ESSAYS ON CARBON AND WATER PRICING

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Diego Soares Cardoso

May 2021

© 2021 Diego Soares Cardoso

ALL RIGHTS RESERVED

ESSAYS ON CARBON AND WATER PRICING

Diego Soares Cardoso, Ph.D.

Cornell University 2021

This dissertation consists of three chapters that are tied together by their study of pricing mechanisms in the context of Environmental and Natural Resource Economics. The first chapter investigates the efficient pricing of carbon emissions under two non-environmental market distortions: market power and distortionary taxes. Also concerning carbon emissions, the second chapter examines how optimal carbon prices change in response to uncertainty about climate tipping consequences. Lastly, the third chapter switches the focus from efficiency to equity; it provides a quantification of water affordability concerns in the United States and studies policy options to alleviate the burden of unaffordable water services on low-income households.

BIOGRAPHICAL SKETCH

Diego Soares Cardoso was born and raised in Juiz de Fora, Brazil. He received a bachelor's degree in Computer Engineering from the University of Campinas ('10) and a master's degree in Economics from the Federal University of São Carlos ('13). Before beginning his Ph.D. studies at Cornell University, he worked as a Sofware Engineer and Data Scientist in the technology industry and was a graduate student at Iowa State University.

This dissertation is dedicated to the memory of my parents, José Eugênio Cardoso and Flávia Maria Gargiulo Soares Cardoso.

ACKNOWLEDGEMENTS

I am deeply grateful to my advisor and dissertation committee chair, Ivan Rudik, for his continuous support, generosity, and dedication in guiding me throughout my entire Ph.D. education. I would also like to thank my dissertation committee members, Catherine Kling and Shanjun Li, for their invaluable guidance and assistance. I have been immensely fortunate to have learned from and worked with three people I admire as scholars and human beings. I further thank Casey Wichman, who has had an active part in this dissertation, for his mentorship and support. I am also grateful to Todd Gerarden, David Keiser, and Joydeep Bhattacharya for their feedback and advice during various stages of my Ph.D.

I want to thank my wife, Mariana, for her unrelenting love and encouragement. She left so much behind to join me in this adventure. If it were not for her on my side, I would have never gotten this far. I thank my brother, Bernardo, for his crucial help and moral support. I am also grateful to my aunts Maria Goretti and Selma (in memoriam) for their enthusiasm and assistance during turbulent times, which made it possible for me to keep pursuing my dreams.

TABLE OF CONTENTS

Biographical Sketch . . . iii
Dedication . . . iv
Acknowledgements . . . v
Table of Contents . . . vi
List of Tables . . . ix
List of Figures . . . xi

1 Introduction 1

2 Optimal carbon taxation under oligopoly: An application to commercial aviation 5
2.1 Introduction . . . 5
2.2 US aviation and climate change . . . 13
2.3 Theoretical framework . . . 17
2.3.1 An illustration of welfare effects . . . 17
2.3.2 Defining social welfare . . . 20
2.3.3 Marginal welfare effects . . . 23
2.3.4 Marginal abatement cost and optimal tax . . . 24
2.4 A model of commercial aviation . . . 26
2.4.1 Definitions . . . 26
2.4.2 Consumers . . . 27
2.4.3 Airlines . . . 29
2.5 Data . . . 31
2.6 Empirical approach and estimation . . . 36
2.6.1 Sufficient statistics . . . 37
2.6.2 Structural parameters . . . 41
2.7 Welfare analyses . . . 49
2.7.1 Marginal welfare changes with a carbon tax . . . 50
2.7.2 Non-marginal welfare changes and the optimal carbon tax 51
2.7.3 Revenue-neutral tax substitution . . . 57
2.7.4 Limitations . . . 61
2.8 Conclusion . . . 62

3 Climate policy under uncertain post-tipping warming 64
3.1 Introduction . . . 64
3.2 Climate tipping points in the geosciences and economics . . . 68
3.2.1 Definitions of tipping points . . . 68
3.2.2 Modeling choices . . . 70
3.3 Optimal plans for an uncertain world . . . 71
3.3.1 Modeling the climate sensitivity uncertainty . . . 71
3.3.2 The Climate-Economy model . . . 78
3.4 Results . . . 81

3.5 Discussion . . . 90
3.6 Conclusion . . . 94

4 Water Affordability in the United States **96**
4.1 Introduction . . . 96
4.1.1 Measuring affordability . . . 100
4.2 Data and Methods . . . 102
4.2.1 Data and calculations . . . 102
4.2.2 Estimation of socioeconomic and demographic conditional correlations . . . 104
4.2.3 Policy simulations . . . 106
4.3 Results . . . 107
4.3.1 High-resolution income data and local prices are critical for measuring household-level water affordability . . . 107
4.3.2 One out of every seven households spends more than 4.5% of annual income on water and sewer bills . . . 108
4.3.3 Households in the lowest income bracket pay 8.1% of annual income on water and sewer bills . . . 110
4.3.4 Water and sewer bills are regressive, but less so, when compared with annual expenditure data . . . 111
4.3.5 Water affordability concerns are pervasive across the US, driven by the local income distribution . . . 114
4.3.6 Water affordability concerns are significantly correlated with select community characteristics . . . 116
4.3.7 Affordability policies that provide lump-sum rebates for low-income households and are funded by income taxes are most effective . . . 120
4.4 Discussion and conclusion . . . 123

A Appendix of Chapter 2 **128**
A.1 Estimation procedures and welfare analyses . . . 128
A.1.1 Derivation of sufficient statistic ε . . . 128
A.1.2 Marginal welfare calculations with sufficient statistics . . . 130
A.1.3 Solving market equilibria and operating costs . . . 132
A.1.4 GMM estimation with high-dimensional fixed effects . . . 135
A.2 Additional regression results and alternative estimators . . . 138
A.2.1 First stage of the sufficient statistics regression . . . 139
A.2.2 Estimating the structural model with two-stage least squares 140
A.2.3 Alternative specifications . . . 142
A.3 Additional details on the data set . . . 147
A.4 Market-specific marginal abatement costs . . . 152

B Appendix for Chapter 3 **154**
B.1 The feedback-implied distributions . . . 154
B.1.1 Derivation of the used distributions . . . 154
B.1.2 Distribution calibration condition . . . 157
B.2 A DICE model with uncertain post-threshold regime . . . 158
B.2.1 Baseline model . . . 158
B.2.2 Solving the model . . . 163
B.3 Additional results . . . 165

C Appendix for Chapter 4 **169**
C.1 Data . . . 169
C.1.1 Water rates and water use . . . 169
C.1.2 Climate zones . . . 174
C.1.3 Socioeconomic-demographic factors . . . 174
C.2 Methods . . . 176
C.2.1 Preliminaries . . . 176
C.2.2 Regression analysis and bootstrapping . . . 178
C.2.3 Policy simulations . . . 183
C.3 Water affordability and socio-economic-demographic factors in US regions . . . 185
C.4 Sensitivity analysis . . . 188
C.4.1 Affordability assessment at the estimated consumption levels . . . 189
C.4.2 Program simulations . . . 192
C.5 Robustness checks . . . 194
C.5.1 Estimation with observed average consumption . . . 195
C.5.2 Estimation with fixed per capita consumption . . . 195
C.6 Comparison to other metrics in the literature . . . 197
C.6.1 National affordability assessment in Mack and Wrase . . . 197
C.6.2 Local affordability assessments using alternative metrics . 201

References **205**

LIST OF TABLES

2.1 Summary statistics. 35
2.2 Estimates for the price elasticity of aggregate demand. 40
2.3 Estimation results for the joint demand-supply estimation. 46
2.4 Average operating revenue, costs, and markups per revenue passenger-mile (RPM). 48
2.5 Differences in average emissions, markups, and taxes for stop and nonstop flights. 58

3.1 Illustration of the next-period distributions for the fixed economy. 93

4.1 Percentage of households who pay more than 4.5 percent of annual household income on combined water and sewer bills by income and consumption data resolution. 109
4.2 Expenditure on water and sewer relative to income by income bracket[a]. 111
4.3 Conditional correlations between water affordability and select socioeconomic characteristics. 118

A.1 First stage regression of the sufficient statistic estimation. 139
A.2 Estimates of demand parameters using 2SLS. 141
A.3 First stage regressions of the demand model using 2SLS. 144
A.4 Estimates of supply parameters using 2SLS. 145
A.5 Estimates of demand parameters under alternative specifications. 146
A.6 Airline groups and types in the dataset. 148

B.1 Parameters of the baseline economic model. 164
B.2 Parameters of the baseline climate model. 164
B.3 Parameters for the climate sensitivity uncertainty cases. 167
B.4 Parameters for alternative specifications. 168

C.1 Summary statistics 180
C.2 Results for the first-stage regression[a] 181
C.3 Conditional correlations between water affordability and select socioeconomic characteristics for subsamples containing states in US regions[a] 186
C.4 Summary statistics for US regions[a] 187
C.5 Percentage of households who pay more than 4.5 percent of annual household income on combined water and sewer bills by income and consumption data resolution under different income elasticity assumptions 190
C.6 Conditional correlations between water affordability and select socioeconomic characteristics[a] 196

C.7 A comparison between Teodoro's AR_{20} and HM metrics for water and sewer affordability and our method for most populous US cities[a] 203

LIST OF FIGURES

2.1 Share and growth trends in US domestic aviation. 14
2.2 Welfare-decreasing Pigouvian tax with multiple market distortions in a polluting monopoly. 18
2.3 Pigouvian and optimal taxes with multiple market distortions in a polluting monopoly. 19
2.4 Passenger traffic in cities and metro areas included in the data set. 33
2.5 Passenger traffic between cities and metro areas. 34
2.6 Predicted changes to prices and quantities from linear extrapolation of marginal effects. 51
2.7 Predicted changes with the introduction of a carbon tax. 53
2.8 Marginal abatement cost and tax wedges at various abatement levels. 54
2.9 Distribution of baseline marginal abatement costs (MACs) across markets. 56
2.10 Distribution of optimal taxes in a market-specific scheme under a social cost of carbon of \$230/ton CO_2. 56
2.11 Predicted changes from removing the sales tax and introducing a carbon tax. 59
2.12 Welfare changes due to sales tax and a revenue-neutral carbon tax under different values for the Social Cost of Carbon. 60

3.1 Calibrated distributions f for each uncertainty case. 76
3.2 Calibrated distributions s for each uncertainty case. 77
3.3 Main components and interactions in the DICE model. Solid, shaded boxes represent the states. Planner's decisions are represented by diamond arrows, with shaded ellipses being the variables of choice. Dashed boxes represent endogenous processes. Dotted lines represent exogenous processes. 79
3.4 Results for the baseline model (with CRRA utility function). . . . 83
3.5 Other results for the baseline model (with CRRA utility function). 84
3.6 Results for Epstein-Zin-Weil preference. 86
3.7 Other results for Epstein-Zin-Weil preference. 87
3.8 Results for the damage function estimated in Howard and Sterner (2017). 88
3.9 Other results for the damage function estimated in Howard and Sterner (2017). 89
3.10 Next-period temperature uncertainty for the fixed economy illustration. 91
3.11 Illustration of the next-period damage uncertainty for the fixed economy. 92
3.12 V_0 at different warming levels at different periods along the optimal trajectory for the RIS-Beta case. 94

4.1 U.S. price indexes and income distribution over time. 98
4.2 Proportion of households above affordability threshold based on water and sewer expenditure-shares of income, based on varying degrees of income data resolution. 110
4.3 Incidence of US water bills using Consumer Expenditure Survey data. 112
4.4 Incidence of water bills by region and race using Consumer Expenditure Survey data. 113
4.5 Geographic distribution of water affordability using estimated levels of consumption to calculate combined water and sewer bills. 115
4.6 Geographic distribution of water affordability in block groups within urban areas. 117
4.7 Results of program simulations. 122

A.1 Fuel efficiency by stage length for the aircraft models most used. 149
A.2 Jet fuel and crude oil prices. 151
A.3 Histograms of baseline market-specific marginal abatement costs (MAC) for each market size quartile. 152
A.4 Cumulative distributions of baseline market-specific marginal abatement costs (MAC) for each market size quartile. 153

B.1 Additional results for the baseline model (with CRRA utility function). 166
B.2 Additional results for a model with EZW utility function. 167
B.3 Additional results for a model with Howard and Sterner (2017) damage function. 168

C.1 Combined water and sewer bills in counties with populations greater than 500,000. 172
C.2 Correlation between the proportion of the combined water and sewer bill (CWSB) that is the fixed access fee (evaluated at 50 gppd) relative to the proportion of households with unaffordable water and sewer service within a county. 173
C.3 Proportion of households above the affordability threshold based on water and sewer expenditure-shares of income for different assumptions of income elasticity. 191
C.4 Average expenditure shares on water and sewer service by income bracket for business-as-usual and each policy option assuming different values of elasticities. 193
C.5 A comparison of the proportion of households above affordability threshold for Mack and Wrase's approach (MW) and our proposed method. 200

C.6 Household annual income distributions in our sample and for the entire United States. 201

CHAPTER 1

INTRODUCTION

In the Environmental and Resource economists' toolkit, most market-based solutions aim to get the prices "right." At the root of many environmental and natural resource issues are external costs that cause misaligned incentives and lead to outcomes that make society as a whole worse-off. As a solution to these issues, market-based policy instruments improve social welfare by implementing mechanisms that realign incentives and result in equilibrium prices that reflect the scarcity and external costs.

In real-world contexts, however, finding the "right price" often requires much more than the simple theoretical intuition can provide. Focusing on economic efficiency (i.e., improving aggregate social welfare) and climate change, the next two chapters of this dissertation calculate the optimal price of carbon emissions in different contexts.

In a short-run, partial equilibrium context, chapter 2 examines how market imperfections affect the optimal carbon price. Economic theory shows that when an environmental externality is the only market imperfection, a corrective environmental equal the value of marginal damage maximizes social welfare. When multiple imperfections exist, however, estimates of marginal damage are not sufficient to set an optimal tax: it is also necessary to understand the market structure and estimate the effect of imperfections such as market power and distortionary taxes. Following this insight, chapter 2 estimates the optimal carbon tax for US domestic aviation by combining a theoretical framework of optimal environmental taxation and structural econometric methods for the study of oligopolies. Based on estimated model parameters and sufficient statistics for

marginal welfare changes, I (i) estimate marginal and total costs of emission abatement via carbon taxation, (ii) calculate the optimal carbon tax in the presence of non-carbon distortions, and (iii) examine the extent to which existing taxes are substitutes for a carbon tax.

Chapter 2 shows that the marginal cost of abatement with carbon taxation starts at \$208/ton CO_2. Thus, if the social cost of carbon (SCC) is smaller than this value, any positive carbon tax would decrease social welfare in the short run. Under a social cost of carbon (SCC) of \$50/ton CO_2, for example, a carbon tax equal to the SCC would reduce emissions by 14% but result in a net loss of \$4.5 billion per year in social welfare (about 5% of the sector's yearly aggregate revenue). Under a higher SCC of \$230/ton CO_2, the optimal carbon tax would be \$40/ton CO_2, much lower than the Pigouvian tax level. This chapter also shows that current taxes on air travel correspond to a carbon tax of approximately \$52/ton CO_2. However, by being proportional to the fare instead of the externality, sales taxes are an imperfect instrument to address the carbon externality, and alternative tax schemes could improve efficiency. Implementing a revenue-neutral carbon tax to replace current taxes would lead to substantial welfare gains. These findings illustrate how the lack of scalable abatement technologies and sizable market imperfections lead to high carbon abatement costs, highlighting a key challenge for climate policy in aviation.

