasked with regulating the sale of marijuana, and the Washington State Liquor and Cannabis Board. The agencies' data provide the location of marijuana dispensaries opened in the states between 2014-2018. By combining this data with the Zillow housing data, we are able to estimate the effect of a dispensary opening on neighborhood home values.

A clear source of endogeneity in a standard difference-in-differences (DiD) approach is that the location of a dispensary is not random; a firm chooses what it believes to be the most profitable location for its dispensary and finds suitable properties to rent or purchase. The firm may rent property in a business district or near transit, which could bias the housing market in the immediate area upward. On the other hand if these are new or inexperienced businesses that have capital constraints, they might locate where property is relatively inexpensive. This would have the opposite effect, as homes in less dense areas are generally on the lower tail of the price distribution.

To account for the endogeneity concern, we use a DiD approach developed in Dronyk-Trosper [2017]. The authors use the local government's construction of public service facilities, such as fire departments and police stations, to identify changes in the local housing market. Control homes are those which maintain their distance from the closest facility throughout the sample period. Treatment homes are those which – at period t_0 – have the same distance as the control group

but at some future period t_s , where $s > 0$, a new facility is constructed that reduces the distance to the nearest option. We modify this approach by substituting the public facilities for marijuana dispensaries. The spatial DiD model is represented by:

$$\log(\text{Price}_i) = \beta_1 \text{Treatment}_i + \beta_2 \text{State}_i + \beta_3 (\text{Treatment}_i \times \text{State}_i) + \gamma X_i + \epsilon_i \quad (3.1)$$

with Treatment_i is an indicator variable which reflects whether a home is in our treatment group – whether a new dispensary has opened closer to home i since period t_0 . State_i is a dummy for whether a home sale occurred before or after the construction of a new closer dispensary, and X_i is a vector of home characteristic controls. β_3 is our variable of interest, which represents the change in home values for treated units following the opening of a new dispensary. Figure 4 demonstrates the buffer zones around marijuana dispensaries in the Denver metropolitan area and the homes that fall within the buffer zone. For the purpose of Model 3.1, only a subset of the homes that appear in Figure 4 will be included in our treatment group.

3.4 Results

3.4.1 *Housing Prices Following Statewide Marijuana Legalization*

Tables 3.3 and 3.4 estimate the effect of recreational marijuana legalization on housing prices using a simple linear model and a fixed effects model, respectively. In these tables and in the rest of the main specifications, the dependent variable is the logged value of home prices. Each column in the two tables includes a single treatment variable with the exception of Column (5), which includes three treatment variables. The treatment variable indicating the date recreational marijuana possession is legalized is excluded in Column (5) because, as indicated in Table 3.2, the gap between the vote and possession dates are typically no longer than a month. If this gap is longer than a month, then the possession date is typically very close to the first legal sales date. We estimate the coefficient for possession separately in Column (2) of Tables 3.3 and 3.4, and as

expected its point estimate falls between the vote and sales points estimates.

In Table 3.3, as in the rest of the tables that follow, each estimation includes variables which control for house characteristics and state economic indicators. Table 3.3 includes city-level clustered standard errors to account for potential correlations of error terms, but does not include any fixed effects indicators. In this simple linear model the estimated coefficients of interest are large and significant, with each point estimate reflecting greater than a eighteen percent appreciation in home prices for the RML variables of interest. Table 3.4 includes city and year-quarter fixed effects for the same five estimations as Table 3.3. This table represents the primary linear cross-state results. Similar to the previous table, we find large and positive estimates for the three RML treatment indicators, again exceeding ten percent when considered individually. A noteworthy difference between the fixed effects and OLS models is the magnitude of the coefficients. Including fixed effects greatly reduced the estimated effect, which is to be expected considering the data is a national sample which features large amounts of heterogeneity in housing and economic characteristics.

The model is designed to identify the effect of RML specifically, but we include the medical coefficient in order to address the potential endogeneity issue of states voting in favor of recreational legalization. Policy treatment represents a selection issue as voters choose whether to vote in favor of marijuana legalization. As seen in Table 3.2, however, there are a large number of states which have legalized medical marijuana but only ten which have legalized recreational marijuana. Due to the limitations of the Zillow housing data discussed in Section 3.2.1, the only states which are in the RML treatment group are California, Colorado, Massachusetts, Nevada, Oregon, and Washington. RML treatment states make up less than a quarter of the MML states as a result. Every state that has enacted RML has enacted MML, but the inverse is not true. By including the medical treatment in our primary model specification, we cannot guarantee the consistency of the medical coefficient but we should recover the marginal effect for the two RML treatment variables.

Column (5) of Table 3.4 demonstrates that once we include city and year-quarter fixed effects into our primary linear model, both Recreational Vote and Dispensaries Date's coefficients retain

large, positive, and significant point estimates. The larger effect happens at the Dispensary Date, when the first dispensary *could* open. This estimate reflects an eleven percent appreciation in home prices. As explained in Section 3.2.2, this is not necessarily the date that the first dispensary opens since each municipality in a given treatment state has different permitting rules for new businesses. As a robustness check, we use the opening date of the first dispensary in a state as the dispensary treatment and find qualitatively similar results. The estimated coefficient for the Recreational Vote treatment meanwhile reflects 5.4 percent price appreciation. Taken together, the two linear models support the hypothesis that RML induces large positive effects in the housing market.

To further test the state-level effect of marijuana legalization on housing prices, we estimate an unconditional quantile regression (UQR) as specified by Firpo et al. [2009]. A UQR has three principle advantages over a traditional linear model despite the fact that it simply recovers the marginal effect of the treatment indicators. First, it is less sensitive to extreme values in the dependent variable. This is unlikely to be an issue in the data used for this paper as the number of observations is substantial, but it is nonetheless a strength of the model. Second, a UQR model accounts for differences across states that could affect the likelihood of a given state passing a marijuana legalization bill, which is a significant concern. Finally it marginalizes the treatment effect across the price distribution, which provides a more complete understanding of the impact of RML on the housing market.

With those advantages in mind, Figures 5 and 6 plot the UQR coefficients for each decile along the distribution. For a more precise view of the estimated coefficients, Appendix Tables 3 and 4 in Appendix B display the point estimates. Again we have estimated two model specifications, one with the Dispensary Date treatment and one with First Dispensary due to the close time proximity of those two variables. A pattern emerges in both cases: there appears to be some significant effect in the Medical Vote or Recreational Vote treatments and a significant, positive, and increasing effect across the Dispensary Date/First Dispensary distributions. The Recreational Vote treatment show some significant appreciation in the top four deciles, but as in the linear models the Medical coefficients should be interpreted conservatively.

The positive effect in the upper deciles for the two Vote treatments range between a three and twelve percent increase in home price. The concentration, especially in $Q_\tau = .80, .90$ could point to the level of liquidity available to those purchasing the most expensive properties. For example, if those wealthy buyers have greater access to credit than buyers lower in the distribution, then their demand for marijuana and in turn housing in RML or MML states could shift immediately upon the success of a ballot measure. This interpretation would be consistent with the economic development hypothesis presented in Section 3.1.3; demand for housing is responsive to employment gains, which itself is a natural byproduct of new business creation, and potential in-migration. The results support the those from the linear fixed effects model estimated in Table 3.4, with the top two deciles dominating the average effect,

The Dispensary Date and First Dispensary treatments differ from the two Vote treatments in that they have large, positive, and significant effects across the price per square foot distribution. These values range from approximately seven percent to nineteen percent, with the point estimates increasing in magnitude until beginning to decrease at the 7th decile. It should be noted that the values in the 8th and 9th deciles have very large confidence intervals and so the point estimates may be overstating the effect. Regardless of the estimated confidence intervals, we can say with some certainty that the two dispensary treatment dates reflect a shift in housing demand in RML and MML states. This large effect again supports the hypothesis that the economic development effect drives the change in the housing market. Once recreational marijuana becomes available to buy easily at a dispensary and tax revenue is generated, there is significant home price appreciation.

3.4.2 *Spatial Model*

To further test whether it is open dispensaries that are driving the increased demand for housing, we estimate the results from a spatial model which identifies the effect of new dispensaries on the value of nearby homes. The model, which is described in Section 3.3 and follows the empirical strategy developed in Dronyk-Trosper [2017], estimates the within-state effect, as opposed to the cross-state effect of the linear and UQR models presented in the previous section. The various

treated groups in this model represent homes which have already been “exposed” to a dispensary by having a dispensary open within a two-mile radius of the property. They are then considered treated when a second dispensary opens geographically closer at a later date. Figure 4 demonstrates this idea graphically.

In order for this empirical strategy to be valid, homes in the treatment groups must not differ from each other in price and house characteristics. Table 3.5 presents the mean and standard deviation values for the four groups. The group “Inside 0.5 Miles” includes all homes sold which were within a half mile of a dispensary at any point in the sample period of 2014-2018 in Colorado and Washington; “Between 0.5 and 1 Mile” includes homes sold which were between a half and one mile of a dispensary at any point in the sample period; “Between 1 and 2 Miles” contains homes sold which were between one and two miles of a dispensary at any point in the sample period; and the “Outside 2 Miles” group includes homes which are outside a two-mile radius of any dispensary.