In contrast with the second chapter, the analysis in chapter 3 looks at a long-run, general equilibrium context in the presence of substantial uncertainty about climate change. Specifically, the risk of an irreversible, discontinuous, and known change in the global climate system has been shown to substantially affect climate policy. Nevertheless, the magnitude of the climate regime changes

caused by climate tipping elements is inherently uncertain. Building on a recursive dynamic model of climate policy, this chapter examines how the optimal carbon tax responds to uncertainty on climate sensitivity in the post-tipping regime. I consider moderate to extreme cases of uncertainty and alternative specifications for risk attitude and damage function. Results show that adding mean-preserving uncertainty about the post-tipping climate sensitivity raises the optimal tax by at most 0.2%. The small extent of the effects is a consequence of the small short-term impact of varying the climate sensitivity.

Even though efficient policies maximize aggregate welfare, the costs associated with policies might fall disproportionally on specific groups of society. The final chapter of this thesis, co-authored with Casey Wichman, switches the focus from efficiency to equity in prices. Affordable water and sewer service is a growing concern in the United States, although the extent of the problem is not known, and the effectiveness of different policy options are underexplored. The cost of water and wastewater services is rising three times faster than inflation, and water infrastructure will require more than 1-trillion USD in investment over the next 20–25 years. In this chapter, we document the extent of water affordability concerns in the US across income, geography, and race. We also examine how policies can be designed to assist low-income households.

Chapter 4 shows that 13.6% of households have water and sewer expenditures greater than 4.5% of annual household income in our sample, and those in the lowest income decile pay on average 8.1% of their annual income on water and sewer service. Our estimates are based on one of the most comprehensive data sets of water and sewer prices to date, capturing approximately 45% of the US population, matched with Census block-group-level socioeconomic

characteristics and typical consumption levels. We demonstrate that using median household income at the county level drastically understates the extent of the water affordability problem relative to using the whole income distribution within Census block groups. Additionally, we find that the number of households facing affordability concerns is positively associated with water and sewer price levels, impoverished residents, and the proportion of black residents even after conditioning on poverty levels. Lastly, we show that self-sufficient water affordability policies that provide a lump-sum rebate to low-income households and are paid for by income taxes are more effective at redistributing the burden borne by low-income customers than policies that change marginal incentives for water and sewer consumption.

CHAPTER 2

OPTIMAL CARBON TAXATION UNDER OLIGOPOLY: AN APPLICATION TO COMMERCIAL AVIATION

2.1 Introduction

"Levy a tax equal to the marginal external cost" is a foundational policy prescription in the economic analysis of externalities. This type of tax, known as Pigouvian taxation, makes agents internalize the external costs, maximizes social welfare, and leads the market to its efficient equilibrium in the absence of other distortions. As such, Pigouvian taxes—and equivalent market-based instruments—enjoy broad support among economists as a policy to mitigate environmental damages.

The optimality of a Pigouvian tax, however, hinges on the assumption that the target externality is the only market imperfection; when this assumption holds, optimal environmental taxation only requires knowledge about marginal damages. But numerous polluting sectors deviate from perfect markets. For instance, several polluting markets are oligopolies. As Buchanan (1969) has demonstrated, optimal taxes differ from marginal environmental costs when firms have market power. Furthermore, most sectors are subject to non-Pigouvian, distortionary taxes that lower equilibrium quantities. These taxes decrease related emissions and act as partial substitutes for a Pigouvian tax. When multiple imperfections exist, estimates of marginal damage are not sufficient for an efficient policy: it is also necessary to understand the market structure and quantify its imperfections. An environmental tax based exclusively on marginal damages can even decrease welfare when the assumption of otherwise

perfect markets does not hold.

This chapter estimates the optimal carbon tax for the US domestic aviation sector. In this context, the optimal carbon tax is the one that achieves second-best equilibrium: it corrects for environmental externality while taking as given market power and existing taxes. To estimate an optimal tax, I combine the theoretical framework of environmental taxation with structural econometric methods for the study of oligopolies. Using air travel and airline financial data from the US Department of Transportation, I estimate sufficient statistics for welfare changes and an oligopoly model for the sector. Based on these estimated parameters, I generate counterfactual scenarios under various levels of taxation to (i) estimate marginal and total welfare costs of emission abatement via carbon taxation, (ii) calculate the second-best carbon tax in the presence of non-carbon distortions, and (iii) examine the extent to which existing air travel taxes are substitutes for a carbon tax and whether replacing these existing taxes for a revenue-neutral carbon tax can improve welfare.

Commercial aviation is a notoriously concentrated sector that has proven challenging for climate policy. A rich literature has documented that airlines have substantial market power as a result of the oligopolistic nature of the sector (e.g., Borenstein, 1989; Ciliberto & Williams, 2010). Adding to market distortions, air travel is also subject to non-Pigouvian, distortionary taxes. Regarding climate-related externalities, the sector is responsible for approximately 3% of global greenhouse gas emissions and 5% of the radiative forcing leading to climate change (Lee et al., 2009). With limited regulation, carbon emissions from aviation are projected to continue growing (Owen et al., 2010) and may account for as much as 22% of global greenhouse gas emissions by 2050 (European Par-

liament, 2015). Despite efforts to curb emissions from international aviation, the scope of policies has been limited, and large domestic air travel markets have not been addressed. Most notably, the US—the largest aviation market—has yet to adopt a comprehensive carbon emission policy for commercial aviation.

In this chapter, the aviation carbon tax is implemented as a volumetric uniform tax on jet fuel. Even though isomorphic alternatives exist, such as tradable emission permits and taxes on air travel tickets, I focus on a jet fuel tax for three practical reasons. First, jet fuel is a homogeneous commodity and jet fuel burn is directly associated with carbon emissions, so a volumetric tax provides a close approximation of the carbon externality. Second, jet fuel is a single-use commodity with limited leakage potential in domestic markets.[1] Third, jet fuel is already taxed in the US; hence, this carbon tax builds on an existing tax structure, which lowers the institutional requirements for its implementation.

I find that any positive carbon tax would decrease social welfare in the short run, based on a social cost of carbon[2] (SCC) of \$50 per metric ton of CO_2. This result follows from the marginal loss of aggregate private surplus exceeding the present value of the marginal damage avoided in the current equilibrium. If a carbon tax of \$50/ton CO_2 were implemented, it would reduce emissions by 14% but would result in a net loss of \$3.5 billion per year in social welfare (about 3.5% of the sector's yearly aggregate revenue). The burden of this carbon tax would fall slightly more on airlines, with a \$4.9B loss in operating profits against a \$4.2B in reduction in consumer welfare; however, an additional \$4.4B raised in taxes could partially compensate for the losses to either side.

[1] Though carrying excess fuel from abroad on international flights may be possible, the additional weight increases the fuel burn rate, limiting the economic feasibility of this strategy.

[2] The social cost of carbon indicates the present discounted value of the stream of climate damages from an additional ton of CO_2 emissions.

Under a higher SCC of \$230/ton CO_2—an upper-bound estimate from Daniel et al. (2019)—I find that the optimal carbon tax is approximately \$40/ton CO_2. Thus, the optimal tax level is less than one-fifth of the standard Pigouvian prescription. The striking difference between the optimal tax and the associated marginal damage can be traced to existing market imperfections. Consistent with financial reports from airlines, the predicted average operating markup is approximately 4 cents per passenger-mile. However, when considering emissions instead of passenger-miles, the average markup exceeds \$240/ton CO_2. By the same account, the current excise tax corresponds to \$69/ton CO_2 on average. These substantial distortions accentuate the loss of private surplus when carbon taxes increase. As a result, if emission reductions are achieved exclusively through demand reduction—and not by long-run technology adjustments such as non-fossil fuels—, the initial marginal abatement cost is \$208/ton CO_2. These findings highlight one of the main challenges for climate policy in the sector: the abatement cost of carbon emissions in aviation is high, at least with limited abatement technologies available in the short run.

Lastly, I find that implementing a revenue-neutral carbon tax to replace the current sales tax would improve social welfare. The current sales tax, set at 7.5% of the fare, partially substitutes for a carbon tax by increasing average ticket prices and reducing demand and emissions. This substitution is inefficient, however, because fares may be inversely correlated with emissions within a market: nonstop flights are shorter and emit less CO_2 per passenger but are priced higher. Eliminating the sales tax and implementing a carbon tax of approximately \$52/ton CO_2 would raise the same tax revenue. The welfare gains from this tax substitution vary from \$250 to \$670 million per year, depending on the SCC scenario. For an SCC of \$50/ton CO_2, these gains correspond to a 18%

reduction in the deadweight loss of taxation. For an SCC of \$230/ton CO_2, a case in which the current sales tax improves welfare relative to having no taxes, this revenue-neutral substitution increases welfare gains from taxation by about 20%.

This chapter makes four contributions to the literature. First, it provides an assessment of the welfare consequences of a carbon tax on US domestic aviation. Previous studies have focused on estimating how hypothetical policies would affect prices and demand, overlooking welfare consequences and the role of non-environmental market imperfections. Among the few studies in this literature, Pagoni and Psaraki-Kalouptsidi (2016) estimate the impact of four carbon tax levels on prices and demand. For instance, they find that a carbon tax of \$50/ton CO_2 would increase prices by 5.9% and decrease air travel demand by 11.2%. Brueckner and Abreu (2017) investigate the determinants of fuel consumption by US commercial airlines, including fuel prices; they estimate that a carbon tax of \$40/ton CO_2 on jet fuel would decrease emissions by 2.2%. Using a quantile regression approach, Fukui and Miyoshi (2017) estimate that an increase of 4.3 cents in the US jet fuel tax would lead to a decrease of 0.14 to 0.18% in emissions in the short run. Finally, Winchester et al. (2013) use an economy-wide model to study the potential impacts of a cap-and-trade program and conclude that such program would not be enough to curb emissions growth; they indicate that the abatement costs in aviation are high and suggest that the best approach would be to subsidize emission reductions in other sectors with lower abatement costs. These studies provide estimates of the impact of carbon taxes on market outcomes but do not consider the welfare consequences of such a policy.

Second, my results contribute to the study of environmental externalities under imperfect competition. In a seminal work, Buchanan (1969) shows how the standard Pigouvian tax can lead to welfare losses under monopoly. Barnett (1980) formalizes this intuition and demonstrates the importance of market structure when considering optimal externality taxes. More recent theoretical results have shown how optimal environmental policy departs from the standard Pigouvian taxation when market imperfections interact.[3]

Empirical analyses of environmental policy in imperfect markets have emerged more recently. Mansur (2007), for example, studies oligopolies in restructured electricity markets and finds that one third of the reductions in emissions can be attributed to market power. Ryan (2012) and Fowlie, Reguant, and Ryan (2016) apply empirical tools frequently used in the industrial organization literature to investigate the interaction of market imperfections; they show how environmental regulations have increased market power and led to welfare losses in the US cement industry.

A series of recent papers have shown that, with market power and incomplete cost pass-through, carbon taxes induce low abatement rates and can even increase emissions. Preonas (2017) investigates how markups in coal-by-rail transportation to power plants respond to changes in natural gas prices. The author finds evidence of incomplete pass-through and shows that a carbon tax would lead to lower abatement, as firms adjust markups and absorb part of the shocks. Leslie (2018) examines the introduction and subsequent repeal of a carbon tax on electricity in Australia; he finds that emissions increased with the carbon tax as a result of market power in the sector. Other papers have documented incomplete cost pass-through in a diverse set of carbon-intensive sec-

[3]Requate (2006) provides an extensive review of advancements in this literature.

tors (e.g., Fabra & Reguant, 2014; Ganapati et al., 2016; Muehlegger & Sweeney, 2017; Lade & Bushnell, 2019), thus indicating that the design of climate change policies must pay special attention to existing market imperfections. While most of this literature has focused on evaluating the incompleteness of existing environmental policy with imperfect competition, in this chapter I incorporate these lessons in the design of optimal environmental policy. My results demonstrate how market imperfections affect policy efficiency and have important implications for environmental policy implementation in concentrated sectors.

Third, this chapter adds to the study of optimal environmental policy in second-best settings. Since the work of Sandmo (1975) on optimal taxation with externalities, an extensive literature has emerged on the interactions between environmental and non-environmental taxes (Goulder, 1995). A significant part of this literature is concerned with the conditions for a double dividend: when the substitution of pre-existing distortionary taxes for environmental policies leads to Pareto improvements by jointly correcting tax distortions and externalities. Most frequently, second-best settings have considered general equilibrium effects with distortionary taxes on labor and other inputs (e.g. Bovenberg & Mooij, 1994; Bovenberg & Goulder, 1996; Cremer et al., 1998; Goulder, 1998; Goulder et al., 1999). Under pre-existing taxation, optimal environmental taxes are frequently smaller than the marginal damage (Bovenberg & Mooij, 1994; Parry, 1995). However, depending on the setting, optimal taxes can theoretically exceed marginal damages (Cremer et al., 1998; West & Williams, 2004; Ren, Fullerton, & Braden, 2011). Furthermore, studies have indicated that distortionary taxes increase the cost of environmental policy (Goulder et al., 1999) and that monopoly power can exacerbate this effect (Fullerton & Metcalf, 2002); these increased costs can be attenuated by revenue-neutral tax substitution,

even when the double dividend fails to materialize (Goulder, 1998). The present analysis contributes to this literature by showing how existing *ad valorem* tax distortions interact with market power in partial equilibrium and increases the cost of reducing emissions. Moreover, applying recently developed empirical methods, this chapter quantifies the individual effects of each distortion.

Finally, this chapter contributes to the literature on empirical welfare economics by demonstrating how sufficient statistics and structural approaches can be combined to assess marginal and non-marginal effects of environmental policies in imperfect markets. The sufficient statistics approach, introduced by Chetty (2009), evaluates marginal welfare changes from a policy based on a small set of key parameters—often expressed as elasticities—that can be identified through reduced-form estimation. These marginal changes are then used to verify whether a policy intervention is granted. Determining the optimal tax, however, involves evaluating non-marginal changes and requires additional structural assumptions (Kleven, 2020). In line with Kleven (2020), this chapter shows how sufficient statistics can be used to quantify the role of non-environmental market distortions, to estimate marginal abatement costs with a carbon tax, and to assess whether such a policy is likely to improve welfare. These statistics offer a consistency check for predictions of a structural model, which can be used to evaluate non-marginal changes and derive optimal taxes.

The remainder of this chapter is organized as follows. Section 2.2 summarizes key characteristics of the US aviation sector and the challenges they present for climate policy. Section 2.3 introduces the chapter's welfare framework and derives expressions to characterize welfare changes, market distortions, and optimal taxes. A model of the US aviation sector is outlined in Section 2.4. Sec-

tion 2.5 describes the data used in this chapter. Section 2.6 explains the estimation procedures and discusses estimated parameters. The estimation of optimal taxes and the impacts of an aviation carbon tax are presented in Section 2.7. Section 2.8 offers concluding remarks.

2.2 US aviation and climate change

This section summarizes the three relevant aspects of the industry: market structure, existing taxes, and carbon emissions. The US domestic commercial aviation market is the largest in the world. As panel (a) in Figure 2.1 shows, the US accounts for approximately 30% of all passengers carried on domestic flights and 45% of domestic passenger-miles served (IATA, 2019). The US domestic aviation market is second in growth only to China (IATA, 2019). After the deregulation of US commercial aviation in the late 1970s, the sector has experienced tremendous expansion in service, growing from around 250 billion revenue passenger-miles (RPM) in the early 1980s to over a trillion RPMs per year in the late 2010s, as shown in Figure 2.1, panel (b).

Sector structure. The industry has also seen changes in its players, with various rounds of entry, bankruptcy, and consolidation.[4] We can organize airlines into two main groups (Belobaba et al., 2015). One group includes the *legacy* airlines, alluding to the fact that these firms have been operating since the pre-deregulation era. This group includes American, Delta, and United Airlines. Some of the distinguishing features of these players are their extensive ser-

[4]Borenstein and Rose (2014) present a comprehensive overview of the US aviation industry, including its history, trends, and unique challenges.

Figure 2.1: Share and growth trends in US domestic aviation.

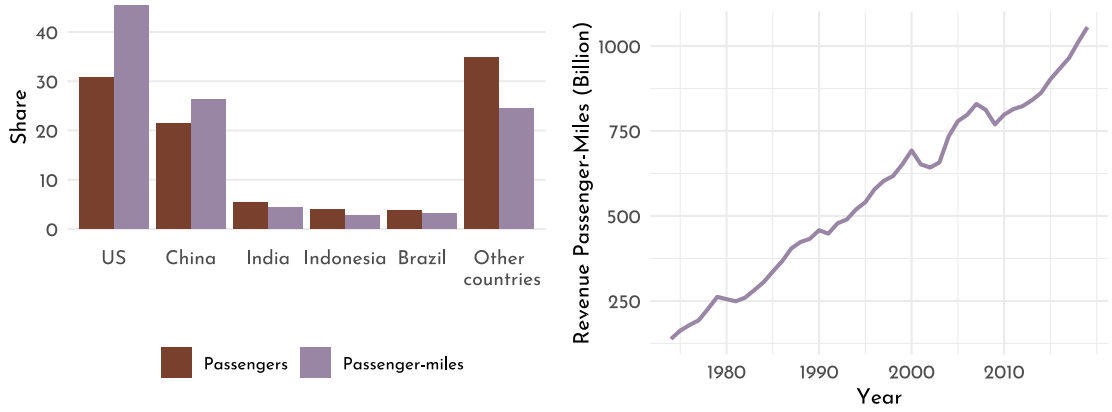

(a) Participation shares of major countries relative to global domestic markets.

(b) Continuous growth in service in the US

Source of data: International Air Transport Association and US Bureau of Transportation Statistics.

vice networks with large hubs; more rigid cost structures with higher levels of unionization; and bundled, higher-quality services (such as meals and in-flight amenities). The other group is formed by *low-cost carriers* (LCCs), which follow the "no-frills" business model successfully implemented by Southwest Airlines. Examples of other airlines in this category are Spirit and Allegiant. As the name suggests, LCCs focus on running cost-efficient operations. This involves, for example, flying point-to-point services from smaller airports instead of maintaining large hubs, keeping a high aircraft utilization rate, and unbundling passenger services by charging extra fees for baggage, reservations, food, and beverages.

In practice, this categorization is more of a conceptual construction than an accurate description of how these airlines operate. The financial success of LCCs has led legacy carriers to adopt some of the LCC practices. On the opposite end, the quest for diversification has also led some LCCs, such as JetBlue, to invest in higher service quality. A third group of airlines can be categorized as

regional carriers; these players are either small, independent companies that run on limited networks, or carriers that operate in partnership with larger airlines to provide connections from hubs to smaller airports—under the brand name of United Express or American Eagle, for example (Belobaba et al., 2015).

With the small number of airlines and high fixed and entry costs, the aviation industry largely operates as an oligopoly. An extensive literature has documented evidence of market power in the US aviation industry. Though a review of this literature is beyond the scope of this brief sector description, prior findings present some common themes. The existence of a hub premium, for example, is a source of market power (Borenstein, 1989, 1991; Lee & Luengo-Prado, 2005; Lederman, 2007, 2008; Berry & Jia, 2010). Other mechanisms generating and maintaining market power are tacit collusion (Evans & Kessides, 1994; Ciliberto & Williams, 2014; Aryal et al., 2019), entry deterrence (Ciliberto & Williams, 2010; Aguirregabiria & Ho, 2012; Ciliberto & Zhang, 2017), and mergers and consolidation (Kim & Singal, 1993). In the opposite direction, increased competition from LCCs, a trend initially attributed to the "Southwest effect", has been found to reduce prices and markups (Morrison, 2001; Goolsbee & Syverson, 2008; Brueckner et al., 2013).

Taxes and fees. In the US, commercial flights are subject to a sales tax, a fuel tax, and various fees. The sales tax corresponds to the US Federal Excise Ticket Tax, set at 7.5% of the base fare. This tax is dedicated to the Airport and Airway Trust Fund (AATF), which most notably funds the Federal Aviation Administration (FAA, 2020). Moreover, jet fuel used for commercial aviation is subject to a federal tax of 4.3 cents, also appropriated by the AATF, plus a 0.1 cent fee, appropriated by the Leaking Underground Storage Tank Trust Fund. There are

three fees for domestic flights in the US: (i) the Federal Security Surcharge, at $11.20 per domestic round-trip itinerary; (ii) the Federal Flight Segment Tax, which charges $4.20 per domestic segment; and (iii) Passenger Facility Charges, costing on average $4.50 per departing airport. My model will capture these taxes and fees, as they will have implications for pricing behavior.

Emissions. Aviation accounts for 2–3% of global CO_2 annual emissions (Owen et al., 2010) and is one of the sectors with the fastest growth in emissions. Between 1990 and 2016, greenhouse gas emissions from aviation grew by 98% (FCCC, 2018). For the first half of the 21st century, these emissions are projected to grow by 200–360% (Owen et al., 2010). With the wide expansion of air travel and limited regulation of its carbon emissions, the sector lags behind other industries in decarbonization. By 2050, aviation may account for 22% of total CO_2 emissions (European Parliament, 2015).

The total contribution of aviation to climate change is larger than its share of CO_2 emissions, accounting for as much as 5% of the radiative forcing leading to global warming (Lee et al., 2009). Jet fuel burn releases other components that affect heat transfer in the atmosphere, including water vapor, nitrous oxides, soot, and contrails. Though some components favor atmospheric cooling, such as aerosols and methane reduction due to nitrous oxides, the net effect of jet fuel burn produces warming above the individual contribution of CO_2 (Lee et al., 2009).

Each gallon of jet fuel burned emits an average of 9.57 kg CO_2 (EIA, 2019). To account for other greenhouse gases, these emissions can be converted to CO_2-equivalent terms. In this chapter, I use a 1.4 conversion factor, which is

the central estimate in Azar and Johansson (2012). This factor relies on a discount rate of 3% per year, which needs to be incorporated in their estimates because different greenhouse gases have different different lifespans in the atmosphere. Hence, one gallon of jet fuel accounts for approximately 13.4 kg of CO_2-equivalent; I use this emission intensity to calculate climate damages.

2.3 Theoretical framework

In this chapter, a carbon tax is considered optimal in a second-best sense: it maximizes social welfare, taking market power and distortionary revenue-raising taxes as given. This section introduces a framework for evaluating the welfare effects of a carbon tax, and then uses that framework to characterize the optimal tax and the role of non-environmental distortions. Throughout the chapter, all analyses are made in partial equilibrium, focusing only on the markets that generate the environmental externality of interest.

2.3.1 An illustration of welfare effects

Before proceeding with a formal analysis, let us start by graphically illustrating a simplified case. As in Buchanan (1969), consider a single-product monopolist. This single product has a constant marginal cost of production MC, and its consumption generates a marginal externality ϕ. Moreover, the demand function for this good is linear.

Figure 2.2 extends the original diagram in Buchanan (1969) by adding a sales tax, which creates a wedge between the inverse demand curve (D) and the price

Figure 2.2: Welfare-decreasing Pigouvian tax with multiple market distortions in a polluting monopoly.

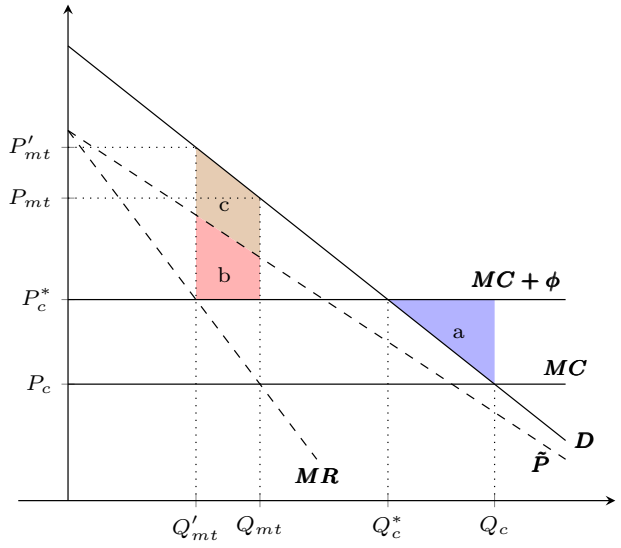

Notes: D is demand, $\tilde{P}$ is price received (before taxes), MR is marginal revenue, MC is marginal private cost, and ϕ is marginal damage from pollution. Area a is welfare gain from a Pigouvian tax under perfect competition. Areas b and c are welfare losses from monopoly power and distortionary taxation, respectively.

received by the monopolist ($\tilde{P}$). If this market satisfies perfect competition, the competitive equilibrium price (P_c) will equal the marginal private cost (MC). Then, the externality will generate a dead-weight loss (area a), which could be corrected by levying a per-unit tax equal to ϕ; i.e., it would achieve the standard Pigouvian setting at the efficient equilibrium (Q^*_c, P^*_c).

When the firm is a monopolist and a sales tax exists, however, the initial market equilibrium is at Q_{mt}, P_{mt}. Introducing a tax ϕ leads the monopolist to decrease supply even further, to Q'_{mt}, with a higher equilibrium price P'_{mt}. As a result, the interaction between the externality tax and other market distortions decreases welfare for two reasons. First, market power leads to the reduction represented by area b; this is the welfare loss identified in Buchanan (1969). Second, sales tax distortion drives the loss represented by area c. Hence, for the

Figure 2.3: Pigouvian and optimal taxes with multiple market distortions in a polluting monopoly.

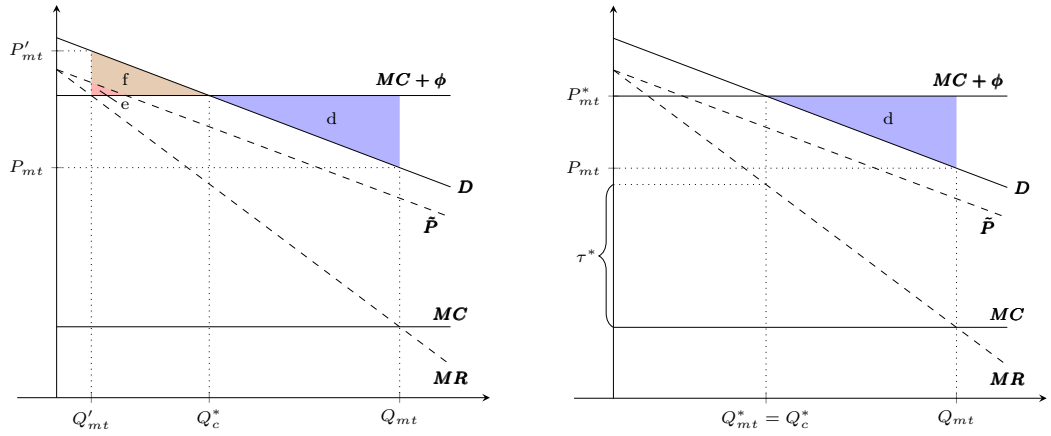

(a) Welfare-improving Pigouvian tax

(b) Optimal tax

Notes: D is demand, $\tilde{P}$ is price before taxes, MR is marginal revenue, MC is marginal private cost, and ϕ is marginal damage from pollution. Area d is welfare gain from correcting the environmental externality. Areas e and f are welfare losses from monopoly power and distortionary taxation, respectively.

case illustrated in Figure 2.2, the standard Pigouvian tax would decrease social welfare. In fact, since $Q_{mt} < Q^*_c$, any positive externality tax would lead to a welfare loss.

The case represented in Figure 2.2 is a particular one: the externality is small relative to other distortions. Panel (a) in Figure 2.3 illustrates a different scenario, where the Pigouvian tax still improves welfare, but not in an efficient manner. In this case, the Pigouvian tax corrects the externality and increases welfare proportionally to the area d. However, this tax leads the market equilibrium to Q'_{mt}, below the efficient level. As a result, there are welfare losses corresponding to areas e and f (analogous to b and c in Figure 2.2). The net effect on welfare can be positive, as long as the gains from the corrected externality exceed losses from the other distortions.

When the externality is large relative to other distortions, an efficient alternative to the Pigouvian tax is a smaller tax τ^*, as illustrated in Figure 2.3, panel (b). This optimal tax leads the monopolist to supply at the efficient market outcome, Q_c^*. Hence, in this particular case, a social planner would be able to efficiently correct the environmental externality under market power and sales tax distortions.

2.3.2 Defining social welfare

The intuitive results obtained from the graphical analysis above can be formalized. Social welfare in a given market m comprises four components: consumer surplus, firms' operating profits, tax revenue,[5] and environmental damages. Throughout the chapter, the term *short-run private surplus* (SRPS) refers to the sum of consumer surplus, operating profits, and tax revenue. Total welfare can be expressed as a function of the emissions tax τ:

$$W_m(\tau) \equiv \underbrace{CS_m(\tau)}_{\text{Consumer surplus}} + \underbrace{\Pi_m(\tau)}_{\text{Profits}} + \underbrace{T_m(\tau)}_{\text{Tax revenue}} - \underbrace{\Phi_m(\tau)}_{\text{Damages}}.$$

There are $K_m + 1$ goods in this market, each indexed by k. Consumption of goods $k = 1, \ldots, K_m$ generates carbon emissions e_k per unit consumed. The other product, indexed by $k = 0$, is a composite consumption good representing the "outside option"; it is assumed that this composite good has a unitary price and does not generate consumption externalities. Furthermore, no emission abatement technologies exist in the short run, so that all abatement is achieved through quantity reductions.

[5] I make no assumptions about how firm profits are distributed to individuals or how tax revenues are recycled, for which reason these elements are kept in separate accounts.

Consumer surplus. There are N_m identical consumers with quasi-linear utility

$$U\left(x_0, x_1, \ldots, x_{K_m}\right) = \alpha x_0 + \sum_{k=1}^{K_m} u_k\left(x_k\right),$$

where x_k represents the quantities consumed and α determines the marginal utility of income. Under standard assumptions about the utility function, utility maximization implicitly defines demands for each good as $x_k^* = x_k\left(\boldsymbol{P}_m, y\right)$, where $\boldsymbol{P}_m = \{1, p_1, \ldots, p_k\}$ is the vector of prices and y is the consumer's income level.

Let $q_k\left(\tau\right)$ and $p_k\left(\tau\right)$ represent the equilibrium quantity and price with emissions tax τ. Then, the aggregate money-metric consumer surplus in this market can be represented as

$$CS_m = \frac{N_m}{\alpha} \sum_{k=1}^{K_m} u_k\left(\frac{q_k\left(\tau\right)}{N_m}\right) - \sum_{k=1}^{K_m} p_k\left(\tau\right) q_k\left(\tau\right) + N_m y. \tag{2.1}$$

Operating profits. There are J_m firms supplying K_m emission-generating goods; the outside option is competitively supplied and can be abstracted from profit considerations. Each product k is subject to a uniform sales tax r and a product-specific lump-sum fee ι_k. Thus

$$p_k = \left(1 + r\right) \tilde{p}_k + \iota_k,$$

where $\tilde{p}_k$ is the pre-tax price received by the firm. In this framework, fees ι_k are infrastructure costs paid by the consumer. In aviation, these costs refer to airport and security services provided by entities other than the airlines. These fees are not taxes, but must be included in the model because they create a significant wedge between the ticket price paid and the amount received by airlines, affecting pricing behavior.