Table 3.6 presents the results for the spatial difference-in-differences models. Like the linear and UQR estimates in the previous section, each of the models have the logged value of price as the dependent variable. Column (1) is a simple fixed effects model, where the point estimates for 1/2 Mile Zone, 1 Mile Zone, and 2 Mile Zone reflect the premium for homes within a two mile radius of a dispensary in Colorado and Washington during our sample period. This model in this column has no causal mechanism and simply estimates the mean difference between homes near (i.e. within two miles) of a dispensary and those outside that bound. Homes within 0.5 miles have a slight premium of 4.5 percent, but homes between 0.5 miles and one mile and homes between one and two miles have a slight discount.

The primary spatial model specifications appear in Columns (2) and (3) of Table 3.6. Both columns follow the identification strategy in Dronyk-Trosper [2017], and so can be interpreted as the causal effect of a marijuana dispensary opening on the local housing market. Column (2) uses homes within two miles of a dispensary as the control group. The two treatment variables – 1/2 Mile Zone and 1 Mile Zone – are indicators for homes which previously were within two miles

of a dispensary and were subsequently sold after a new dispensary opens. The sold homes are newly situated within a half mile or between a half mile and a mile of a dispensary, respectively. The coefficients for 1/2 Mile Zone and 1 Mile Zone represent the premium for these homes. Both treatment zones experience an appreciation in price after the construction of a new dispensary. The 1 Mile Zone homes increase in value by slightly under one percent and the 1/2 Mile Zone homes increase by slightly over seven percent. Column (3) is the same specification, except now the only treated homes are those within a half mile of a new dispensary. The homes in 1 Mile Zone that were previously considered part of our treatment group in Column (2) are now included in the control group. Again the estimated coefficient for the half mile group is significant and positive with an eight percent appreciation. In order to guarantee that the results are not being driven by one of the two state's effect dominating the other, we separate the sample into tables for Colorado and Washington as a robustness check. Appendix tables 5 and 6 appear in Appendix B. The results are similar between the two states and between the individual state estimates and the combined estimates, suggesting that this effect is not due to one state's influence.

Dronyk-Trosper [2017] find that the effect of municipal government service buildings, such as police stations and firehouses, increases the value of homes at a decreasing rate. Those homes closest to the government buildings actually decrease in value, likely as a response to the increased traffic and noise associated with those services. Our results imply the opposite; when a dispensary opens nearby, homes closest to it appreciate in price the most. This is consistent with our interpretation that new dispensaries act as amenities in the local housing market. Since the spatial model is restricted to Washington and Colorado – the first two states to legalize recreational marijuana – we cannot guarantee that these results generalize to each subsequent state that legalizes. However, together with the cross-state models presented in the previous section, it is clear that recreational marijuana legalization has large positive effects on the housing market of states that legalize and municipalities which allow dispensaries to open in their communities.

3.4.3 *Robustness Checks*

There are two primary robustness check categories we employ. First, we use the home price per square foot as the dependent variable rather than home price. Geographic heterogeneity in our sample suggests that simply using house price as the dependent variable could bias the results since treatment homes are in high-price states. By using house price per square foot as the dependent variable, we can ensure that this potential source of bias is accounted for. Second, we include the First Dispensary treatment in place of the Dispensary Date variable for the reasons outline in Section 3.2.2. If the primary mechanism in our cross-state models is the economic development effect, then it is possible that the impact is only felt once the first dispensaries open and a large volume of marijuana sales take place, thereby generating tax revenue.

Appendix Table 1 uses the log value of house price per square foot as the dependent variable in the two linear cross-state models. In this table, Dispensary Date is still the right-hand side treatment variable of choice. As in the price per square foot results, the OLS model in the first five columns shows large positive results for all four treatment variables, including the Medical Vote treatment. Again, these results should be interpreted carefully as the Medical Vote treatment is likely absorbing a large amount of the effect due to the lack of time fixed effects. That being said, the point estimates are very similar to those presented in Table 3.3. The same can be said for the fixed effects results in columns (6) through (10). The Recreational Vote variable is still significant and positive, as is the Dispensary Date. The point estimates are large and positive, as in the original specification.

Next, we check our results using First Dispensary as our treatment variable of interest rather than Dispensary Date. For some states these dates are the same, so we would expect the results to be very similar. Appendix table 2 presents the estimates, and indeed that is what we find. The results are consistent with the Dispensary Date results. Once again, there are positive effects for each of the two RML variables, Recreational Vote and First Dispensary, just as in our primary results. The magnitude of the First Dispensary estimates are similar to those for Dispensary Date presented in Table 3.4. Appendix table 2 also presents the original model specification with various

levels of controls. Excluding house characteristic and local economic variables do not affect the magnitude or significance of the estimated models.

3.5 Conclusion

Uncertainty regarding the costs and benefits of marijuana legalization, along with marijuana's status on the federal level as a Schedule 1 drug, have made the public reluctant to support policies which liberalize its use and distribution. To help fill this information gap, this research demonstrates that there is a large positive spillover effect on the housing market following legalization. We further support these findings with a spatial approach which shows that within states that legalize recreational marijuana use, homes experience a positive valuation shock when a dispensary opens nearby. The results are robust to a number of specifications, including a different (but temporally similar) date for the actual sale of marijuana at dispensaries. Taken together, the inter and intra-state results suggest that preferences for public services – derived from a new source of tax revenue – and dispensaries as a commercial amenity create largely positive effects following the legalization of recreational marijuana.

The impact of legalization on the housing market is supported by two models. First, a fixed effects model demonstrates a five percent appreciation in home prices following the passage of RML and an eleven percent appreciation once sales of marijuana products begin. Extending this logic to an unconditional quantile regression approach, we find positive effects across the home price distribution following the date that dispensaries are allowed to open. Differences across the price distribution can likely be thought of as heterogeneous preferences among different levels of wealth. The promise of future funding to schools and other public infrastructure as a result of legalization supports a long literature showing a positive relationship between home prices and local economic development.

To approximate the effect of dispensaries we estimate a spatial model in Colorado and Washington. The results again show price appreciations for homes as the distance to the nearest dispen-

sary decreases. This demonstrates that it is not simply the benefits of increased tax revenue, but also the existence of the dispensaries themselves, that is driving the price increases. The dispensaries act as commercial amenities that the public puts a premium on being nearby.

Without the benefit of foresight, our research is not able to determine whether the positive effect will persist. For example if immigration inflows are the primary cause of our results, then we would expect that states would experience diminishing returns to legalization. The first cohort of states which legalized recreational marijuana would draw those that valued legalization most, and each successive state should not expect a similar inflow. Additionally, more research on marijuana legalization is required to fill in the remaining knowledge gaps. We do not estimate some of the other second-order effects, such as the impact on policing and the outcomes for minority communities that were previously convicted for marijuana possession at a disproportionate rate. Future research would be well served to approach these questions, as it will better inform the public and policy makers with respect to the reclassification of recreational drugs.

3.6 Tables

Summary Statistics					
Variable	Mean	Std. Dev.	Min.	Max.	N
Transaction Prices (\$)					
House Price	330,342	364,989	10,838	9,999,181	38,145,054
log(House Price)	12.35	0.86	9.29	16.12	38,145,054
Price per Sq. Foot	180	179	1.24	23,088	38,145,054
log(Price per Sq. Foot)	4.90	0.80	0.21	10.05	38,145,054
Property Characteristics					
Bedrooms	3.1	0.9	1.0	7.0	38,145,054
Bathrooms	2.1	0.8	0.25	7.0	38,145,054
Sq. Feet	1,948	1,065	420	10,228	38,145,054
log(Sq. Feet)	7.6	0.4	6.0	9.2	38,145,054
Year Built	1976	29	0.00	2018	38,145,054
State Characteristics					
GDP (Millions \$)	787,941	706,135	36,281	2,968,117	38,145,054
Population	15,535,151	12,272,350	567,136	39,557,045	38,145,054
Land (Acres)	77,264	50,066	61	261,797	38,145,054
Density	2.83	6.04	0.19	114.41	38,145,054
log(GDP)	13.18	0.92	10.50	14.90	38,145,054
log(Population)	16.21	0.89	13.25	17.49	38,145,054
log(Land)	10.96	0.96	4.12	12.48	38,145,054
Treatment Indicators					
Recreational Vote	0.07	0.25	0	1	38,145,054
Recreational Possession	0.06	0.25	0	1	38,145,054
Dispensary Date	0.04	0.20	0	1	38,145,054
First Dispensary	0.04	0.20	0	1	38,145,054
Medical	0.45	0.50	0	1	38,145,054

Housing variables are at the individual property transaction level ist , where i is a single property in state s . t reflects the date of transaction. The Price and Price per Sq. Foot variables represent unique transaction prices and are deflated using the 2018 Consumer Price Survey. The home characteristics Bedrooms, Bathrooms, Sq. Feet, and Year Built are unique to a given property but not necessarily unique to the dataset if a given property was sold more than once during the sample period. State characteristic variables are yearly at the state level s . GDP is the gross domestic product in a given year, Population is the state's total population, Land is the total land area of state s in acres, and Density is Population divided by Land which represents how concentrated a state's population is geographically. Treatment indicators are those indicators described in Section 3.2.2.