Let $\tilde{p}_k(\tau)$ represent equilibrium pre-tax prices and c_k represent the respective marginal operating costs of production, here assumed constant. Then, total operating profit[6] can be written as

$$\Pi_m(\tau) = \sum_{k=1}^{K_m} \left[\tilde{p}_k(\tau) - c_k - \tau e_k\right] q_k(\tau). \tag{2.2}$$

Tax revenue. There are two sources of tax revenue: a sales tax r and an emissions tax τ. Total tax revenue, then, is simply given by

$$T_m(\tau) = \tau \sum_{k=1}^{K_m} q_k(\tau) e_k + r \sum_{k=1}^{K_m} q_k(\tau) \tilde{p}_k(\tau). \tag{2.3}$$

Environmental damage. Each unit of emission produces a constant environmental damage ϕ. For carbon emissions, these can be understood as the present value of the stream of future damages, i.e., the social cost of carbon. Under this setting, environmental damages in this market are given by

$$\Phi_m(\tau) = \phi \sum_{k=1}^{K_m} q_k(\tau) e_k. \tag{2.4}$$

Combining all four components yields the expression for social welfare in a market:

$$W_m(\tau) = \sum_{k=1}^{K_m} \left\{ \frac{N_m}{\alpha} u_k\left(\frac{q_k(\tau)}{N_m}\right) - \left[\iota_k + c_k + \phi e_k\right] q_k(\tau) \right\} + N_m y \tag{2.5}$$

As usual, tax revenues are transfers and get canceled out. Thus, social welfare is a function of the utility derived from each good and the private and external costs of providing these goods. Note that the rightmost term in (2.5) is constant and does not affect the assessment of welfare changes.

[6]These are, in fact, variable operating profits. Entry and exit decisions driven by tax changes are not considered in this analysis, for which reason I omit fixed costs.

2.3.3 Marginal welfare effects

With social welfare defined, it is useful to evaluate how welfare and its components change given a marginal change in tax τ. In this exercise, I assume a new equilibrium exists but remain agnostic about the specific changes in price ($\frac{dp_k}{d\tau}$) and quantity ($\frac{dq_k}{d\tau}$). Differentiating each component with respect to τ yields:[7]

$$\frac{dCS_m}{d\tau} = -(1+r)\sum_{k=1}^{K_m} q_k \frac{d\tilde{p}_k}{d\tau} \tag{2.6}$$

$$\frac{d\Pi_m}{d\tau} = \sum_{k=1}^{K_m} \left\{ \mu_k \frac{dq_k}{d\tau} + \left[\frac{d\tilde{p}_k}{d\tau} - e_k \right] q_k \right\} \tag{2.7}$$

$$\frac{dT_m}{d\tau} = \sum_{k=1}^{K_m} \left[q_k e_k + \tau e_k \frac{dq_k}{d\tau} \right] + r \sum_{k=1}^{K_m} \left[q_k \frac{d\tilde{p}_k}{d\tau} + \tilde{p}_k \frac{dq_k}{d\tau} \right] \tag{2.8}$$

$$\frac{d\Phi}{d\tau} = \phi \sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}. \tag{2.9}$$

Equation (2.6) uses the first order optimality condition $u_k'(x_k^*) = \alpha p_k$, for $k \in \{1 \dots K_m\}$. In (2.7), the term $\mu_k \equiv \tilde{p}_k - c_k - \tau e_k$ denotes the operating markup. In the same equation, $\frac{d\tilde{p}_k}{d\tau} - e_k$ captures the marginal change in markup; this term is equal to zero when there is complete tax pass-through. Combining (2.6)–(2.9) gives the expression for marginal welfare change:

$$\frac{dW_m}{d\tau} = \sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k + (\tau - \phi) e_k \right] \frac{dq_k}{d\tau}. \tag{2.10}$$

In equation (2.10), the terms within square brackets demonstrate how all three market imperfections affect welfare. In particular, the terms referring to sales tax and markup distortions are analogous to the usual Harberger triangle terms (Kleven, 2020). For these terms, the welfare losses are equal to the "mechanical variation" in tax revenue and operating profit (i.e., holding prices

[7] For the remainder of this section, function arguments are omitted for notation clarity.

constant). The third term inside the square brackets refers to the environmental externality and its correction mechanism.

The standard Pigouvian taxation arises as a special case in (2.10), when there is no sales tax ($r = 0$) or market power ($\mu_k = 0$). In this case, the usual prescription applies: setting the environmental tax to marginal damages ($\tau = \phi$) maximizes social welfare.[8] In contrast, if $\tau = \phi$ but other distortions exist, then

$$\left.\frac{dW_m}{d\tau}\right|_{\tau=\phi} = \sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k\right] \frac{dq_k}{d\tau},$$

which is likely negative because $\frac{dq_k}{d\tau} < 0$ for most goods, if not all. This demonstrates that, when other distortions exist, the second-best tax is smaller than the standard Pigouvian tax.

2.3.4 Marginal abatement cost and optimal tax

Equations (2.6)–(2.8) can be used to represent the marginal change in SRPS following a change in τ. Combining the change in SRPS with the marginal change in aggregate emissions ($\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}$), the marginal abatement cost (MAC) is given by

$$MAC\left(\tau\right) \equiv \tau + \frac{\sum_{k=1}^{K_m} \left[r\tilde{p}_k + \mu_k\right] \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}}. \tag{2.11}$$

Recall that it was assumed no abatement technologies are available. Thus, the abatement costs in this framework refer to the private welfare losses following emission reductions induced by environmental tax τ.

[8] With concave utility, function W is locally concave at $\tau = \phi$ with no other distortions. A sufficient condition of global concavity is $\frac{d^2 q_k}{d\tau^2} \geq 0$ for all goods; this condition cannot be guaranteed for every equilibrium adjustment process, however.

Equation (2.10) can be used to characterize the optimal environmental tax with other distortions. Assuming W is globally concave (see footnote 8), at the optimal tax τ^* it follows that

$$\tau^* = \phi - \underbrace{\left\{ \frac{\sum_{k=1}^{K_m} \mu_k \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}} + \frac{\sum_{k=1}^{K_m} r\tilde{p}_k \frac{dq_k}{d\tau}}{\sum_{k=1}^{K_m} e_k \frac{dq_k}{d\tau}} \right\}}_{\text{Tax wedges}}, \tag{2.12}$$

thus making explicit how non-environmental distortions create a wedge between the optimal environmental tax and marginal damages. In particular, this wedge is determined by the marginal changes in each distortion relative to the marginal change in emissions. Throughout the text, I refer to these two components as *tax wedges*, measured in \$/ton CO_2. Combining (2.11) and (2.12), a standard result follows: $MAC\left(\tau^*\right) = \phi$. That is, the MAC at the optimal emission level is equal to marginal damage—even though the optimal tax rate is smaller than the marginal damage.

Note that tax optimality in this framework is expressed in a second-best sense. Besides the constraints of existing distortionary taxes and market power, this characterization also assumes a uniform tax. In theory, a product-specific tax schedule would weakly improve the outcome of this market, leading it closer to the first-best equilibrium. Nevertheless, it is difficult in practice to implement taxes that are specific to each firm or product, especially for input taxes such as the one studied in this chapter; that is why I emphasize uniform taxes.

2.4 A model of commercial aviation

Marginal welfare analyses can shed light on the effects of small changes in the jet fuel tax. These effects can be approximated using a sufficient statistics approach, as described in section 2.6.1. To estimate the optimal tax, however, it becomes necessary to evaluate non-marginal tax changes—this is where the sufficient statistics approach proves limited (Kleven, 2020). Non-marginal changes can be estimated by parameterizing the market equilibrium with structural modeling. In this section, I outline a model for US domestic aviation. This model builds on previous studies on the aviation sector. Most notably, it draws from the models used in Berry et al. (2006), Berry and Jia (2010), Aguirregabiria and Ho (2012), and Pagoni and Psaraki-Kalouptsidi (2016).

2.4.1 Definitions

In this model, time is discrete and each period represents a quarter. A *locations* is a city or metropolitan area with one or more airports. A *market*, indexed by m, is a directional pair of locations (of the form $origin \rightarrow destination$). This market definition follows Berry et al. (2006) and Aguirregabiria and Ho (2012). Pagoni and Psaraki-Kalouptsidi (2016) have a similar approach, but define locations as a cluster of airports within a radius. In contrast, other studies have defined markets as directional airport pairs (e.g. Borenstein, 1989; Ciliberto & Tamer, 2009; Berry & Jia, 2010). In this chapter, I define markets over locations because doing so allows consumers in large metropolitan areas to choose between airports; this allows the model to capture competition between flights departing from airports in close proximity.

A *segment* is an ordered pair between two airports. A *route* r is a sequence of up to four segments forming a round trip. Routes are represented by a four-tuple $(a_o, a_{c1}, a_d, a_{c2})$ of airports: the origin, the outbound connection (if any), the destination, and the inbound connection (if any).[9]

A *product* in this industry is a route r operated by airline i at time period t. For simplicity of notation, let $k = (r, i, t)$ index products; this short-hand definition follows Aguirregabiria and Ho (2012). Let $\mathcal{K}$ represent the set of available products, using subscripts to indicate partitions. For example, $\mathcal{K}_{mt}$ is the set of products in market m at time t, while $\mathcal{K}_{imt}$ contains only airline i's products in $\mathcal{K}_{mt}$.

2.4.2 Consumers

There are N_m consumers in market m. At each period, consumers decide to purchase at most one of the products available in this market. For consumer n, purchasing product k yields a (money metric, indirect) utility of

$$u_{nk} = \underbrace{X_k^D \beta^D - \alpha p_k + \xi_k}_{V_k} + \nu_n(\lambda) + \lambda \epsilon_{nk},$$

where X_k^D is a vector of observed characteristics for product k, p_k is the ticket price, and ξ_k represents unobserved (in the data) product characteristics; α, λ, and β^D are model parameters. To simplify notation, $V_k \equiv X_k^D \beta^D - \alpha p_k + \xi_k$ represents the average consumer surplus for product k. The average surplus for the choice of not purchasing a product, indexed by $k = 0$, is normalized to zero. This specification is similar to Pagoni and Psaraki-Kalouptsidi (2016).

[9] This model does not consider flights within the same city, with disjoint segments, or with more than one connection each way. In the data used for estimation, these excluded flights correspond to less than 3% of all domestic enplanements.

Consumer-specific tastes are represented by the additive error term $\nu_n(\lambda) + \lambda\epsilon_{nk}$, which yields the nested logit discrete choice model (McFadden, 1978). All flights are grouped in a single nest, denoted by g. The outside option is the single choice available in a separate nest. In this specification, $\nu_n(\lambda)$ is constant across all products and accounts for the correlation of tastes across flights (but not with the outside option). The term ϵ_{nk} represents the consumer-specific taste for product k. The distribution of the error term is determined by parameter $\lambda \in [0, 1]$. When $\lambda = 1$, there is no correlation of random tastes across flights, and the model becomes the standard logit model. When λ approaches 0, the correlation of random tastes goes to 1.

Among all options available in a market, a consumer will choose the one that yields the highest indirect utility. Assuming the distribution of the error is Type I Extreme Value, it is possible to integrate over the error term to derive the probability of a consumer choosing product k. This probability corresponds to the expected market share of the product, s_k. Then, the aggregate demand of k in market m is $q_k = s_k \times N_m$. The share of product k within all flights follows the standard logit demand form, given by

$$Pr(u_{nk} \geq u_{nj}, \forall j \in \mathcal{K}_{mt}) \equiv s_{k|g} = \frac{\exp(V_k/\lambda)}{\sum_j \exp(V_j/\lambda)}. \tag{2.13}$$

Define $D_g \equiv \sum_j \exp(V_j/\lambda)$, which can be interpreted as a measure of the expected utility of choosing to purchase a product in nest g. Then, the probability of choosing nest g is given by

$$s_g = \frac{D_g^\lambda}{1 + D_g^\lambda}, \tag{2.14}$$

where the first term in the denominator comes from $\exp(V_0/\lambda) = 1$, representing

the outside option nest. Combining (2.13) and (2.14), it follows that

$$s_k = s_{k|g} \times s_g = \frac{\exp\left(V_k/\lambda\right)}{D_g^{1-\lambda}\left(1 + D_g^{\lambda}\right)}. \tag{2.15}$$

Under the nested logit specification, the expected consumer surplus in market m at period t is given by (Train, 2009)

$$CS_{mt} = \frac{1}{\alpha} \ln\left(1 + D_g^{\lambda}\right) + \kappa, \tag{2.16}$$

where κ is a constant term that is eliminated when evaluating welfare changes.

2.4.3 Airlines

There are I airlines, each indexed by i. The set of airlines is fixed and exogenously given. In each period and market, airlines maximize operating profits by setting prices for each route they operate in that market. When setting prices, airlines take as given a vector of exogenous demand, cost variables, and the set of routes they operate. Defining the set of routes as given has an important implication for this chapter, as it limits all analyses to short-run effects only. Over a longer horizon, airlines make plans that affect their networks and the markets they serve. Beyond dynamic profit maximization, these decisions take into account strategic considerations and numerous technical and non-technical constraints (Belobaba et al., 2015). Modeling airline decisions to assess long-term network changes is difficult, for which reason long-run effects are left outside the scope of this chapter.

Besides the sales tax (r), ticket prices also include product-specific fees (ι_k); these fees are regarded as the cost of infrastructure services paid for by the con-

sumer. Taxes and fees create a wedge between the ticket price p_k, which observed in the data, and the price received by the airline $\tilde{p}_k$ (see section 2.2 for an overview of these charges). Prior literature has largely overlooked this feature of the sector. Nevertheless, most studies are interested in structural characteristics, such as hub premia (Borenstein, 1989, 1991; Ciliberto & Williams, 2010), or large sector shocks, such as mergers (Ciliberto & Tamer, 2009; Aguirregabiria & Ho, 2012), for which the wedges between prices paid and received are likely inconsequential. In contrast, the present chapter is concerned with price adjustments from smaller changes in costs; for this reason, explicitly including taxes and fees in the model is essential for better capturing pricing decisions.

Firms maximize variable operating profits by choosing pre-tax prices (or base fares) for each flight they operate in a market. Thus, the pre-tax price vector chosen by an airline, $\tilde{P}_{imt} = (\tilde{p}_{k_1}, \tilde{p}_{k_2}, \ldots, \tilde{p}_{k_n}), k_1, k_2, \ldots, k_n \in \mathcal{K}_{imt}$, maximizes

$$\Pi_{imt} = \sum_{k \in \mathcal{K}_{imt}} (\tilde{p}_k - c_k) s_k.$$

The marginal cost per passenger is given by $c_k = \tilde{c}_k - (w_t + \tau) f_k$, where $\tilde{c}_k$ is the constant marginal cost per passenger excluding fuel costs, w_k is the jet fuel cost per gallon, and f_k is the volumetric fuel consumption per passenger.

The first-order optimality condition for each product k is given by

$$s_k + \sum_{j \in \mathcal{K}_{imt}} (\tilde{p}_j - \tilde{c}_j - (w_j + \tau) f_j) \frac{\partial s_j}{\partial \tilde{p}_k} = 0. \tag{2.17}$$

The resulting pre-tax price and share vectors, $\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}$ and $\boldsymbol{S}_{\boldsymbol{mt}}$, satisfy a Nash-Bertrand equilibrium. Stacking all first-order conditions from (2.17), the market equilibrium is a solution to the system of equations

$$\boldsymbol{\mu}_{\boldsymbol{mt}} \equiv \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} - \boldsymbol{C}_{\boldsymbol{mt}} = -J_{mt}^{-1} \boldsymbol{S}_{\boldsymbol{mt}}, \tag{2.18}$$

where $\boldsymbol{\mu_{mt}}$ is the vector of operating markups, $\boldsymbol{C_{mt}}$ is the vector of marginal operating costs, and J_{mt} is a matrix with partial derivatives of quantities with respect to prices, multiplied element-wise by an indicator matrix of product ownership (that is, cell ij is equal to 1 if the products i and j are offered by the same airline and 0 otherwise).[10]

2.5 Data

Data sources. The data set used in this research combines seven data sources from four providers. US aviation data are sourced from the Bureau of Transportation Statistics (BTS) of the US Department of Transportation. I query four BTS databases in this chapter. First, the Origin and Destination Survey (DB1B) provides quarterly data on domestic air travel—including origins, destinations, connections, and ticket prices—based on a 10% sample of all tickets (BTS, 2019d). Second, Table T-100 of the Form 41 Traffic Database contains monthly data on air travel operations by segment, aircraft, and airline (BTS, 2019b); from this table, I collect the number of available seats, passengers transported, and ramp-to-ramp time by segment, aircraft model, and airline; I also collect the use share of each aircraft model by segment and airline. Third, I gather data on the number of departures and delays by segment and airline from the On-Time Performance Database (BTS, 2019c). Fourth, I collect aggregate operating revenues and costs by airline and quarter from the Form 41 Financial Database, Schedule P-1.2 (BTS, 2019a).

Average fuel burn by aircraft model and stage length (i.e., flight segment dis-

[10] Appendix A.1.3 presents details on how to construct this matrix and solve equilibria given by equation (2.18).

tance) are collected from the International Civil Aviation Organization's carbon emissions calculator documentation (ICAO, 2018). Jet fuel prices come from the US Energy Information Administration (EIA, 2018); I collect monthly prices of sales to end users by region. Finally, city and metropolitan area populations are collected from the US Bureau of Economic Analysis' Regional Economic Accounts (BEA, 2019).[11]

Sample criteria. The sample used in this research is constructed based on the four quarters of 2018. The data set includes all 73 cities or metropolitan areas in the contiguous US with at least 50,000 passengers surveyed in 2018, which jointly account for 92% of all domestic traffic. These locations include a total of 98 airports. Figure 2.4 displays the geographical distribution of locations and their traffic, aggregating all airports in metro areas. This map shows that most traffic is concentrated on both coasts and in dense urban areas scattered around the country.

Following the data reliability criteria adopted in the literature (Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016), itineraries in the DB1B data are excluded if any of the following conditions apply: (i) is operated by a non-US airline; (ii) is not a round trip; (iii) has more than one stop in either direction; (iv) has fare credibility issues flagged by the BTS; (v) has extreme fares (below $50 or above $3,000); or (vi) has fewer than 3 tickets surveyed in a quarter.

Itineraries selected from the DB1B are assigned to the reporting airline, which is the same as the operating and ticketing airlines in the majority of the

[11] Additional details on the construction of the data set are presented in Appendix A.3.

Figure 2.4: Passenger traffic in cities and metro areas included in the data set.

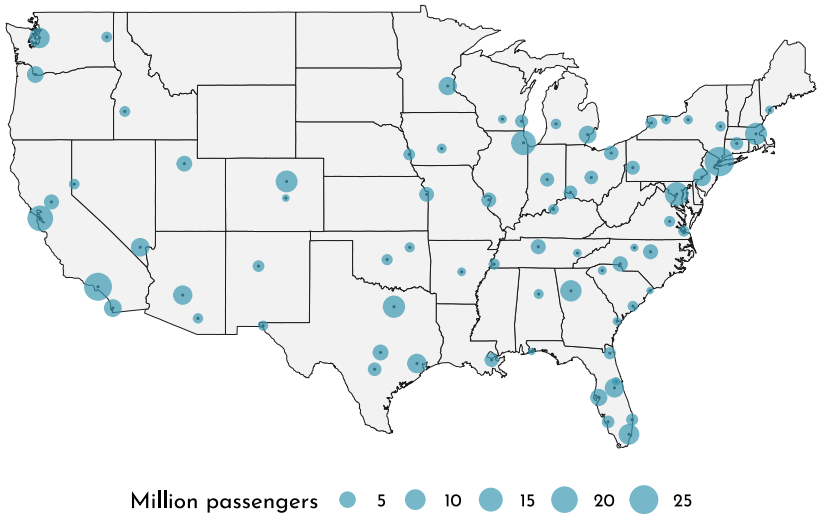

Source of data: US Bureau of Transportation Statistics. Note: traffic is measured in passengers enplaned in domestic flights from all airports within a city or metro area.

cases. Moreover, small and regional carriers with exclusive service for, or acquired by, another airline are grouped with their controlling airline.

Covariates. The resulting data set has 267,967 observations, each representing one product. These products are offered by 10 airline groups in 5,018 markets. This data set includes approximately 48% of all domestic passenger-trips surveyed in DB1B. Figure 2.5 shows the distribution of traffic across different markets. Lines connecting locations indicate the total number of passengers flying round trips, with thicker lines indicating a higher number of passengers. These lines connect only the end points of a market and are not a representation of the actual routes (i.e., they do not show connections). This map highlights the fact that the largest markets connect dense urban areas on the same coast (such as Los Angeles–San Francisco and New York–Miami), while a few are coast-to-coast, or from a coast to a large city in the middle of the country (such as

Figure 2.5: Passenger traffic between cities and metro areas.

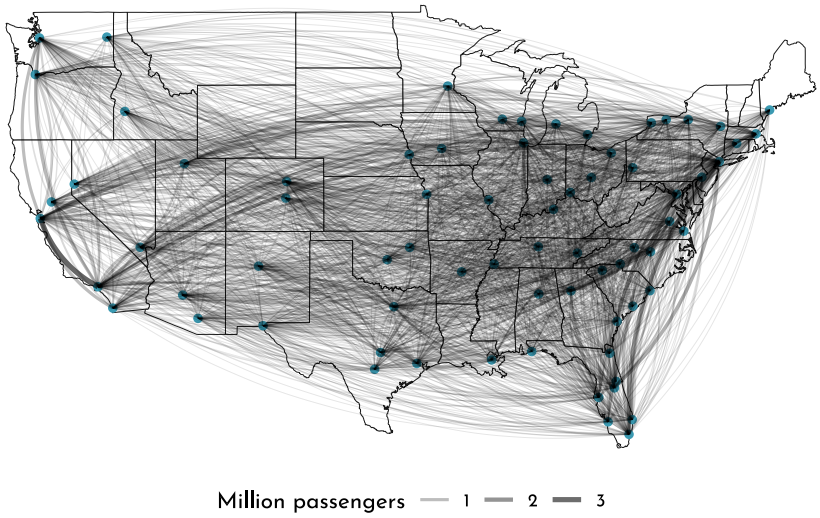

Notes: each segment connects only the end points (city or metro area) of a round trip, regardless of any connections in between. Traffic measures the number of passengers enplaned in round-trip flights between any airport at the end points in either direction.

Chicago, Denver, and Houston).

Table 2.1 presents summary statistics for the covariates used in this analysis. Some covariates require additional details. *Shares* are calculated based on *market size*, which is defined as the geometric mean of the origin and destination populations (as in Berry and Jia (2010) and Pagoni and Psaraki-Kalouptsidi (2016)). *Passengers* is the number of passengers for a product surveyed in DB1B, multiplied by 10 (the survey weight). *Market distance* is the great circle distance between origin and destination airports. *Connection extra distance* is the travel distance added by having connections; it is calculated as the difference between total travel distance and twice the market distance. *Departures per week* is assigned to the minimum number of departures across each segment in a route. *Delayed flights* indicates the percent of departures in the previous quarter with delays above 15 minutes. *Destinations from origin* indicates the number of des-

Table 2.1: Summary statistics.

Statistic	Mean	St. Dev.	Min	Median	Max
Share (%)	0.01	0.06	0.0002	0.002	2.04
Share within nest (%)	7.37	15.41	0.01	1.41	100.00
Passengers	509.25	2,261.27	30	60	62,640
Market Size ($\times 10^{-6}$)	3.47	2.40	0.43	2.76	16.02
Price ($)	479.06	157.41	56.97	465.33	2,041.33
Number of stops	1.60	0.64	0	2	2
Market distance (miles)	1,371.16	640.40	67.13	1,256.22	2,724.08
Connection extra distance (miles)	273.36	295.94	0	181.0	2,993
Departures per week	15.85	13.54	0.08	13.00	145.69
Delayed flights (%)	17.28	5.77	0.00	16.88	100.00
Destinations from origin	25.09	19.40	1	18	84
Airlines in market	5.38	1.48	1	5	9
Rivals' products in market	57.34	69.44	0	34	604
Rivals' % of nonstop flights	5.00	9.44	0.00	2.53	100.00
Potential legacy entrants	0.05	0.25	0	0	3
Potential LCC entrants	3.01	1.06	0	3	5
Compl. segment density ($\times 10^{-3}$)	48.69	28.42	0.02	44.25	256.40
Fuel expenditure ($/avail. seat)	97.46	31.77	6.44	94.93	276.49
Total ramp-to-ramp time (h)	8.36	2.75	1.50	8.03	18.40
Observations (products)			267,967		
Routes			103,720		
Time periods (quarters)			4		
Airlines			10		
Markets			5,018		
Cities or metropolitan areas			73		
Airports			98		

tinations an airline serves from the origin airport. *Total ramp-to-ramp time* sums the time taken in each segment of a route.

Seven covariates are used as excluded instruments for estimation, as explained in section 2.6. This set of instruments was chosen based on previous studies on the aviation sector (especially Berry et al., 2006; Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016), Five of these instruments measure the degree of competition: the number of *airlines in market*, the number of *rivals' products in market*, the percentage of those rivals' products

that are nonstop flights, and the number of legacy and low-cost *potential entrants*. Potential entrants are identified as airlines currently not offering flights in a market but operating in the origin or destination cities (a definition similar to that used in Goolsbee and Syverson (2008)). *Complementary segment density* indicates the sum of passengers from other markets who are transported on each segment of a route; this variable is a measure of the scale of operations along a route. Finally, *fuel expenditure* is the sum of fuel consumption per available seat along each segment, multiplied by the respective fuel price.[12]

2.6 Empirical approach and estimation

This chapter undertakes two complementary approaches to assess the impact of a carbon tax in aviation. First, using estimated sufficient statistics, I evaluate marginal changes to prices, quantities, and welfare following a small incremental change in the current jet fuel tax. As outlined in section 2.3, these marginal changes measure the relative size of market power and tax distortions. Without requiring extensive structural assumptions, sufficient statistics then indicate when an increase in the jet fuel tax is welfare-improving. This first approach, however, has limited use for deriving optimal taxes (Kleven, 2020). To calculate the optimal tax, the second approach relies on estimated structural parameters to characterize non-marginal changes to market equilibria. Specifically, I simulate counterfactual equilibria for various levels of carbon tax to find the optimal tax level and analyze its components.

[12]This covariate accounts for the length of each segment, as well as fuel efficiency and use share of each aircraft model in each segment of a route. Jet fuel prices are assigned based on the departing airport in each segment. See Appendix A.3 for details.

In this section, I describe the methods used for estimation of parameters and present the estimation results. The welfare analyses using estimated statistics and parameters are presented in the next section.

2.6.1 Sufficient statistics

The right-hand side on marginal change equations (2.6)–(2.9) shows three types of terms that are not directly observed in the data: product-specific markups (μ_k), marginal pre-tax price changes ($\frac{d\tilde{p}_k}{d\tau}$), and marginal quantity changes ($\frac{dq_k}{d\tau}$). These product-specific terms cannot be directly estimated from the data for two reasons. First, there are no publicly available data on product-level costs to inform markup calculations. Second, there have not been any changes to the jet fuel tax in the data set period. To circumvent issues with product-specific estimation, I focus instead on aggregate marginal effects using averages across products and markets.

For markups, I use airline average operating revenue and costs per revenue passenger-mile (RPM). These system averages, calculated using Form 41 Financial data, are widely used in the aviation sector to analyze airline performance. Average markups per airline are shown in Table 2.4. Since these metrics are relative to flight distance, in this approach different flights from the same airline are assigned different markups.

Estimating marginal changes in equilibrium outcomes requires further assumptions. Following the usual route taken in the sufficient statistics approach, I express marginal changes in terms of elasticities. In doing so, I make the simplifying assumption of a sequential shock propagation: a marginal increase in

fuel costs leads to a proportional change in air travel fares, which then affects equilibrium quantities. For small changes in the fuel tax, I assume that changes to markup and non-fuel costs are negligible, so that additional costs are fully passed on to consumers.[13] Then, the first shock can be represented by an elasticity of pre-tax ticket prices with respect to fuel costs:

$$\frac{\partial \tilde{p}_k}{\partial \tau} = \frac{\partial \tilde{p}_k}{\partial F_k}\frac{\partial F_k}{\partial w_k}\frac{\partial w_t}{\partial \tau} = \frac{\tilde{p}_k}{w_k}\eta_k, \tag{2.19}$$

where $F_k = (w_k + \tau) f_k$ is the fuel expenditure per revenue passenger (other variables are defined in section 2.4.3) and $\eta_k \equiv \frac{\partial \ln \tilde{p}_k}{\partial \ln F_k}$ is a pass-through elasticity. Under the assumption about markups, η_k is equal to the ratio of fuel cost share to pre-tax price and can be calculated from the data set.[14]

In equilibrium, changes in quantities are functions of the vector of all price changes. As in the previous case, it is possible to express these changes in terms of elasticities

$$\frac{dq_k}{d\tau} = \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j}\frac{dp_j}{d\tilde{p}_j}\frac{d\tilde{p}_j}{d\tau} = (1+r)\,\eta_k \frac{q_k}{w_k}\sum_{j \in \mathcal{K}_{mt}} \varepsilon_{kj}\frac{\tilde{p}_j}{p_j}, \tag{2.20}$$

where ε_{kj} is the elasticity of demand for product k with respect to the price of product j. Estimating ε_{kj} for every pair of products in each market is infeasible with the data set at hand. Instead of working with individual own- and cross-price elasticities, I focus on aggregate quantities in each market. To do so, I estimate how changes in average fuel costs affect the aggregate demand for flights (Q_{mt}). This approach assumes that, for small changes in the fuel tax, the relative market shares of flights are not affected; thus, quantities change at the

[13] In contrast, an alternative approach would be to estimate the absolute pass-through rate of cost shocks; however, such an approach would be complicated by fuel cost hedging strategies, which are unlikely to be representative of responses to a certain and (assumed) permanent tax shock.

[14] This can be demonstrated as follows: consider pre-tax prices can be decomposed as fuel costs, non-fuel costs, and markup: $\tilde{p}_k = F_k + \tilde{c}_k + \mu_k$. Then, holding $\tilde{c}_k$ and μ_k constant, it follows that $d \ln \tilde{p}_k = \frac{F_k}{\tilde{p}_k} d \ln F_k$.

same proportion: $dq_k = s_{k|g} dQ_{mt}$. Flights in a market can then be treated as a composite flight $Q_{mt} = \sum_{k \in \mathcal{K}_{mt}} q_k$, with average prices ($\tilde{p}_{mt}$), fuel use ($f_{mt}$), and markups ($\mu_{mt}$) weighted by market shares ($s_{k|g}$). Changes in aggregate quantities can then be expressed as[15]

$$\frac{dQ_{mt}}{d\tau} = \sum_{k \in \mathcal{K}_{mt}} \frac{dq_k}{d\tau} = (1 + r) \frac{\eta_{mt}}{w_t} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon, \tag{2.21}$$

where $\eta_{mt} = f_{mt} w_{mt} / \tilde{p}_{mt}$ is the market-average cost shock elasticity and $\varepsilon = \frac{\partial Q_{mt}}{\partial P_{mt}}$ is the elasticity of aggregate demand with respect to the market average ticket price. In this formulation, ε is the sufficient statistic to be estimated.

Estimation. The elasticity of aggregate demand is estimated in a reduced-form approach with the following equation

$$\ln Q_{mt} = \varepsilon \ln P_{mt} + \gamma_{ot}^{(\varepsilon)} + \gamma_{dt}^{(\varepsilon)} + \nu_{mt}^{(\varepsilon)}, \tag{2.22}$$

where $\gamma_{ot}^{(\varepsilon)}$ and $\gamma_{dt}^{(\varepsilon)}$ represent fixed effects for origin and destination locations by quarter; $\nu_{mt}^{(\varepsilon)}$ is an idiosyncratic demand shock. The fixed effects are added to capture demand components affected by characteristics of the end point locations and seasonality. As usual in demand estimation, market price P_{mt} is potentially endogenous. To address this source of bias, I construct instruments based on the aviation literature (e.g. Berry et al., 2006; Pagoni & Psaraki-Kalouptsidi, 2016) that measure variations in competition and costs. Competition instruments include number of airlines, number of products, percent of products that are nonstop flights, number of potential legacy entrants, and number of potential low-cost entrants.[16] The cost-shifting instrument is the average fuel expen-

[15] See Appendices A.1.1 and A.1.2 for details on the derivation of sufficient statistics and how they are used to calculate marginal welfare changes.