Marijuana Legalization Laws

State	Vote	Possession	Dispensary Date	First Dispensary	Medical
Alaska	Nov 4, 2014	Feb 24, 2015	Feb 24, 2015	Oct 31, 2016	Mar 4, 1999
Arizona					Nov 2, 2010
Arkansas					Nov 9, 2016
California	Nov 8, 2016	Nov 9, 2016	Jan 1, 2018	Jan 1, 2018	Nov 6, 1996
Colorado	Nov 6, 2012	Dec 6, 2012	Jan 1, 2014	Jan 1, 2014	Jun 1, 2001
Connecticut					May 31, 2012
Delaware					Jul 1, 2011
Florida					Jan 3, 2017
Hawaii					Jun 14, 2000
Illinois	Jun 25, 2019	Jan 1, 2020	Jan 1, 2020	Jan 1, 2020	Jan 1, 2014
Louisiana					1978
Maine	Nov 8, 2016	Jan 30, 2017	May 2, 2018	Spring 2020 (Expected)	Dec 22, 1999
Maryland					Jun 1, 2014
Massachusetts	Nov 8, 2016	Dec 15, 2016	Jul 1, 2018	Nov 20, 2018	Jan 1, 2013
Michigan	Nov 6, 2018	Dec 6, 2018	Dec 1, 2019	Dec. 1, 2019	Dec 4, 2008
Minnesota					May 30, 2014
Missouri					Dec 6, 2018
Montana					Nov 2, 2004
Nevada	Nov 8, 2016	Jan 1, 2017	Jan 1, 2017	Jul 1, 2017	Oct 1, 2001
New Hampshire					Jul 23, 2013
New Jersey					Jul 1, 2010
New Mexico					Jul 1, 2007
New York					Jul 5, 2014
North Dakota					Apr 18, 2017
Ohio					Sep 8, 2016
Oklahoma					Jul 26, 2018
Oregon	Nov 4, 2014	Jul 1, 2015	Oct 1, 2015	Oct 1, 2015	Dec 3, 1998
Pennsylvania					May 17, 2016
Rhode Island					Jan 3, 2006
Utah					Dec 1, 2018
Vermont	Jan 22, 2018	Jul 1, 2018			Jul 1, 2004
Washington	Nov 6, 2012	Dec 6, 2012	Jul 8, 2014	Jul 8, 2014	Nov 3, 1998
Washington DC	Nov 4, 2014	Feb 26, 2015			Jun 20, 2010
West Virginia					Jul 1, 2018
Total	12	12	10	10	34

Note: Vermont and Washington D.C. have passed laws allowing for the possession and cultivation of recreational marijuana, but have yet to allow for sales at retail locations as of this writing in February 2020. The data was derived from legislative and ballot acts, which are compiled nationally at the Marijuana Policy Project – <https://www.mpp.org/>

Effect of Marijuana Legalization on House Price per Sq. Foot (OLS)

	log (House Price)				
	(1)	(2)	(3)	(4)	(5)
Medical	0.414*** (0.035)				0.409*** (0.034)
Recreational Vote		0.180*** (0.029)			0.110*** (0.020)
Recreational Possession			0.186*** (0.029)		
Dispensary Date				0.152*** (0.035)	-0.024 (0.029)
R-squared	0.322	0.281	0.281	0.280	0.323
Observations	38,145,054				

Note: (i) The Possession dummy is excluded in the main column (5) since the time gap between Recreational Vote and Possession or Possession and the Dispensary Date are typically quite small. (ii) Both house characteristics – which includes bedrooms, bathrooms, the year built – and state characteristics such as state per capita GDP and density are controlled for in each model. (iii) City level clustered standard errors in parenthesis to take into account potential correlation in the error terms. (iv) As a robustness check we use house price per square foot as the dependent variable, which can be seen in Table A1 in Appendix B.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Effect of Marijuana Legalization on Home Price (Fixed Effects)

	log (Price)				
	(1)	(2)	(3)	(4)	(5)
Medical	0.039** (0.019)				0.061*** (0.020)
Recreational Vote		0.106*** (0.014)			0.054*** (0.013)
Recreational Possession			0.107*** (0.014)		
Dispensary Date				0.138*** (0.015)	0.111*** (0.010)
City FE	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.610	0.610	0.610	0.610	0.611
Observations	38,144,444				

Note: All models include city and year-quarter fixed effects. Beside our typical house characteristic controls (number of bedrooms, bathrooms, age), we also include local economic indicators at the state level. These include per capita GDP and population density. City level clustered standard errors are in parentheses to account for potential correlation in the error terms. As a robustness check we use house price per square foot as the dependent variable, which can be seen in Table A1 in Appendix B.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$.

Summary Statistics by Spatial Difference-in-Difference Treatment

	Inside 0.5 Miles		Between 0.5 and 1 Mile		Between 1 and 2 Miles		Outside 2 Miles	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
House Price (\$)	413,997	373,542	364,451	338,805	378,964	355,227	368,097	296,083
Price per Sq. Foot (\$)	255	193	214	149	198	132	186	128
Sq. Feet	1,711	873	1,723	846	1,941	963	2,046	964
Bedrooms	2.9	0.9	3.0	0.9	3.1	0.9	3.1	0.9
Bathrooms	2.0	0.9	2.1	0.8	2.3	0.8	2.4	0.9
Age of House at Sale (Years)	41.5	31.3	43.3	28.9	33.3	24.7	27.1	22.8
Observations	382,937		134,337		150,123		218,436	

The sample for the spatial difference-in-differences (SDD) model includes all home transactions in Colorado and Washington from 2014-2018. Each grouping represents the distance a home is from a dispensary, so for example homes in the first group are less than a half mile away from the nearest dispensary.

Spatial Difference-in-Differences

	log (Price)		
	(1)	(2)	(3)
1/2 Mile Zone	0.045*** (0.002)	0.072*** (0.002)	0.082*** (0.002)
1 Mile Zone	-0.028*** (0.002)	0.009*** (0.002)	
2 Mile Zone	-0.034*** (0.002)		
Control Group	Outside 2 Miles	Within 2 Miles	Within 1 Mile
Observations	885,833	650,437	565,923
R-squared	0.406	0.425	0.431

The sample includes transactions in the period between 2014 and 2018 in Colorado and Washington. Logged county level data such as county number of employees, wage, and the county employment ratio (county employees/state total employees), as well as home characteristics including the number of bedrooms, the square value of bedrooms, the age of the home, the number of bathrooms, and the square footage of the home, are used in the regression to control for differences across the states. Column 1 is an OLS model where treatment homes are homes that fall within 2 miles or closer of a dispensary and control homes are home that are not within 2 miles of a dispensary. Column 2 is the spatial difference in difference model where the control group becomes all homes that fall within 2 miles of a dispensary and the treatment group are homes that start off within 2 miles of a location and move within .5 or 1 mile of a dispensary. Column 3 is the same but now control are home starting off 1 mile and moving within .5 miles of a dispensary. Robust standard errors in parenthesis.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Figures




Figure 3.1: Trend of Average House Price Per Sq. Foot (CO, WA)

Note: (i) Control states are divided into three groups – high, middle, and low – based on their average home price per square foot. The low group is composed of Alabama, Florida, New Hampshire, Rhode Island, South Carolina, Tennessee, Texas, and West Virginia. The middle group consists of Georgia, Iowa, Kentucky, Mississippi, Montana, Nebraska, North Carolina, and Pennsylvania. The high group is made of Connecticut, Washington D.C., Delaware, Illinois, Minnesota, New Jersey, Virginia, and Wisconsin. (ii) The vertical line reflects the recreational marijuana legalization date for Colorado and Washington, 2012.




Figure 3.2: Trend of Average House Price Per Sq. Foot (OR)

Note: The control grouping is the same as in Figure 1. The vertical line reflecting the RML treatment date is 2014 for Oregon.





Figure 3.3: Trend of Average House Price Per Sq. Foot (CA, MA, NV)

Note: The control grouping is the same as in Figures 1 and 2. The vertical line reflecting the RML treatment date is 2016 for California, Massachusetts, and Nevada.