[16] As in Goolsbee and Syverson (2008), potential entrants are airlines that operate in any of the end points of a specific market. Hence, their entry costs would be lower in theory because they would have to expand their network to one additional node to enter a specific market.

Table 2.2: Estimates for the price elasticity of aggregate demand.

	OLS ln(Agg. Passengers)	2SLS ln(Agg. Passengers)
ln(Agg. ticket price) [ε]	-1.929	-2.308
	(0.056)	(0.077)
Fixed effects		
Origin city-by-quarter	Yes	Yes
Destination city-by-quarter	Yes	Yes
Observations	267,967	267,967
First-stage F-statistic		1422
First-stage Conditional F-statistic		632

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The conditional F-statistic is calculated following Sanderson and Windmeijer (2016). The first stage regression for 2SLS is reported in Appendix A.2.1.

diture per available seat.

Estimation results. Table 2.2 shows the results of estimating equation (2.22) with Ordinary Least Squares (OLS) and Two-Stage Least Squares (2SLS). Results from both estimators indicate negative elasticities; 2SLS estimates are larger in absolute value. For comparison, Berry and Jia (2010), one of the few papers reporting aggregate price elasticities, find that these elasticities have been increasing in absolute value over time, reflecting structural and demand changes including quality of service and consumer behavior. Based on estimated structural parameters of an industry model, they calculate a sector-aggregate elasticity of -1.55 for 1999 and -1.67 for 2006. The present chapter's estimates are larger than those reported in Berry and Jia (2010). Part of the difference in estimates can be attributed to a continuation of the trends identified in their paper. Another difference comes from the fact that ε is an intra-market elasticity, which does not account for possible substitution across markets; for this reason, this parameter is expected to be larger than a sector-wide elasticity.

2.6.2 Structural parameters

Demand specification. Manipulating equations (2.13)–(2.15), it is possible to derive the estimating equation for the nested logit demand (Berry, 1994)

$$\ln s_k - \ln s_0 = X_k^D \beta^D - \alpha p_k + (1 - \lambda) \ln s_{k|g} + \xi_k, \tag{2.23}$$

from which parameters β^D, α, and λ are estimated. Vector X_k^D includes the following observed product characteristics: (i) service frequency in departures per week, (ii) number of stops, (iii) market distance (i.e., between endpoint cities) and its square, (iv) travel distance added due to connections and its square, (v) percent of delayed departures in the previous quarter, (vi) number of destinations offered by airline from the origin city, and (vii) fixed effects for airline, origin city-by-quarter, and destination city-by-quarter.

Three variables in equation (2.23) are potentially correlated with unobserved characteristics (ξ_k) and are thus endogenous: prices, within-nest shares, and flight frequency (Berry & Jia, 2010). To address this endogeneity, I construct instruments following the aviation literature (Berry et al., 2006; Berry & Jia, 2010; Aguirregabiria & Ho, 2012; Pagoni & Psaraki-Kalouptsidi, 2016). There are four groups of instruments. First, the competition-shifting instruments are: (i) number of airlines in a market, (ii) number of products offered by competitors, (iii) share of competitors' products that are nonstop flights, (iv) number of potential legacy entrants, and (v) number of potential low-cost entrants. Second, (vi) fuel cost per available seat is a cost-shifter. Third, (vii) complementary density along segments measures the number of passengers from other markets transported on the same segments of a route; this instrument indicates the scale of operations that are complementary to a product and affect both costs and frequency of service. The fourth group includes all exogenous variables in equation (2.23).

Supply specification. Using estimated parameters $\widehat{\alpha}$ and $\widehat{\lambda}$, observed prices, and observed market shares, we can calculate predicted marginal operating costs by rewriting equation (2.18) as

$$\widehat{\boldsymbol{C}}_{\boldsymbol{mt}} = \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} + \widehat{J}_{mt}^{-1}\boldsymbol{S}_{\boldsymbol{mt}}.$$

These predicted costs are then used to estimate the supply side of the model.

Understanding the cost structure of aviation services is relevant for a correct specification of the supply-side equation. I present next a brief overview of this structure, as described in Belobaba et al. (2015). Flight operating costs can be mapped onto five categories based on their respective unit of variation. First, there are *costs per block hour*. This category includes all aircraft operating costs, which are directly proportional to the time an airplane is used; it also includes passenger service costs, such as flight attendant wages, entertainment, and food, which are proportional to the duration of a flight. Second, there are *costs per departure*, which are primarily aircraft servicing costs; these include cleaning, fueling, and related ground operations. Third, there are *costs per enplaned passenger*, which account for traffic servicing costs, such as passenger and baggage processing. Fourth, there are *costs per distance*, which reflect primarily fuel costs. Fifth and finally, there are *indirect and overhead costs*; this category includes sales, advertising, management, and other categories that are not clearly mapped to any specific units of the flight operation.

The specification of the cost equation builds on the different categories described above:

$$\widehat{c}_k = \rho F_k + \beta_i^S \text{Ramp-to-ramp}_k + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k, \tag{2.24}$$

where F_k is fuel expenditure per available seat; Ramp-to-ramp$_k$ is the flight duration measured in hours; $\gamma_{i,o}, \gamma_{i,c_1}, \gamma_{i,d}$, and γ_{i,c_2} are fixed effects of each airport

along a route interacted with airline; γ_t is a quarter fixed effect; and ω_k is an idiosyncratic cost shock.

The key parameter in equation (2.24) is ρ: it informs how the implied cost (used for pricing) varies across the different levels of fuel expenditure. The exogeneity of fuel expenditures relies on two factors. First, jet fuel tracks closely crude oil prices, as shown in Figure A.2. Hence, since jet fuel accounts for a small fraction of refined oil products, shocks to jet fuel prices largely reflect shocks that are exogenous to the aviation sector. Second, fleet and network composition affecting fuel use reflect long term decisions which are unlikely to respond quickly to recent fuel price shocks. Thus, concerns over responses to small shocks are mitigated by the fact that the data set considers only a short period. For this reason, the relevant variation in this estimation is cross-sectional. The additional terms in equation (2.24) map the other cost categories described above. Ramp-to-ramp$_k$ is a proxy for costs per block hour, with parameter β_i^S accommodating differences across airlines. A rich set of fixed effects captures how average costs and costs per enplaned passenger and departure vary for each airline at each airport. A time fixed effect captures average variation in costs across quarters.

Estimation. In the aviation literature, there is often a trade-off between the dimensionality of the characteristics space and the estimation procedure. Even though most papers work with large data sets, it is common to use a small set of proxies and dummy variables instead of a flexible set of fixed effects, especially when applying a maximum likelihood or generalized method of moments (GMM) estimator. For example, many papers have used average temperatures and dummies for traditional touristic destinations (e.g. Reiss & Spiller, 1989;

Berry et al., 2006; Berry & Jia, 2010) and dummy variables for whether an airport is slot-controlled or a hub (e.g. Berry & Jia, 2010; Pagoni & Psaraki-Kalouptsidi, 2016). One exception is Aguirregabiria and Ho (2012), which specifies demand and cost equations with several fixed effects; that paper, however, performs separate 2SLS estimations for demand and supply, thus adding the assumption that error terms across both equations are uncorrelated.

Specifications with a high-dimensional characteristic space create additional issues, especially for GMM estimation. Since the joint demand-supply system is nonlinear in parameters α and λ, fixed effects cannot be directly factored out of the equations. Leaving a large number of dummy variables raises computer memory requirements and computation time, both of which can increase exponentially with the number of covariates. Moreover, numerical minimization routines become more computationally challenging when there is a large number of moment conditions (Bennett, Kallus, & Schnabel, 2019).

I address the limitation described above using a method proposed in Conlon and Gortmaker (2020).[17] This method first modifies the estimating equations so that nonlinear terms are absorbed in the left-hand side. Then, fixed effects are factored out using the method of alternating projections (Bauschke, Deutsch, Hundal, & Park, 2003). With these transformations, estimating linear parameters becomes computationally simpler. Furthermore, linear parameters can be expressed as functions of nonlinear parameters. These steps result in a much faster estimation, since the numerical optimization routine only searches over a 2-dimensional space, with linear parameters calculated in the inner loop.

Identification in this estimation relies on the exogeneity of the instruments

[17] A detailed description of the estimation procedure is presented in Appendix A.1.4.

in each equation. These identification assumptions can be arranged in a vector of moment conditions of the form

$$\begin{bmatrix} E\left(Z_k^D \xi_k\right) \\ E\left(Z_k^S \omega_k\right) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{2.25}$$

where Z_k^D is the vector of demand instruments and Z_k^S is the vector of supply instruments (equal to the vector of supply covariates, in this case). After factoring out fixed effects—corresponding to approximately 3,000 dummy variables—these conditions form a system of 22 equations and 16 unknowns. This system is the base of the 2-step GMM estimation used in this chapter. Standard errors are robust and clustered by non-directional city pairs, so as to allow for correlation across markets in opposite directions.

Estimation results. The results for the joint estimation procedure are shown in Table 2.3, with demand and supply coefficients in the left and right columns, respectively. The last row in this Table shows the value of the GMM objective function; this value corresponds to the statistic of the Hansen's test of overidentifying restrictions (Hansen, 1982). The small value of this test statistic provides an indication of the validity of the instruments used.[18]

The key demand coefficients estimated in Table 2.3 are largely in line with the literature. Estimates for α in more recent studies generally vary between -1.36 (Aguirregabiria & Ho, 2012) and -0.45 (Pagoni & Psaraki-Kalouptsidi, 2016). Berry and Jia (2010) differentiate leisure and business travelers; based on 2006 data, they estimate the price parameter as -1.05 for the first group (estimated as 63% of consumers) and -0.10 for the second group. Parameter λ

[18] Appendix A.2.2 presents further evidence of instrument validity in a two-stage least squares estimation.

Table 2.3: Estimation results for the joint demand-supply estimation.

Demand	$ln(s_k/s_0)$	*Supply*	$\widehat{c}_k$
Price ($100) [$-\boldsymbol{\alpha}$]	-0.805	Fuel expenditure/avail. seat [$\boldsymbol{\rho}$]	0.759
	(0.110)		(0.123)
ln(share within nest) [$\mathbf{1-\lambda}$]	0.368	Total ramp-to-ramp time (h)	0.145
	(0.046)		(0.014)
Departures per week	0.036	$\times$ *American*	-0.041
	(0.002)		(0.007)
Number of stops	-0.848	$\times$ *Delta*	0.046
	(0.067)		(0.008)
Market distance (100 mi.)	0.077	$\times$ *United*	-0.018
	(0.015)		(0.008)
Market distance squared	0.0001	$\times$ *Alaska*	-0.014
	(0.0003)		(0.032)
Connection extra distance (100 mi.)	-0.089	$\times$ *JetBlue*	-0.016
	(0.008)		(0.010)
Connection extra distance squared	0.004	$\times$ *Other low-cost*	-0.127
	(0.001)		(0.007)
Share of delayed departures (%)	-0.634		
	(0.099)		
Destinations from origin	0.011		
	(0.002)		
Observations	267,967		
Objective Function	3.653×10^{-5}		

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The demand equation includes fixed effects for airline, origin airport-by-quarter, and destination airport-by-quarter. The supply equation includes fixed effects for quarter and for each route airport-by-airline. Southwest is the base (omitted) airline in the interaction with ramp-to-ramp time.

depends on the nesting assumption; among papers that also group all flights in the same nest, estimates of $(1-\lambda)$ are usually in the 0.3–0.4 range.

Other demand coefficients in Table 2.3 also present the expected signs and are in line with the literature. Consumers value a higher frequency of service, meaning more opportunities for convenient travel times. A larger number of offered destinations increases the value of frequent flier programs, thus making consumers more willing to travel with airlines that have more destination options. Moreover, the positive coefficient on market distance captures the value of air travel.

The results also confirm that consumers dislike connecting flights, insofar as they increase both the number of stops and the total distance traveled; the positive coefficient on extra distanced squared indicates that this marginal disutility decreases with distance. Finally, consumers also avoid those flights that were more frequently delayed in the previous quarter, though this effect is quite small.

The supply side of Table 2.3 shows that a \$1 increase in fuel costs per available seat translates, on average, to a \$0.76 increase in implied costs. It is worth mentioning that ρ is not *per se* a cost pass-through parameter but, instead, a parameter that flexibly captures how marginal costs used for pricing vary. In equilibrium, the realized cost pass-through also depends on each market's demand and structure. Other cost parameters indicate an expected pattern: the cost per block hour and passenger among legacy carriers is generally greater than those among low-cost carriers.

Out-of-sample model predictions. To assess the validity of the model, I compare predicted outcomes with out-of-sample values reported by airlines in the Form 41 Financial database. This database contains quarterly aggregate indicators reported by airlines, including revenue passenger-miles (RPM) and total operating costs and revenues. These data are not used to estimate model parameters, so they are a good candidate for performing a sanity check on the model. In particular, I evaluate the model's ability to generate reported patterns in average revenues, costs, and markups per RPM by airline. These metrics are not only widely used in the aviation sector but also capture the focal point of the model's application: markups.

Table 2.4: Average operating revenue, costs, and markups per revenue passenger-mile (RPM).

	Revenue (¢/RPM)		Cost (¢/RPM)		Markup (¢/RPM)		Market share (%)
Airline	Predicted	Reported	Predicted	Reported	Predicted	Reported	
Southwest	16.13	16.42	11.15	11.31	4.98	5.11	27.88
American	17.27	17.35	13.26	12.28	4.01	5.07	19.26
Delta	18.61	18.37	14.39	12.66	4.22	5.72	18.93
United	16.21	14.70	12.77	12.55	3.44	2.16	14.30
Other LCCs	4.93	9.78	1.28	7.88	3.66	1.90	8.39
JetBlue	12.73	14.28	9.20	11.19	3.53	3.10	5.62
Alaska	11.94	13.70	8.27	9.99	3.67	3.71	5.61
Average	15.38	15.43	11.28	11.34	4.10	4.09	

Notes: monetary values are displayed in cents per revenue passenger-mile. Predicted averages result from the estimated sector model. Reported values are calculated based on the BTS Form 41 Financial database. Market shares are calculated aggregating all markets in the data set.

Table 2.4 shows how model outcomes compare with reported financials, with airlines ordered by aggregate market share (displayed in the rightmost column). The last row shows that predictions for revenue, cost, and markups are remarkably close to the values reported by airlines. For individual airlines, however, the quality of the predictions vary.

Differences in predicted versus reported revenues have two explanations. The first reason is related to selection: the data set used for prediction does not include all domestic markets and, thus, may not be perfectly representative of the average revenue in the whole network. The second reason is that the model does not capture product unbundling, so predicted revenue comes from ticket sales only. In practice, baggage, reservation, and cancellation fees make up a small fraction of airline revenues; however, for low-cost carriers (LCCs) these sources of revenue can represent a large share of total operating revenues (Belobaba et al., 2015; Brueckner et al., 2015). As a consequence, the model has limited ability to reproduce the business of LCCs and is least accurate when predicting their revenues.

Errors in revenue prediction, however, are not fully passed onto markup prediction errors. The reason for this partial correction is that the model predicts costs that rationalize pricing choices via the Nash-Bertrand equilibria. Especially for LCCs, predicted average costs are very low. Nevertheless, these predictions result in markups that are closer to the reported values. Hence, even though these equilibria may not capture all relevant components of pricing decisions, they do reproduce important patterns in the reported data.

2.7 Welfare analyses

Based on the estimated sufficient statistics and structural parameters, this section performs welfare analyses for the introduction of a carbon tax in the US domestic aviation industry. Section 2.7.1 uses sufficient statistics to evaluate marginal impacts from the current equilibrium. Next, section 2.7.2 moves towards non-marginal changes to analyze optimal taxation under distortions. Section 2.7.3 considers the effects of substituting the current sales tax for a revenue-neutral carbon tax. Lastly, section 2.7.4 discusses limitations of these analyses.

The welfare analyses consider three values of the SCC to estimate damages of carbon emissions. The low SCC scenario is set at \$50/ton CO_2, a reference value commonly used in policy and in other studies in the literature. The medium SCC of \$125/ton is taken from a recent estimate by Daniel et al. (2019), which finds declining carbon prices under Epstein-Zin preferences. The high SCC scenario is set at \$230/ton, reflecting an upper bound estimate in that paper. As indicated in section 2.2, each burned gallon of jet fuel emits the equivalent of 13.4 kg of CO_2.

2.7.1 Marginal welfare changes with a carbon tax

Marginal changes to average prices and aggregate quantities are calculated using equations (2.19) and (2.21). Figure 2.6 presents these marginal variations extrapolated linearly for tax values of up to \$100. Though the upper limit of this range is not a small variation, it allows for comparisons with other methods and studies. Note that the current jet fuel tax (4.4 cents/gallon) corresponds to a carbon tax of about \$3.3/ton CO_2, so the baseline does not start at the zero intercept.

Figure 2.6 shows that raising the carbon tax to \$50 would increase prices by 6.5%, on average, and decrease demand by 15%. The changes in demand are one-to-one with the reduction in emissions, as products are aggregated with fixed shares in each market and there is no possibility of substitution. Despite the strong assumptions on the linearity of effects and market aggregation, these predictions are somewhat close to those estimated in Pagoni and Psaraki-Kalouptsidi (2016): an increase of 5.9% in prices and a decrease of 11.2% in demand following a \$50 carbon tax.

Marginal changes to short-run private surplus (SRPS) and damages can be calculated with equations (2.6)–(2.9). The marginal loss of SRPS amounts to \$55M per dollar added to the carbon tax. Based on equation (2.11), this loss translates into a marginal abatement cost (MAC) of \$233/ton CO_2. The tax wedge due to market power is approximately \$184/ton CO_2, while the sales tax wedge is \$49/ton CO_2. The marginal damages avoided per ton of CO_2 correspond to the SCC. Hence, based on the sufficient statistics approach, a marginal increase in the carbon tax above the baseline (\$3.3/ton CO_2) would lead to social welfare losses when the SCC is below \$233/ton CO_2 (the MAC).

Figure 2.6: Predicted changes to prices and quantities from linear extrapolation of marginal effects.

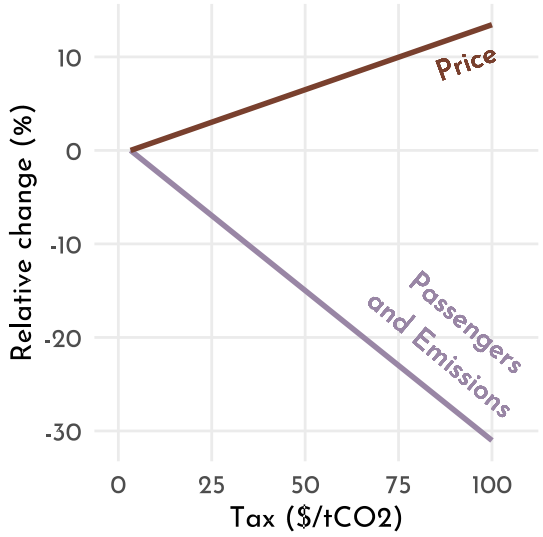

Note: marginal changes are computed using the sufficient statistics approach described in section 2.6.1.

2.7.2 Non-marginal welfare changes and the optimal carbon tax

With a structural model of the sector, it is possible to evaluate non-marginal increases in the jet fuel tax and recalculate the equilibrium outcomes in each market. This is done by increasing costs $\widehat{c}_k$ and solving for prices and shares in the equilibrium equation (2.18).

Panel (a) in Figure 2.7 displays the changes in average ticket prices and aggregate demand and emissions. If a tax of \$50/ton CO_2 were implemented, ticket prices would increase by 4.4%; demand would decrease by 11.9% and emissions by 14.2%. In contrast with the previous subsection, carbon taxes affect products differently: they are more stringent on more polluting flights. For this reason, emission reductions are larger than traffic reductions, as demand for more polluting flights decreases faster. For the \$50/ton CO_2 tax, less pol-

luting flights (those in the highest decile of fuel efficiency) would experience an average price increase of 3.2% and a decrease of 9.9% in passengers. In contrast, more polluting flights (in the lowest decile of fuel efficiency) would face an average price increase of 6.1% and a decrease of 14.7% in passengers.

Translating equilibria changes into welfare consequences, panel (b) in Figure 2.7 shows that social welfare decreases under a low or medium SCC. Hence, no positive optimal tax exists when the SCC is \$50 or \$125. In a high SCC scenario, the optimal tax is \$40/ton CO_2. At only 17.3% of the marginal damage, this optimal tax is much lower than the standard Pigouvian tax prescription. The difference between the \$230 SCC and the \$40 optimal tax indicates a \$190 tax wedge, of which approximately \$140 are due to market power and \$50 to sales tax.

To understand the mechanisms driving these welfare results, panel (d) in Figure 2.7 decomposes variations into private surplus and external damages. Marginal damage avoided is analogous to the social benefit of a carbon tax, whereas marginal loss of SRPS is analogous to the MAC (in terms of private welfare); however, these marginal values are here measured in terms of tax levels rather than emissions. This panel shows that other market distortions lead to a high marginal loss of SRPS. If damages are low relative to these losses, as in the case of low and medium SCCs, social welfare decreases. When marginal damage is higher than the marginal SRPS loss at the baseline, the optimal tax is found at the intersection of these curves—this is the standard marginal cost equals to marginal benefit result.

Panel (c) in Figure 2.7 displays a decomposition of SRPS losses and shows that the burden of the carbon tax is slightly larger for airlines than for con-

Figure 2.7: Predicted changes with the introduction of a carbon tax.

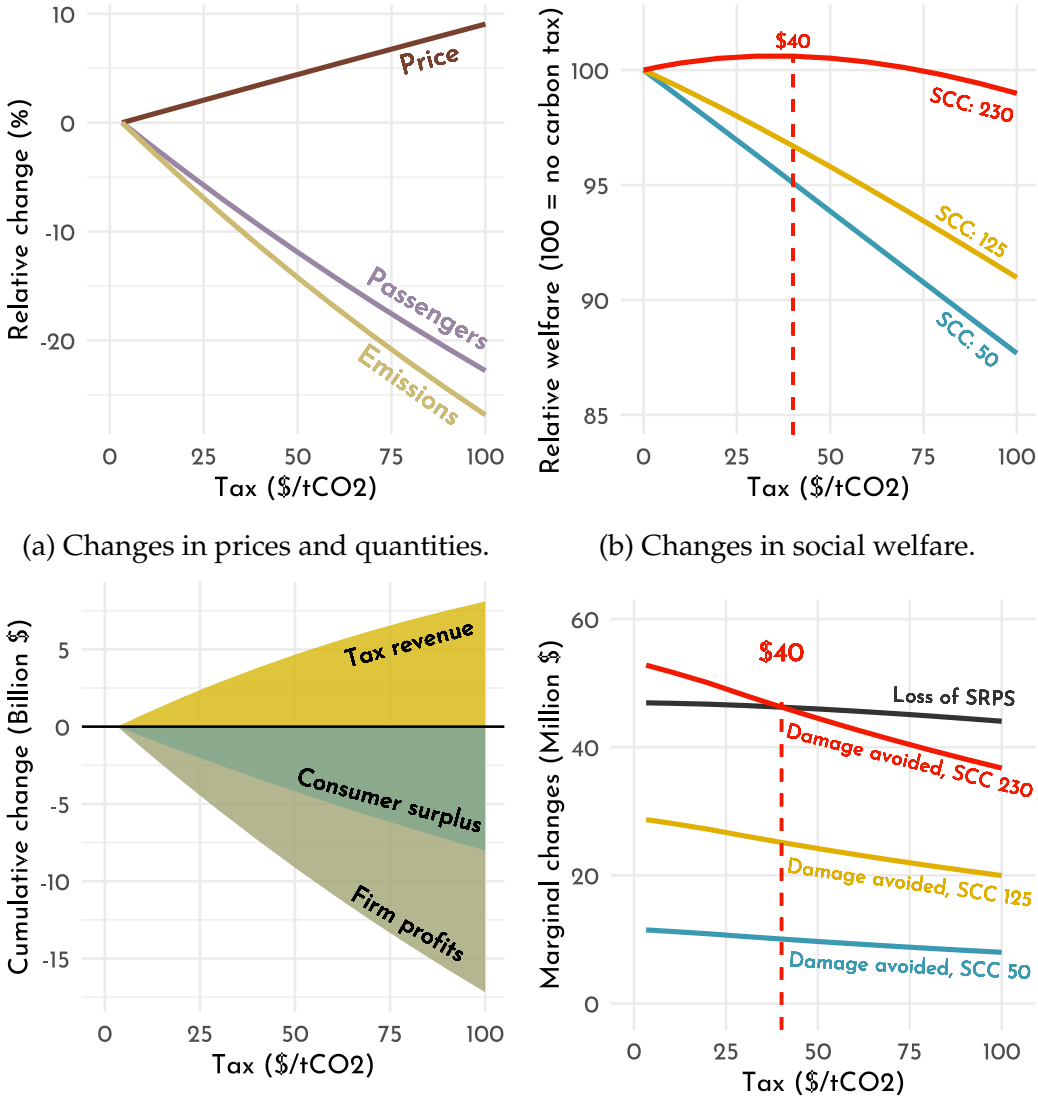

(a) Changes in prices and quantities.

(b) Changes in social welfare.

(c) Decomposition of changes in Short-Run Private Surplus (SRPS).

(d) Marginal changes in social welfare.

Notes: predictions are obtained by re-computing market equilibria with different tax levels using the estimated model, as described in section 2.6.2. In panel (d), *marginal damage avoided* is analogous to the marginal benefit from lower emissions, which depends on the social cost of carbon (SCC) and the emission reductions at each tax level. *Loss of SRPS* (short-run private surplus) is analogous to the private cost of reducing emissions with a carbon tax.

sumers. On average, declining profits account for 54% of the losses and consumer surplus loss for 46%. For a reference carbon tax of \$50/ton CO_2, consumer surplus would fall by \$4.2B and airline profits by \$4.9B. An increase of

Figure 2.8: Marginal abatement cost and tax wedges at various abatement levels.

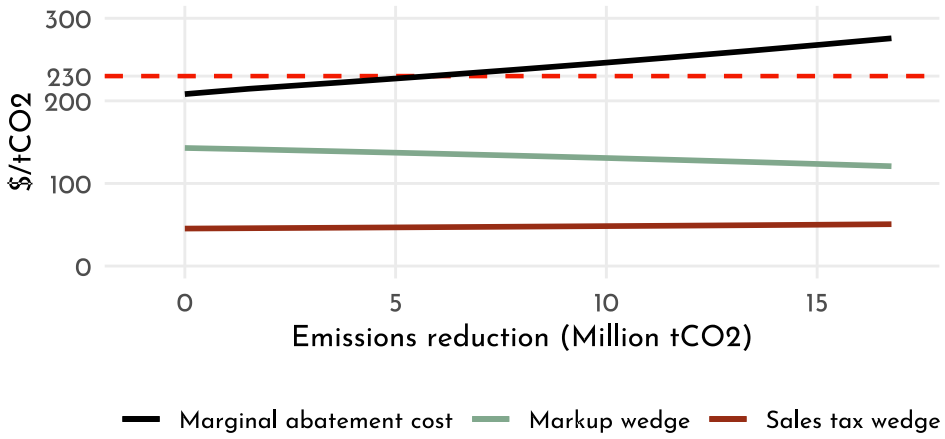

Notes: marginal abatement costs and wedges are defined in section 2.3.4. These calculations are based on predictions using the estimated model, as described in section 2.6.2.

$4.4B in tax revenues, however, could be allocated to partially offset losses to either side.

Changes in SRPS and emissions allow us to evaluate the costs of abating emissions via a carbon tax. Figure 2.8 illustrates how these costs and tax wedges vary with abatement levels. This figure shows that the MAC from the baseline is approximately \$208/ton CO_2. With low and medium SCCs, the private welfare costs of abatement exceed the avoided damages at the margin, thus providing a complementary perspective on the welfare results described above. When damages are high and exceed the initial MAC, a positive optimal carbon tax exists where these curves intersect. Moreover, the wedge due to markups initially accounts for about two thirds of the distortion wedges. As the abatement level increases with a higher carbon tax, incomplete pass-through tends to decrease markups, thus reducing the markup wedge. In contrast, with more expensive tickets due to a higher carbon tax, the distortion from the sales tax increases.

The baseline MAC in this approach is about 10% smaller than the \$233/ton

CO_2 estimated with sufficient statistics. Further examination shows that this difference is primarily driven by differences in market shares, which were held fixed in the previous approach, and in markups, which were set at airline averages instead of being product specific. With more flexibility for market shares to adjust based on externality charges, abatement costs under the structural approach become smaller. Nevertheless, the difference between estimates is relatively small, showing that the results from these methods offer a sanity check on each other.

The value of \$208/ton CO_2 corresponds to the marginal aggregate abatement cost based on a uniform carbon tax applied to all markets. Nevertheless, market power and ticket prices vary across markets, so tax wedges and abatement costs are heterogeneous. Figure 2.9 illustrates the distribution of baseline MACs. Panel (a) in this Figure displays a histogram of the abatement costs per market; panel (b) shows cumulative distributions for the case where markets have equal weights (unweighted) and for the case where markets are weighted by their emissions. These graphs confirm the intuition that while existing distortions vary substantially, MACs are high even among the markets with the lowest abatement costs: the minimum MAC is \$105/ton CO_2. Thus, under the benchmark SCC of \$50/ton CO_2, there are no markets where a positive carbon tax would increase welfare. Emissions with relatively low abatement costs are mostly from large markets connecting dense urban areas, where competition leads to lower markups, fares, and, thus, tax wedges.[19]

The heterogeneity in costs also indicates inefficiencies arising from the use of a uniform carbon tax, as indicated in Section 2.3.4. Even though the second-best uniform tax of \$40/ton CO_2 (under a high SCC) improves aggregate wel-

[19]See Appendix A.4 for details.

Figure 2.9: Distribution of baseline marginal abatement costs (MACs) across markets.

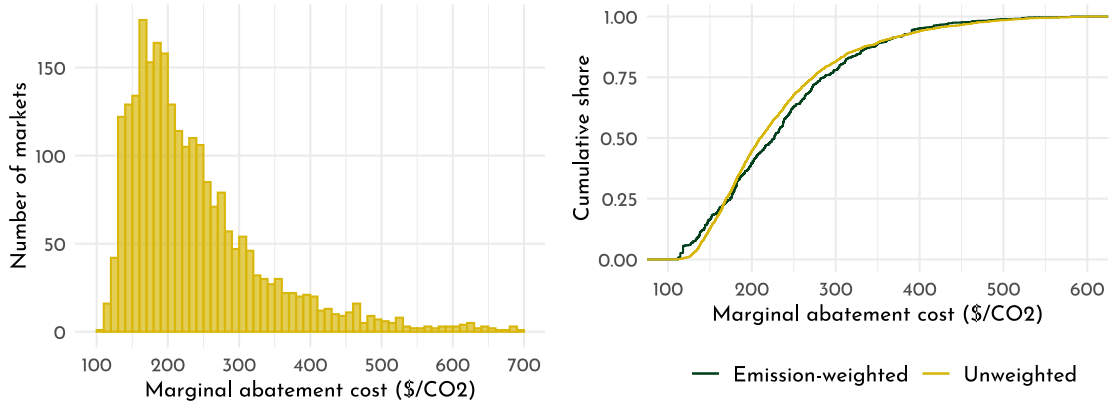

(a) Histogram.

(b) Cumulative distribution.