Legend

- | | | |
|-------------------|----------------------|-------------------|
| ◆ 2014 Dispensary | — 2014 1 Mile Zone | ■ Control Homes |
| ◆ 2015 Dispensary | — 2015 1/2 Mile Zone | ■ Treatment Homes |

Figure 3.4: Illustration of Spatial Difference in Difference Model in Denver, Colorado

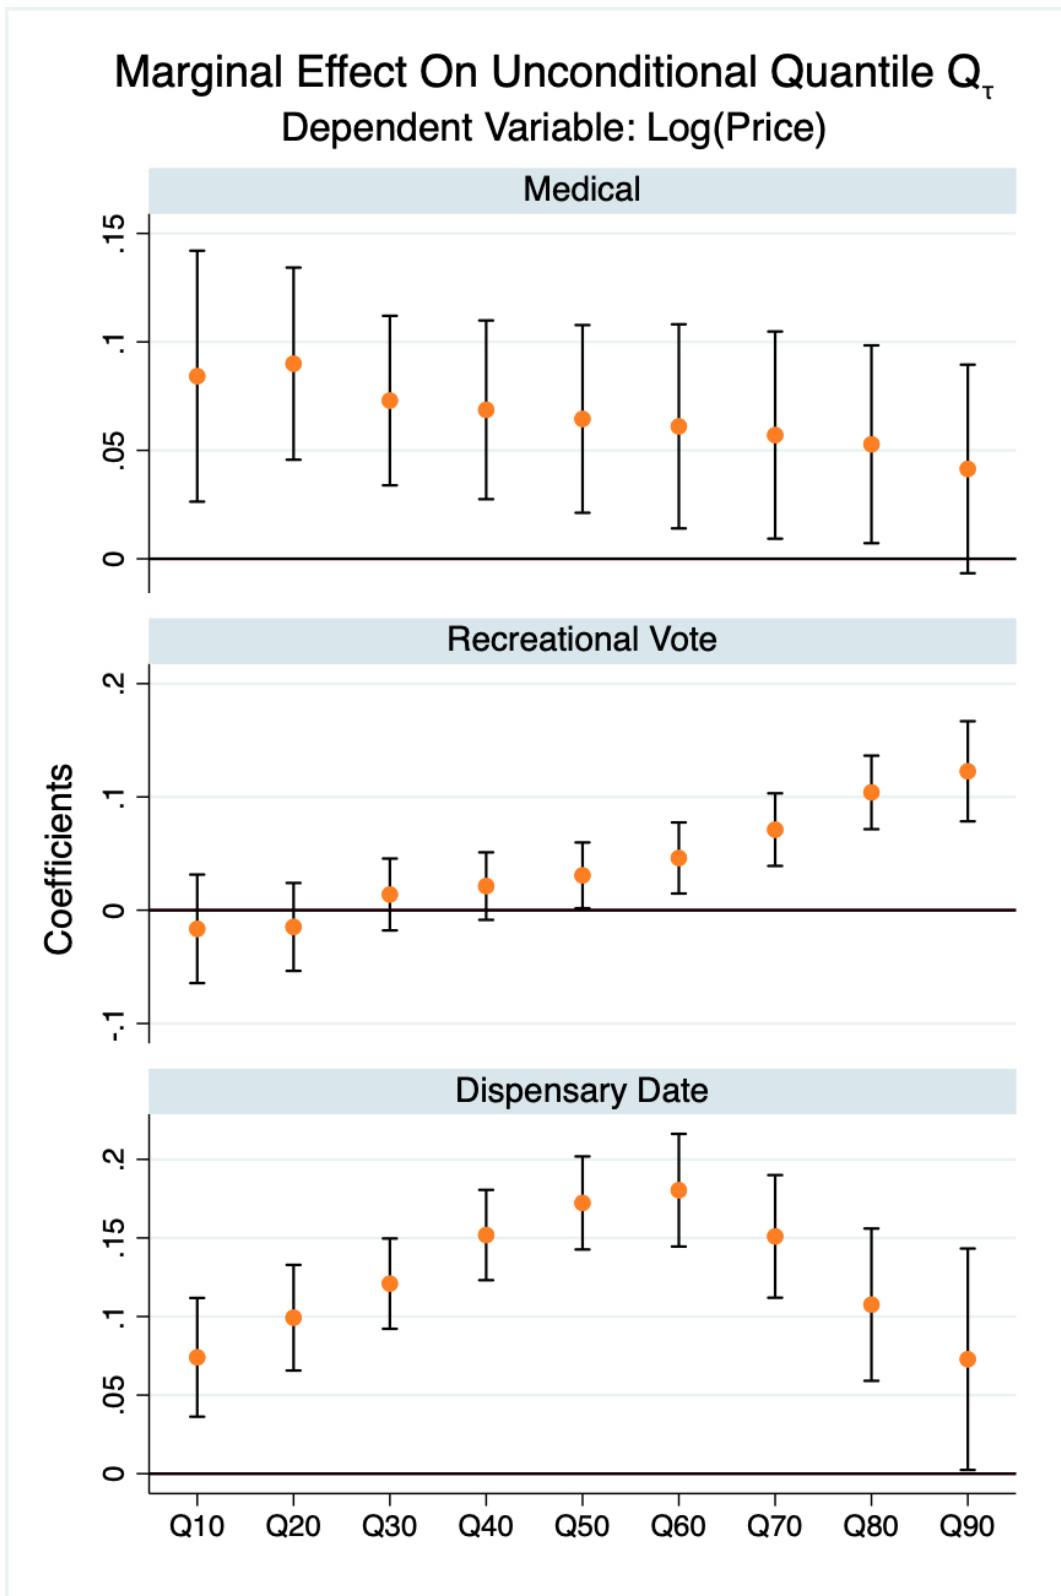


Figure 3.5




Figure 3.6

Bibliography

- Robert Mendelsohn, William D. Nordhaus, and Daigee Shaw. The Impact of Global Warming on Agriculture: A Ricardian Analysis. American Economic Review, 84(4):753–771, 1994.
- Olivier Deschênes and Michael Greenstone. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us. American Economic Journal: Applied Economics, 3(4):152–185, 2011.
- Garth Heutel, Nolan H Miller, and David Molitor. Adaptation and the mortality effects of temperature across us climate regions. Technical report, National Bureau of Economic Research, 2017.
- Alan Barreca, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. Journal of Political Economy, 124(1):105–159, 2016.
- Lucas W Davis and Paul J Gertler. Contribution of air conditioning adoption to future energy use under global warming. Proceedings of the National Academy of Sciences, 112(19):5962–5967, 2015.
- Erin T. Mansur, Robert Mendelsohn, and Wendy Morrison. Climate change adaptation: A study of fuel choice and consumption in the US energy sector. Journal of Environmental Economics and Management, 55(2):175–193, 2008. ISSN 00950696. doi: 10.1016/j.jeem.2007.10.001.
- Maximilian Auffhammer. Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data. 2018.
- Achyuta Adhvaryu, Namrata Kalaz, and Nyshadham Nyshadham. The Light and the Heat: Pro-

- ductivity Co-Benefits of Energy-Saving Technology. The Review of Economics and Statistics, pages 1–36, 2019. ISSN 0034-6535. doi: 10.1162/rest_a_00886.
- Jonathan Colmer. Temperature, Labor Reallocation, and Industrial Production: Evidence from India. Under Review, 2019.
- Joshua Graff Zivin and Matthew Neidell. Temperature and the allocation of time: Implications for climate change. Journal of Labor Economics, 32(1):1–26, 2014. ISSN 0734306X. doi: 10.1086/671766.
- Daniel Kahneman and Alan B Krueger. Developments in the measurement of subjective well-being. Journal of Economic perspectives, 20(1):3–24, 2006.
- Nathan W Chan and Casey J Wichman. Climate change and recreation: Evidence from north american cycling. Environmental and Resource Economics, pages 1–33, 2020.
- Edward Arens and Peter Bosselmann. Wind, sun and temperature—predicting the thermal comfort of people in outdoor spaces. Building and Environment, 24(4):315–320, 1989.
- Peter Höppe. Different aspects of assessing indoor and outdoor thermal comfort. Energy and buildings, 34(6):661–665, 2002.
- Theodore Stathopoulos, Hanqing Wu, and John Zacharias. Outdoor human comfort in an urban climate. Building and Environment, 39(3):297–305, 2004.
- Reuben Gronau. The allocation of time of israeli women. Journal of Political Economy, 84(4, Part 2):S201–S220, 1976a.
- Reuben Gronau. Economic approach to value of time and transportation choice. Transportation Research Record, 587(1), 1976b.
- Gary S Becker. A theory of the allocation of time. The Economic Journal, pages 493–517, 1965.
- M Bruce Johnson. Travel time and the price of leisure. Economic Inquiry, 4(2):135, 1966.
- Maximilian Auffhammer and Wolfram Schlenker. Empirical studies on agricultural impacts and adaptation. Energy Economics, 46:555–561, 2014. ISSN 01409883. doi: 10.1016/j.eneco.2014.09.010. URL <http://dx.doi.org/10.1016/j.eneco.2014.09.010>.
- Ethan E Butler and Peter Huybers. Adaptation of us maize to temperature variations. Nature

Climate Change, 3(1):68–72, 2013.

John T Abatzoglou. Development of gridded surface meteorological data for ecological applications and modelling. International Journal of Climatology, 33(1):121–131, 2013.

Michael Dorman, Johnathan Rush, Ian Hough, Dominic Russel, and Charles F.F. Karney. ngeo: k-Nearest Neighbor Join for Spatial Data, 2020. URL <https://cran.r-project.org/package=ngeo>. R package version 0.3.7.

Jeffrey Hollister and Tarak Shah Shah. elevatr: Access Elevation Data from Various APIs, 2020. URL <https://cran.r-project.org/package=elevatr>. R package version 0.2.0.

Wolfram Schlenker and Michael J Roberts. Nonlinear temperature effects indicate severe damages to us crop yields under climate change. Proceedings of the National Academy of sciences, 106(37):15594–15598, 2009.

Anin Aroonruengsawat and Maximilian Auffhammer. Impacts of climate change on residential electricity consumption: evidence from billing data. In The Economics of climate change: Adaptations past and present, pages 311–342. University of Chicago Press, 2011.

BM Sanderson and MF Wehner. Model weighting strategy. Climate Science Special Report: Fourth National Climate Assessment, 1:436–442, 2017.

Keywan Riahi, V Krey, S Rao, V Chirkov, G Fischer, P Kolp, G Kindermann, N Nakicenovic, and P Rafai. Rcp-8.5: exploring the consequence of high emission trajectories. Climatic Change. doi, 10:1007, 2011.

Gerald A Meehl, Catherine A Senior, Veronika Eyring, Gregory Flato, Jean-Francois Lamarque, Ronald J Stouffer, Karl E Taylor, and Manuel Schlund. Context for interpreting equilibrium climate sensitivity and transient climate response from the cmip6 earth system models. Science Advances, 6(26):eaba1981, 2020.

Maximilian Auffhammer and Erin T. Mansur. Measuring climatic impacts on energy consumption: A review of the empirical literature. Energy Economics, 46:522–530, 2014. ISSN 01409883. doi: 10.1016/j.eneco.2014.04.017. URL <http://dx.doi.org/10.1016/j.eneco.2014.04.017>.

- Marshall Burke, Solomon M. Hsiang, and Edward Miguel. Climate and Conflict. Annual Review of Economics, 7:577–617, 2015. ISSN 1943345X. doi: 10.1146/annurev-economics-080614-115430.
- Solomon M Hsiang. Climate Econometrics. Annual Review of Resource Economics, 8(1):43–75, 2016.
- Emanuele Massetti and Robert Mendelsohn. Measuring climate adaptation: Methods and evidence. Review of Environmental Economics and Policy, 12(2):324–341, 2018.
- Diana E Bowler, Lisette M Buyung-Ali, Teri M Knight, and Andrew S Pullin. A systematic review of evidence for the added benefits to health of exposure to natural environments. BMC public health, 10(1):456, 2010.
- David G Pearson and Tony Craig. The great outdoors? exploring the mental health benefits of natural environments. Frontiers in psychology, 5:1178, 2014.
- Caoimhe Twohig-Bennett and Andy Jones. The health benefits of the great outdoors: A systematic review and meta-analysis of greenspace exposure and health outcomes. Environmental research, 166:628–637, 2018.
- Mark Aguiar and Erik Hurst. Measuring trends in leisure: The allocation of time over five decades. Quarterly Journal of Economics, 122(3):969–1006, 2007. ISSN 00335533. doi: 10.1162/qjec.122.3.969.
- Valerie A Ramey and Neville Francis. A century of work and leisure. American Economic Journal: Macroeconomics, 1(2):189–224, 2009.
- Alan B. Krueger and Andreas I. Mueller. Time use, emotional well-being, and unemployment: Evidence from longitudinal data. American Economic Review, 102(3):594–599, 2012. ISSN 00028282. doi: 10.1257/aer.102.3.594.
- Dean Knox, Will Lowe, and Jonathan Mummolo. Administrative Records Mask Racially Biased Policing. American Political Science Review, 114(3):619–637, August 2020. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055420000039. URL https://www.cambridge.org/core/product/identifier/S0003055420000039/type/journal_article.

- John M MacDonald and Jeffrey Fagan. Using shifts in deployment and operations to test for racial bias in police stops. In AEA Papers and Proceedings, volume 109, pages 148–51, 2019.
- Felipe Goncalves and Steven Mello. A few bad apples? racial bias in policing. Racial Bias in Policing (June 15, 2020), 2020.
- Melissa LoPalo. Temperature, worker productivity, and adaptation: evidence from survey data production. Available at SSRN, 2020.
- Matthew E Kahn and Pei Li. The effect of pollution and heat on high skill public sector worker productivity in china. Technical report, National Bureau of Economic Research, 2019.
- Dwight A Hennessy and David L Wiesenthal. Traffic congestion, driver stress, and driver aggression. Aggressive Behavior: Official Journal of the International Society for Research on Aggression, 25(6):409–423, 1999.
- Alan S Reifman, Richard P Larrick, and Steven Fein. Temper and temperature on the diamond: The heat-aggression relationship in major league baseball. Personality and Social Psychology Bulletin, 17(5):580–585, 1991.
- Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, et al. A large-scale analysis of racial disparities in police stops across the united states. Nature human behaviour, 4(7):736–745, 2020.
- Paul M Collier. Police performance measurement and human rights. Public Money and Management, 21(3):35–39, 2001.
- Nigel Fielding and Martin Innes. Reassurance policing, community policing and measuring police performance. Policing & society, 16(02):127–145, 2006.
- Jeffrey Grogger and Greg Ridgeway. Testing for racial profiling in traffic stops from behind a veil of darkness. Journal of the American Statistical Association, 101(475):878–887, 2006.
- Douglas T Kenrick and Steven W MacFarlane. Ambient temperature and horn honking: A field study of the heat/aggression relationship. Environment and behavior, 18(2):179–191, 1986.
- Daniel S Hamermesh et al. Workdays, workhours, and work schedules: Evidence for the united

- states and germany. Books from Upjohn Press, 1996.
- John Lamberth. Measuring the racial/ethnic make up of traffic: The how, what and why. In workshop titled, "Confronting Racial Profiling in the 21st Century: Implications for Racial Justice," Boston, Mass.: Institute on Race and Justice, 2003.
- Sébastien Annan-Phan and Bocar A. Ba. Hot Temperatures, Aggression, and Death at the Hands of the Police: Evidence from the U.S. SSRN Electronic Journal, 2019. ISSN 1556-5068. doi: 10.2139/ssrn.3356045. URL <https://www.ssrn.com/abstract=3356045>.
- Peng Zhang, Olivier Deschenes, Kyle Meng, and Junjie Zhang. Temperature effects on productivity and factor reallocation: Evidence from a half million chinese manufacturing plants. Journal of Environmental Economics and Management, 88:1–17, 2018.
- Achyuta Adhvaryu, Namrata Kala, and Anant Nyshadham. The light and the heat: Productivity co-benefits of energy-saving technology. Technical report, National Bureau of Economic Research, 2018.
- Olli Seppanen, William J Fisk, and QH Lei. Effect of temperature on task performance in office environment. 2006.
- Tord Kjellstrom, David Briggs, Chris Freyberg, Bruno Lemke, Matthias Otto, and Olivia Hyatt. Heat, human performance, and occupational health: a key issue for the assessment of global climate change impacts. Annual review of public health, 37:97–112, 2016.
- Jisung Park. Hot temperature and high stakes exams: Evidence from new york city public schools. URL: <https://scholar.harvard.edu/files/jisungpark/files/papernycaejep.pdf>, 2017.
- Matthew Ranson. Crime, weather, and climate change. Journal of Environmental Economics and Management, 67(3):274–302, 2014.
- Murat Iyigun, Nathan Nunn, and Nancy Qian. Winter is coming: The long-run effects of climate change on conflict, 1400-1900. Technical report, National Bureau of Economic Research, 2017.
- Craig A Anderson, Kathryn B Anderson, and William E Deuser. Examining an affective aggression framework weapon and temperature effects on aggressive thoughts, affect, and attitudes. Personality and social psychology Bulletin, 22(4):366–376, 1996.

- Robert A Baron. The reduction of human aggression: A field study of the influence of incompatible reactions 1. Journal of Applied Social Psychology, 6(3):260–274, 1976.
- Richard P Larrick, Thomas A Timmerman, Andrew M Carton, and Jason Abrevaya. Temper, temperature, and temptation: Heat-related retaliation in baseball. Psychological Science, 22(4): 423–428, 2011.
- Robert J. Hijmans. terra: Spatial Data Analysis, 2021. URL <https://CRAN.R-project.org/package=terra>. R package version 1.1-4.
- Roland Stull. Wet-bulb temperature from relative humidity and air temperature. Journal of applied meteorology and climatology, 50(11):2267–2269, 2011.
- Casey J Wichman. Interpreting nonlinear semi-elasticities in reduced-form climate damage estimation. Climatic Change, 148(4):641–648, 2018.
- Greg Ridgeway. Cincinnati police department traffic stops: Applying RAND’s framework to analyze racial disparities, volume 914. Rand Corporation, 2009.
- Danna Thomas and Lin Tian. Hits from the Bong: The Impact of Recreational Marijuana Dispensaries on Property Values. 2017.
- James Conklin, Moussa Diop, and Herman Li. Contact high: The external effects of retail marijuana establishments on house prices. Real Estate Economics, 48(1):135–173, 2020.
- Justin Tyndall. Getting high and low prices: Marijuana dispensaries and home values. Real Estate Economics, 2019.
- Jesse Burkhardt and Matthew Flyr. The effect of marijuana dispensary openings on housing prices. Contemporary Economic Policy, 37(3):462–475, 2019.
- Cheng Cheng, Walter J. Mayer, and Yanling Mayer. The Effect of Legalizing Retail Marijuana on Housing Values: Evidence From Colorado. Economic Inquiry, 56(3):1585–1601, 2018. ISSN 14657295. doi: 10.1111/ecin.12556.
- Salomeh Keyhani, Stacey Steigerwald, Julie Ishida, Marzieh Vali, Magdalena Cerdá, Deborah Hasin, Camille Dollinger, Sodahm R. Yoo, and Beth E. Cohen. Risks and benefits of marijuana use a national survey of U.S. Adults. Annals of Internal Medicine, 2018. ISSN 15393704. doi:

10.7326/M18-0810.

Maria L. Roditis and Bonnie Halpern-Felsher. Adolescents' Perceptions of Risks and Benefits of Conventional Cigarettes, E-cigarettes, and Marijuana: A Qualitative Analysis. Journal of Adolescent Health, 2015. ISSN 18791972. doi: 10.1016/j.jadohealth.2015.04.002.

E. Kondrad and A. Reid. Colorado Family Physicians' Attitudes Toward Medical Marijuana. The Journal of the American Board of Family Medicine, 2013. ISSN 1557-2625. doi: 10.3122/jabfm.2013.01.120089.

Beatriz H. Carlini, Sharon B. Garrett, and Gregory T. Carter. Medicinal Cannabis: A Survey Among Health Care Providers in Washington State. American Journal of Hospice and Palliative Medicine, 2017. ISSN 10499091. doi: 10.1177/1049909115604669.

Mary Ann Fitzcharles, Peter A. Ste-Marie, Daniel J. Clauw, Shahin Jamal, Jacob Karsh, Sharon Leclercq, Jason J. McDougall, Yoram Shir, Kam Shojania, and Zach Walsh. Rheumatologists lack confidence in their knowledge of cannabinoids pertaining to the management of rheumatic complaints. BMC Musculoskeletal Disorders, 2014. ISSN 14712474. doi: 10.1186/1471-2474-15-258.

Ilana M. Braun, Alexi Wright, John Peteet, Fremonta L. Meyer, David P. Yuppa, Dragana Bolcic-Jankovic, Jessica LeBlanc, Yuchiaio Chang, Liyang Yu, Manan M. Nayak, James A. Tulskey, Joji Suzuki, Lida Nabati, and Eric G. Campbell. Medical oncologists' beliefs, practices, and knowledge regarding marijuana used therapeutically: A nationally representative survey study. Journal of Clinical Oncology, 2018. ISSN 15277755. doi: 10.1200/JCO.2017.76.1221.

Jaren C. Pope. Fear of crime and housing prices: Household reactions to sex offender registries. Journal of Urban Economics, 64(3):601–614, 2008. ISSN 00941190. doi: 10.1016/j.jue.2008.07.001.

Paolo Buonanno, Daniel Montolio, and Josep Maria Raya-Vílchez. Housing prices and crime perception. Empirical Economics, 45(1):305–321, 2013. ISSN 03777332. doi: 10.1007/s00181-012-0624-y.

Jeffrey Brinkman and David Mok-Lamme. Not in my backyard? not so fast. the effect of marijuana

- legalization on neighborhood crime. Regional Science and Urban Economics, 78:103460, 2019.
- Robert G. Morris, Michael TenEyck, J. C. Barnes, and Tomislav V. Kovandzic. The effect of medical marijuana laws on crime: Evidence from state panel data, 1990-2006. PLoS ONE, 9 (3), 2014. ISSN 19326203. doi: 10.1371/journal.pone.0092816.
- Arthur Huber, Rebecca Newman, and Daniel LaFave. Cannabis control and crime: Medicinal use, depenalization and the war on drugs. B.E. Journal of Economic Analysis and Policy, 16(4), 2016. ISSN 19351682. doi: 10.1515/bejeap-2015-0167.
- David A Makin, Dale W Willits, Guangzhen Wu, Kathryn O DuBois, Ruibin Lu, Mary K Stohr, Wendy Koslicki, Duane Stanton, Craig Hemmens, John Snyder, et al. Marijuana legalization and crime clearance rates: Testing proponent assertions in colorado and washington state. Police quarterly, 22(1):31–55, 2019.
- Benjamin Hansen, Keaton Miller, and Caroline Weber. Early evidence on recreational marijuana legalization and traffic fatalities. Economic inquiry, 58(2):547–568, 2020.
- Bradley J Bartos, Carol Newark, and Richard McCleary. Marijuana medicalization and motor vehicle fatalities: a synthetic control group approach. Journal of Experimental Criminology, pages 1–18, 2018.
- Joseph J. Sabia and Thanh Tam Nguyen. The Effect of Medical Marijuana Laws on Labor Market Outcomes. The Journal of Law and Economics, 61(3):361–396, 2018. ISSN 0022-2186. doi: 10.1086/701193.
- Lauren Hersch Nicholas and Johanna Catherine Maclean. The Effect of Medical Marijuana Laws on the Health and Labor Supply of Older Adults: Evidence from the Health and Retirement Study. Journal of Policy Analysis and Management, 38(2):455–480, 2019. ISSN 15206688. doi: 10.1002/pam.22122.
- Diego Zambiasi and Steven Stillman. The pot rush: Is legalized marijuana a positive local amenity? Economic Inquiry, 58(2):667–679, 2020.
- Kathrin Degen and Andreas M Fischer. Immigration and Swiss House Prices. Swiss Journal of Economics and Statistics, 163(1):15–36, 2017. ISSN 00318256. doi: 10.1016/j.tetlet.2008.08.

010. URL <http://ideas.repec.org/p/cpr/ceprdp/7583.html>.
- Filipa Sá. Immigration and House Prices in the UK. Economic Journal, 125(587):1393–1424, 2015. ISSN 14680297. doi: 10.1111/eoj.12158.
- David Card and Alan B. Krueger. School Resources and Student Outcomes. The Annals of the American Academy of Political and Social Science, 559(1):39–53, 1998. ISSN 1059-6011. doi: 0803973233.
- C. Kirabo Jackson, Rucker C. Johnson, and Claudia Persico. The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms. Quarterly Journal of Economics, 133(1):157–218, 2016. doi: 10.1093/qje/qjv036.Advance.
- Paco Martorell, Kevin Stange, and Isaac McFarlin. Investing in schools: capital spending, facility conditions, and student achievement. Journal of Public Economics, 140:13–29, 2016. ISSN 00472727. doi: 10.1016/j.jpubeco.2016.05.002. URL <http://dx.doi.org/10.1016/j.jpubeco.2016.05.002>.
- Stephanie Riegg Cellini, Fernando Ferreira, and Jesse Rothstein. The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design. The Quarterly Journal of Economics, 125(1):215–261, 2010.
- Christopher A. Neilson and Seth D. Zimmerman. The effect of school construction on test scores, school enrollment, and home prices. Journal of Public Economics, 120:18–31, 2014. ISSN 00472727. doi: 10.1016/j.jpubeco.2014.08.002. URL <http://dx.doi.org/10.1016/j.jpubeco.2014.08.002>.
- Sergio Firpo, Nicole M Fortin, and Thomas Lemieux. Unconditional Quantile Regressions. Econometrica, 77(3):953–973, 2009. ISSN 0012-9682. doi: 10.3982/ecta6822.
- Wallace E Oates. The Effects of Property Taxes and Local Public Spending on Property Values : An Empirical Study of Tax Capitalization and the Tiebout Hypothesis Author (s): Wallace E . Oates Source : Journal of Political Economy , Vol . 77 , No . 6 (Nov . - Dec . , 1969 . Journal of Political Economy , 77(6):957–971, 1969.
- John E. Anderson. Property taxes and the timing of urban land development. Regional Science

- and Urban Economics, 16(4):483–492, 1986. ISSN 01660462. doi: 10.1016/0166-0462(86)90019-0.
- Trey Dronyk-Trosper. Searching for Goldilocks: The Distance-Based Capitalization Effects of Local Public Services. Real Estate Economics, 45(3):650–678, 2017. ISSN 15406229. doi: 10.1111/1540-6229.12171.
- Rosalie Pacula, Jamie Chriqui, and Joanna King. Marijuana Decriminalization: What does it Mean in the United States? 2003. URL <http://www.nber.org/papers/w9690>.
- Andrew Henley. Residential Mobility , Housing Equity and the Labour Market. The Economic Journal, 108(447):414–427, 1998.
- Andrew Foote. The effects of negative house price changes on migration: Evidence across U.S. housing downturns. Regional Science and Urban Economics, 60:292–299, 2016. ISSN 18792308. doi: 10.1016/j.regsciurbeco.2016.08.001. URL <http://dx.doi.org/10.1016/j.regsciurbeco.2016.08.001>.
- Jeffrey E. Zabel. Migration, housing market, and labor market responses to employment shocks. Journal of Urban Economics, 72(2-3):267–284, 2012. ISSN 00941190. doi: 10.1016/j.jue.2012.05.006. URL <http://dx.doi.org/10.1016/j.jue.2012.05.006>.
- Gizem Koşar, Tyler Ransom, and H. Wilbert van der Klaauw. Understanding Migration Aversion Using Elicited Counterfactual Choice Probabilities. Federal Reserve Bank of New York Staff Reports, 883, 2019. doi: 10.2139/ssrn.3368408.
- Yong Chen and Stuart S. Rosenthal. Local amenities and life-cycle migration: Do people move for jobs or fun? Journal of Urban Economics, 64(3):519–537, 2008. ISSN 00941190. doi: 10.1016/j.jue.2008.05.005. URL <http://dx.doi.org/10.1016/j.jue.2008.05.005>.
- Rui Baptista and Miguel Torres Preto. New firm formation and employment growth: Regional and business dynamics. Small Business Economics, 36(4):419–442, 2011. ISSN 0921898X. doi: 10.1007/s11187-009-9254-y.
- Martin Andersson and Florian Noseleit. Start-ups and employment dynamics within and across sectors. Small Business Economics, 36(4):461–483, 2011. ISSN 0921898X. doi: 10.1007/

s11187-009-9252-0.

Zheng Liu, Jianjun Miao, and Tao Zha. Land prices and unemployment. Journal of Monetary Economics, 80:86–105, 2016. ISSN 03043932. doi: 10.1016/j.jmoneco.2016.05.001. URL <http://dx.doi.org/10.1016/j.jmoneco.2016.05.001>.

Alan K. Reichert. The impact of interest rates, income, and employment upon regional housing prices. The Journal of Real Estate Finance and Economics, 3(4):373–391, 1990. ISSN 08955638. doi: 10.1007/BF00178859.

William T Bogart. How Much Is a Neighborhood School Worth? Journal of Urban Economics, 47:280–305, 2000.

Paul Cheshire and Stephen Sheppard. Capitalising the value of free schools: The impact of supply characteristics and uncertainty. Economic Journal, 114(499), 2004. ISSN 00130133. doi: 10.1111/j.1468-0297.2004.00252.x.

Vladimir Bajic. The Effects of a New Subway Line on Housing Prices in Metropolitan Toronto. Urban Studies, 20(2):147–158, 1983. ISSN 1360063X. doi: 10.1080/00420988320080291.

D. N. Dewees. The effect of a subway on residential property values in Toronto. Journal of Urban Economics, 3(4):357–369, 1976. ISSN 00941190. doi: 10.1016/0094-1190(76)90035-8.

Donald J. Epp and K. S. Al-Ani. The Effect of Water Quality on Rural Nonfarm Residential Property Values. American Journal of Agricultural Economics, 61(3):529–534, 1979. ISSN 00029092. doi: 10.2307/1239441.

C. Edwin Young and Frank A. Teti. The Influence of Water Quality on the Value of Recreational Properties Adjacent to St. Albans Bay, Vermont. Technical report, United States Department of Agriculture, Economic Research Service, 1984.

Christopher G. Leggett and Nancy E. Bockstael. Evidence of the effects of water quality on residential land prices. Journal of Environmental Economics and Management, 39(2):121–144, 2000. ISSN 00950696. doi: 10.1006/jcem.1999.1096.

David R. Bowes and Keith R. Ihlanfeldt. Identifying the impacts of rail transit stations on residential property values. Journal of Urban Economics, 50(1):1–25, 2001. ISSN 00941190. doi:

10.1006/juec.2001.2214.

Stephen Gibbons and Stephen Machin. Valuing rail access using transport innovations. Journal of Urban Economics, 57(1):148–169, 2005. ISSN 00941190. doi: 10.1016/j.jue.2004.10.002.

Daniel P. McMillen and John McDonald. Reaction of house prices to a new rapid transit line: Chicago’s midway line, 1983-1999. Real Estate Economics, 32(3):463–486, 2004. ISSN 10808620. doi: 10.1111/j.1080-8620.2004.00099.x.

Daryl A. Hellman and Joel L. Naroff. The Impact of Crime on Urban Residential Property Values. Urban Studies, 16(1):105–112, 1979. ISSN 0042-0980. doi: 10.1080/713702454.

Stephen Gibbons and Stephen Machin. Valuing school quality, better transport, and lower crime: Evidence from house prices. Oxford Review of Economic Policy, 24(1):99–119, 2008. ISSN 0266903X. doi: 10.1093/oxrep/grn008.


Zillow. ZTRAX: Zillow Transaction and Assessor Dataset, 2017-Q4, 2017.

F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. Robust Statistics: The Approach Based on Influence Functions. 2005. doi: 10.2307/2987975.

Appendices


Appendix A: Time Use Robustness Checks

Figure .0A1: Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time with 3-Day Temperature Lags



Note: The solid orange line represents the median response, and each shaded green area represents the surrounding high and low decile. For example, the darkest shaded area is the estimated coefficients that are in the 40th to 60th percentile of the coefficient distribution. The coefficients are interpreted as the change in time outside as a result of being in a particular temperature bin relative to a day in the 60-70bin.

Figure .0A2: Distribution of Stage 1 Temperature Response Coefficients for Outdoor Time with No Fixed Effects



Note: The solid orange line represents the median response, and each shaded green area represents the surrounding high and low decile. For example, the darkest shaded area is the estimated coefficients that are in the 40th to 60th percentile of the coefficient distribution. The coefficients are interpreted as the change in time outside as a result of being in a particular temperature bin relative to a day in the 60-70bin.

Appendix B: Data Cleaning Description

Zillow Housing Data

Considering the size and scope of the Zillow ZTRAX repository, it is necessary to document the data cleaning process used for this research. However, in order to create a dataset that is both national and representative, some adjustments were made to the import process. In general, the effort follows Zillow's own script which creates a hedonic dataset.¹³ The end product results in a dataframe in which each row is a home transaction and each column reflects home and transaction characteristics. The files are initially imported state-by-state and then appended together to make a master file.

The process goes as follows. First, three tables are imported from the Assessment repository: Main, Building, and BuildingAreas. These three tables combine to provide house characteristics, as well as information about the type of property exchanged in a given transaction. For example, the variable "PropertyLandUseStdCode" in the Building table details whether a property is a single-family residence, used in industry, is a farm, et cetera. We erred on the side of inclusivity when filtering for these variables during import, as reporting standards across counties and states vary widely. The properties included are described as follows in Zillow's documentation:

1. Residential General
2. Single-Family Residences
3. Rural Residences
4. Mobile Home
5. Townhouse
6. Cluster Home
7. Condominium

8. Cooperative
9. Row House
10. Planned Unit Development
11. Residential Common Area
12. Seasonal, Cabin, Vacation Residence
13. Bungalow
14. Zero Lot Line
15. Manufactured, Modular, Prefabricated Homes
16. Patio Home
17. Garden Home
18. Condominium
19. Inferred Single-Family Residential

Also, following the logic described by Zillow, we filter the “BuildingAreaStdCode” from the BuildingAreas table in order to get as accurate a measure of total square footage as possible. Again, different counties have different reporting standards as to what is included in their square footage calculations, so to ensure consistency we have included only those options which enumerate the buildings on the property, not the land itself. These two filters – for “PropertyLandUseStdCode” and “BuildingAreaStdCode” – are the only two at this point in the process. Once this is complete, the three assessment tables are merged to create a single assessment file with all the necessary housing characteristic variables to be used in analysis.

The second set of data comes from the Transaction repository. Included are the PropertyInfo and Main tables. All the information provided here reflects the transaction itself, not any characteristics of the home. This includes variables like the price of the transaction, the date of transfer,

and the type of transfer. The only filtering that occurs in this step is in regard to the variable “DataClassStdCode,” which details the type of transaction occurring. Since the subject of study are property transactions, only deed transfers and deed transfers with concurrent mortgages are included. This excludes other types of transactions, including foreclosures and inter-family transfers as in the case of inheritances. These two tables are appended together to make a single transaction file. Finally, the transaction and assessment files are combined to make a single master file for a given state. The states files are then appended together to make a national-level dataset which is then used for analysis.

The master file is filtered to exclude extreme observations, as well as define the period of study. To ensure that results are not being driven by incorrect or implausible observations, we drop transactions which had sales prices below \$10,000 and above \$10,000,000, similar to Cheng et al. [2018]. On the lower end it is unlikely that transactions with prices below \$10,000 occurred on the market, and may have slipped through the “DataClassStdCode” filter. Prices above \$10,000,000 are extraordinary and in some cases are likely the result of data entry errors. Similarly, house characteristics are filtered to exclude observations that are in the top thousandth or top ten-thousandth percentile. Doing so, for example, eliminated an observation with over 1000 bedrooms. This process removed a large number of observations in states which do not require counties to report the home characteristics, leaving small states like Maine with just 11,000 transaction observations. To guarantee a representative sample, we then dropped states which did not have at least 100,000 observations. That is an arbitrary standard, but by doing so we can more confidently argue that each states’ market is properly represented. Finally, prices were adjusted to reflect 2018 prices using the Federal Reserve’s Consumer Price Index.

Appendix C: Additional Model Specifications

Effect of Recreational Marijuana Legalization on House Price per Sq. Foot

	log (Price per Sq. Foot)									
	OLS					Fixed Effects				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Medical	0.466*** (0.035)				0.461*** (0.035)	0.046** (0.020)				0.069*** (0.021)
Recreational Vote		0.187*** (0.030)			0.088*** (0.020)		0.108*** (0.014)			0.055*** (0.014)
Recreational Possession			0.193*** (0.031)					0.109*** (0.015)		
Dispensary Date				0.169*** (0.037)	0.001 (0.030)				0.141*** (0.015)	0.116*** (0.011)
Bedrooms	-0.434*** (0.035)	-0.458*** (0.036)	-0.458*** (0.036)	-0.459*** (0.036)	-0.433*** (0.036)	-0.102*** (0.020)	-0.102*** (0.020)	-0.102*** (0.020)	-0.102*** (0.020)	-0.102*** (0.020)
Bedrooms ²	0.044*** (0.004)	0.047*** (0.004)	0.047*** (0.004)	0.047*** (0.004)	0.044*** (0.004)	0.004* (0.002)	0.004* (0.002)	0.004* (0.002)	0.004* (0.002)	0.004* (0.002)
Bathrooms	0.188*** (0.011)	0.192*** (0.011)	0.192*** (0.011)	0.192*** (0.011)	0.187*** (0.011)	0.044*** (0.006)	0.044*** (0.006)	0.044*** (0.006)	0.045*** (0.006)	0.045*** (0.007)
Age	0.001** (0.000)	0.001 (0.000)	0.001 (0.000)	0.001 (0.000)	0.001** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)	-0.002*** (0.000)
Age ²	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
GDP Per Capita	0.356*** (0.053)	0.828*** (0.062)	0.826*** (0.062)	0.870*** (0.061)	0.314*** (0.059)	1.323*** (0.109)	1.273*** (0.114)	1.273*** (0.114)	1.302*** (0.113)	1.252*** (0.111)
Density	0.004*** (0.001)	-0.004 (0.003)	-0.003 (0.003)	-0.004 (0.003)	0.005*** (0.001)	-0.003 (0.004)	-0.002 (0.004)	-0.002 (0.004)	-0.002 (0.003)	-0.001 (0.003)
City FE	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes
City Clustered S.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.161	0.099	0.099	0.098	0.162	0.528	0.529	0.529	0.529	0.530
Observations	38,145,054	38,145,054	38,145,054	38,145,054	38,145,054	38,144,444	38,144,444	38,144,444	38,144,444	38,144,444

Note: (i) The dependent variable is the log of house price per square foot while the first half columns are OLS results and the latter half are FE results. (ii) Possession dummy is excluded in our main columns (5) and (10) since the time gap between vote and possession, or sale and possession are too small to capture significantly valuable variations. (iii) City level clustered standard errors in parenthesis to take into account the correlations of error terms.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Effect of Recreational Marijuana Legalization on House Price (Robustness)

	log (Price)				
	(1)	(2)	(3)	(4)	(5)
Medical	0.094*** (0.023)	0.101*** (0.023)	0.095*** (0.023)	0.102*** (0.023)	0.062*** (0.020)
Recreational Vote	0.168*** (0.013)	0.154*** (0.013)	0.161*** (0.013)	0.146*** (0.013)	0.052*** (0.013)
Dispensary Date	0.063*** (0.013)	0.074*** (0.013)			
First Dispensary			0.078*** (0.011)	0.091*** (0.011)	0.120*** (0.011)
Bedrooms		0.023 (0.021)		0.023 (0.021)	0.035* (0.020)
Bedrooms ²		-0.005** (0.002)		-0.005** (0.002)	-0.007*** (0.002)
Bathrooms		0.136*** (0.007)		0.136*** (0.007)	0.135*** (0.007)
log(Sq. Feet)		0.654*** (0.017)		0.654*** (0.017)	0.645*** (0.016)
Age		-0.002*** (0.000)		-0.002*** (0.000)	-0.002*** (0.000)
Age ²		0.000* (0.000)		0.000* (0.000)	0.000** (0.000)
GDP Per Capita					1.282*** (0.104)
Density					-0.001 (0.002)
R-squared	0.427	0.601	0.428	0.601	0.611
Observations	38,144,444				

Note: (i) Various levels of controls are used to ensure that the models are not misspecified. (ii) City level clustered standard errors in parenthesis to take into account the correlations of error terms.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Heterogeneous Effect of Marijuana Legalization on House Price across Q_τ

	Log(Price per Sq. Foot)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Q10	Q20	Q30	Q40	Q50	Q60	Q70	Q80	Q90
Recreational Vote	-0.016 (0.024)	-0.015 (0.020)	0.014 (0.016)	0.021 (0.015)	0.031** (0.015)	0.046*** (0.016)	0.071*** (0.016)	0.104*** (0.017)	0.123*** (0.023)
Dispensary Date	0.074*** (0.019)	0.099*** (0.017)	0.121*** (0.015)	0.152*** (0.015)	0.172*** (0.015)	0.180*** (0.018)	0.151*** (0.020)	0.108*** (0.025)	0.073** (0.036)
Medical	0.084*** (0.030)	0.090*** (0.023)	0.073*** (0.020)	0.069*** (0.021)	0.064*** (0.022)	0.061** (0.024)	0.057** (0.024)	0.053** (0.023)	0.041* (0.025)
R-squared	0.060	0.123	0.179	0.218	0.243	0.254	0.247	0.220	0.163
Number of Cities	10,640								
Observations	38,145,054								

Note: (i) Possession dummy is excluded since the time gap between vote and possession, or sale and possession are quite small. (ii) House characteristics such as the number of bedrooms, bathrooms, year built and state characteristics such as state GDP, population, land area, and density are controlled in the regressions. (iii) City level clustered standard errors in parenthesis to take into account the correlations of error terms.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Heterogeneous Effect of Marijuana Legalization on House Price across Q_τ

	Log(Price per Sq. Foot)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Q10	Q20	Q30	Q40	Q50	Q60	Q70	Q80	Q90
Recreational Vote	-0.015 (0.024)	-0.014 (0.019)	0.015 (0.016)	0.022 (0.015)	0.030** (0.015)	0.043*** (0.016)	0.068*** (0.017)	0.100*** (0.017)	0.119*** (0.022)
First Dispensary	0.076*** (0.017)	0.104*** (0.016)	0.126*** (0.014)	0.159*** (0.015)	0.184*** (0.016)	0.195*** (0.018)	0.165*** (0.019)	0.120*** (0.024)	0.084** (0.036)
Medical	0.085*** (0.030)	0.091*** (0.023)	0.074*** (0.020)	0.070*** (0.021)	0.066*** (0.022)	0.063*** (0.024)	0.058** (0.024)	0.054** (0.023)	0.042* (0.025)
R-squared	0.060	0.123	0.179	0.218	0.243	0.254	0.247	0.220	0.164
Number of Cities	10,640								
Observations	38,145,054								

Note: (i) First Dispensary is used in place of Dispensary Date for the purpose of a robustness check. (ii) The Possession dummy is excluded since the time gap between Recreational Vote and Recreational Possession, or First Dispensary and Recreational Possession are quite small. (iii) House characteristics such as the number of bedrooms, bathrooms, year built and state characteristics such as state per capita GDP, and density are controlled in the regressions. (iv) City level clustered standard errors in parenthesis to take into account the correlations of error terms.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Spatial Difference-in-Differences: Colorado Subsample

	log (Price)		
	(1)	(2)	(3)
1/2 Mile Zone	0.059*** (0.003)	0.114*** (0.003)	0.123*** (0.003)
1 Mile Zone	-0.036*** (0.003)	0.041*** (0.003)	
2 Mile Zone	-0.067*** (0.003)		
Control Group	Outside 2 Miles	Within 2 Miles	Within 1 Mile
Observations	447,501	256,699	218,605
R-squared	0.411	0.414	0.413

The results in this table are from the same model specification as in the Spatial Difference-in-Differences Table 6, but limited to the observations in the Colorado subsample. House characteristics and county-level economic data are used as controls with robust standard errors.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$

Spatial Difference-in-Differences: Washington Subsample

	log (Price)		
	(1)	(2)	(3)
1/2 Mile Zone	0.061*** (0.002)	0.065*** (0.002)	0.061*** (0.002)
1 Mile Zone	-0.015*** (0.003)	-0.006*** (0.002)	
2 Mile Zone	-0.013*** (0.003)		
Control Group	Outside 2 Miles	Within 2 Miles	Within 1 Mile
Observations	438,332	393,738	347,318
R-squared	0.491	0.510	0.519

The results in this table are from the same model specification as in the Spatial Difference-in-Differences Table 6, but limited to the observations in the Washington subsample. House characteristics and county-level economic data are used as controls with robust standard errors.

*** : $p < 0.01$

** : $p < 0.05$

* : $p < 0.1$