Notes: these calculations are based on the estimated model, as described in section 2.6.2. In panel (b), the *emission-weighted* distribution weights each market by its total emissions. This can be interpreted as the distribution of MACs with respect to total emissions, so that each point along the line indicates the share of emissions with a MAC at or below that level. In the *unweighted* distribution, shares are relative to the total number of markets.

Figure 2.10: Distribution of optimal taxes in a market-specific scheme under a social cost of carbon of \$230/ton CO_2.

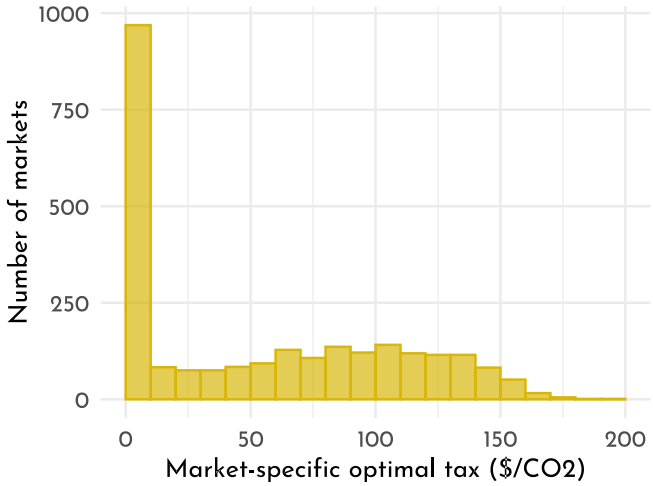

Note: these calculations are based on the estimated model, as described in section 2.6.2.

fare, it does so by taxing markets where the MAC is above the SCC of \$230—at the baseline, about a quarter of emissions have MACs greater than \$230. Conversely, this uniform tax also undertaxes markets with low abatement costs.

Market-specific carbon taxes would be difficult to implement in practice, especially considering that the optimal levels are based on firm market power. However, this hypothetical policy is helpful for gauging the inefficiencies of tax uniformity. For an SCC of $230, moving from a uniform tax to optimal taxes chosen for each market would result in a welfare gain three times as large.[20] As shown in Figure 2.10, these additional gains follow from not taxing markets with high MACs and taxing more heavily those with lower MACs.

2.7.3 Revenue-neutral tax substitution

Even though market distortions act in the direction of reducing aggregate emissions, they are imperfect substitutes for an externality tax: sales taxes and markups do not necessarily match the externalities generated by the carbon emissions of each product. Table 2.5 illustrates how this mismatch happens by comparing nonstop and stop flights. For an appropriate comparison of flights across markets, the first three rows are normalized by market distance. Nonstop flights are more valued by customers, showing higher average markups and fares; therefore, non-stop flights are proportionally more taxed than stop flights. However, nonstop flights travel shorter distances and emit less CO_2 per passenger than stop flights. As a result, markup and sales tax per ton of CO_2 are significantly higher in nonstop flights, thus more strongly disincentivizing the less polluting of the two flight types.

In theory, it would be possible to improve the efficiency of taxation as a pollution instrument. As outlined in section 2.2, the existing sales tax raises rev-

[20] For the low and medium SCC scenarios, market-specific taxes would result in small or no welfare gains. Since baseline MACs are large even at the lower end of the distribution, the optimal tax would still be zero for most markets.

Table 2.5: Differences in average emissions, markups, and taxes for stop and nonstop flights.

	Nonstop	Stop
Fare per market distance ($/mi)	0.315	0.309
Sales tax per market distance ($/mi)	0.024	0.023
Markup tax per market distance ($/mi)	0.089	0.063
Emissions per market distance (Kg CO_2/mi)	0.328	0.381
Sales tax by emissions ($/t$CO_2$)	71.87	60.87
Markup tax by emissions ($/t$CO_2$)	271.46	165.85
Share of passengers (%)	83.36	16.64

Notes: stop flights are flights with at least one connection. Market distance is the great circle distance between the end points of a round trip.

enues for several government operations, including the Federal Aviation Administration. In this sense, replacing the sales tax with a revenue-neutral carbon tax that raises the same amount of revenue can be Pareto improving. Welfare improvements, in this case, are still second best, focusing on environmental policy and taking market power as given.

Panel (a) in Figure 2.11 shows the effects of removing the sales tax and adding a carbon tax at various levels. With no sales or carbon taxes, prices would decrease by 4.7% on average. This variation is below the 7.5% sales tax, thus reflecting an incomplete pass-through with market power. Accordingly, demand would increase by 15.7% and emissions by 18.3%. In this scenario, however, no taxes to fund government operations would raised in this sector.

Panel (b) in Figure 2.11 focuses on the $50–60 range to show that the revenue-neutral carbon tax is about $52/ton CO_2; that is, eliminating the sales tax and replacing it with such a carbon tax would raise the same aggregate tax revenues. This tax substitution would lead to an average decrease in prices of 0.2%; demand would increase by 1.2% and emissions by 0.7%. Panel (b) also highlights an emission-neutral tax level, at which aggregate emissions stay at the same

Figure 2.11: Predicted changes from removing the sales tax and introducing a carbon tax.

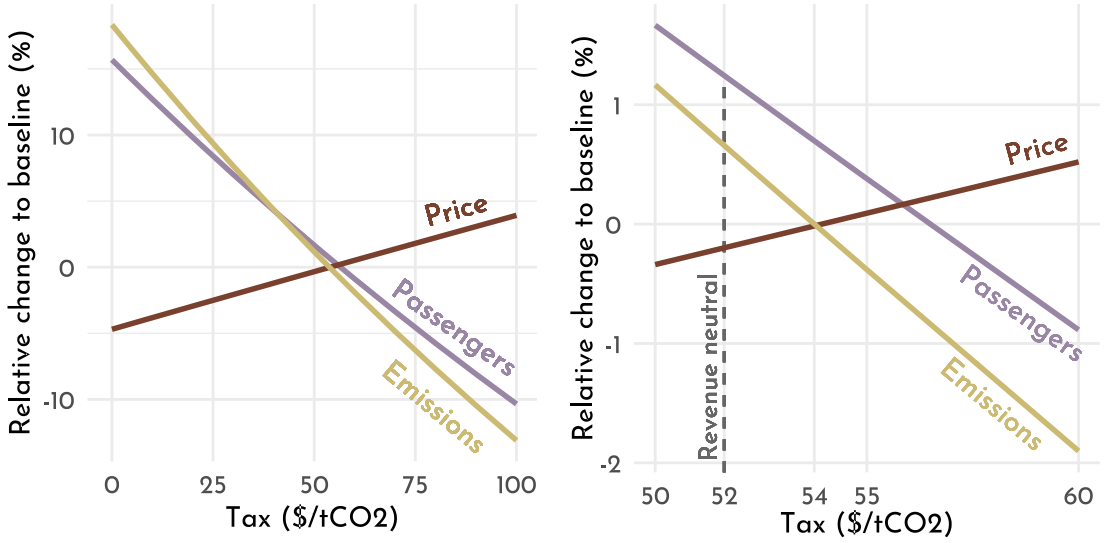

(a) Changes to prices and quantities.

(b) Zooming in on the $50–60 range.

Notes: predictions are obtained from re-computing market equilibria at different tax levels using the estimated model, as described in section 2.6.2. In panel (b), *revenue neutral* is the level of carbon tax that raises the same revenue as the baseline sales tax. *Emission neutral* is the level of carbon tax for which total emissions equals baseline emissions.

level. This exercise also illustrates the role of the current tax in curbing emissions: in the current state of the sector, the sales tax is equivalent to a carbon tax of \$54/ton CO_2.

Regardless of SCC value, a sales tax substitution for a revenue-neutral carbon tax would lead to substantial welfare gains relative to the tax distortions. To calculate welfare changes relative to the tax scheme, I first evaluate aggregate social welfare without any taxes. Then, I introduce either of the two options that raise the same tax revenue in equilibrium: (i) a 7.5% sales tax or (ii) %52 carbon tax. The efficiency gains from the carbon tax are calculated based on the difference between the welfare in each case.

The welfare gains from tax substitution are illustrated in Figure 2.12. In abso-

Figure 2.12: Welfare changes due to sales tax and a revenue-neutral carbon tax under different values for the Social Cost of Carbon.

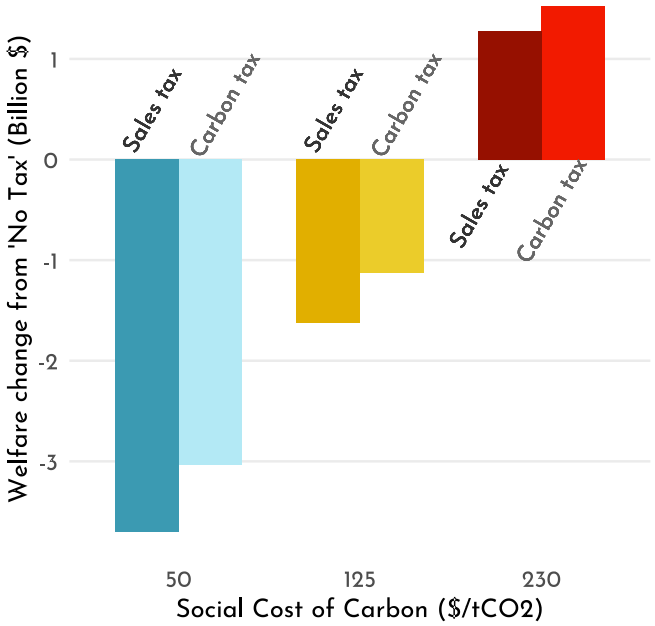

Note: these calculations are based on the estimated model, as described in section 2.6.2.

lute value, these gains vary between \$250–670M, with higher gains associated with lower SCC values. For SCC values of \$50 and \$125/ton CO_2, additional taxation leads to a welfare reduction—a taxation deadweight loss (DWL), which follows from the distortions exceeding the externality in both cases. However, the carbon tax achieves the same level of tax revenue while addressing the externality in a one-to-one relation, resulting in a smaller tax deadweight loss. For an SCC of \$50, these efficiency gains correspond to a reduction of 18% in the tax DWL; for an SCC of \$125, there is an analogous reduction of 30%. Under an SCC of \$230/ton CO_2, the current sales tax improves welfare, since it does part of the job of a carbon tax and must be paired with a positive optimal jet fuel to achieve the second best. In this high-SCC scenario, a revenue-neutral tax substitution would further increase the welfare gains of taxation by an additional 20%, thus improving the efficiency of taxation with regard to environmental concerns.

2.7.4 Limitations

The analyses presented in this chapter are subject to a number of limitations. First, all results concern short-run effects, as airlines' networks and fleets are held fixed. For this reason, the size of distortions can be underestimated when considering longer periods. This is because higher fuel taxes may deteriorate profitability and lead to firm exits that would increase market power for the remaining players. Nevertheless, the effect on exit decisions is likely to be limited, as the cost shocks considered in this chapter are relatively small; a $50 carbon tax, for example, would increase average operating costs by 7.9%. Hence, I would expect the probability of exit decisions to be minimally affected. Furthermore, entry and exit decisions in this sector involve other strategic considerations beyond pure profitability (Belobaba et al., 2015). In addition, over longer periods, average fuel efficiency tends to increase, as older planes are retired or transferred to other markets, and newer and more efficient aircraft models are put in operation.

Second, all analyses are in partial equilibrium, insofar as the chapter considers only the effects on the domestic aviation sector. In doing so, it overlooks any effects on other transportation modes: its projections of emission abatement are intra-sector only. In practice, more expensive flights may lead travelers to substitute away from air transportation for some or all parts of a trip, thus switching to driving or taking buses or trains. For instance, a consumer in a small location within a driving distance from a large airport may substitute across markets, choosing to drive the first segment of their trip. These substitution possibilities create leakage opportunities, which the models used in this chapter do not capture.

Finally, climate change is the only environmental externality considered in this chapter. Aviation has other important environmental consequences, especially those with local impacts. For instance, fuel burn at ground level and low altitudes increases local air pollution, with significant health impacts for populations next to the airports (Schlenker & Walker, 2016). Moreover, take-off, landing, and other airport operations can result in high noise levels, generating a disamenity and affecting property values in the neighborhood of airports (Nelson, 2004).

2.8 Conclusion

This chapter has studied how oligopoly market power and existing distortionary taxes affect environmental policy. Building on seminal work in environmental economics (Buchanan, 1969; Barnett, 1980), I have shown how market imperfections affect optimal environmental taxes, and how welfare effects can be decomposed and attributed to each market imperfection. Based on this theoretical framework, I have evaluated the impact of a carbon tax on aviation. In doing so, I have used sufficient statistics to calculate marginal effects and a structural approach to calculate non-marginal effects and optimal taxes.

My main findings indicate that existing distortions are large and, at the margin, exceed the climate damages from aviation in scenarios where the social cost of carbon is below $200. As a consequence, there is no positive optimal carbon tax in this sector unless the social cost of carbon is high. Even if a positive optimal tax exists, it is only a fraction of the marginal damage, and is thus below the standard Pigouvian tax prescription. I find that the wedge between marginal

damage and optimal tax is primarily driven by market power, which accounts for about three-quarters of the wedge.

These results illustrate a key challenge for aviation: abatement via demand reduction has a high cost in terms of private welfare. Existing taxes and market power already drive equilibrium quantities down, so further reductions in demand come with a significant welfare cost. These features suggest that alternative policies may be more adequate in the short run; multi-sector emission permits and offsets are particularly interesting, as they take advantage of lower abatement costs in other sectors. However, in the long run, more ambitious efforts to curb aviation emissions are likely to depend on technological advancements towards fuel alternatives, among which bio jet fuels and electric and hydrogen-powered planes are potential candidates in development.

CHAPTER 3

CLIMATE POLICY UNDER UNCERTAIN POST-TIPPING WARMING

3.1 Introduction

From an economic point of view, climate change is a market failure caused by a pollution externality. Unlike classic cases of pollution externality, a fundamental feature of climate change is the significant uncertainty over some of its key elements. This uncertainty poses significant challenges for climate policy, as it hinders precise forecasts.

Policymakers, nevertheless, must act in the face of uncertainties while relying on scientific advancements to improve knowledge and reduce margins of error. In this context, one could argue that a 'wait and see' approach is adequate: as we learn more about the phenomenon, we are able to respond more precisely and less wastefully. However, as the geophysics literature warns, abrupt and irreversible changes in the climate regime can be triggered and intensify warming (Alley et al., 2003; Overpeck & Cole, 2006; Lenton et al., 2008; Kopp et al., 2016). These changes significantly increase the costs of inaction. As recent reviews on the economics of climate change indicate, much research is still needed until models are able to provide a credible analysis of tipping points and the risks of climate catastrophes (Greenstone et al., 2013; Metcalf & Stock, 2017; Weyant, 2017).

Tipping points are relevant not only because they alter the climate dynamics, but mostly due to the potential catastrophes that can follow them. For this reason, a fundamental concern related to tipping points is precisely the un-

certainty over what can change after a regime shifts. Despite its importance, this uncertainty has been overlooked in the literature. Previous papers have acknowledged the fact that the location of tipping points can be unknown (e.g. Lemoine & Traeger, 2014; Cai et al., 2015; Lontzek et al., 2015; Belaia et al., 2017). These papers, however, assume that policymakers have full knowledge of the post-tipping regime. Nevertheless, the lack of consensus in expert judgment (Kriegler et al., 2009; Oppenheimer et al., 2016) demonstrates that full knowledge might be an unrealistic assumption. Robust policies must be designed to model and embed uncertainties (Kopp et al., 2016). This chapter takes a step towards addressing this gap in designing robust climate policies.

In this chapter, I study the effect of uncertainty on the climate sensitivity[1] after an abrupt climate regime change. I assess these effects in two steps. First, I compare optimal policies under certainty to optimal policies under three uncertainty scenarios. Second, I consider alternative specifications by changing the baseline Constant Relative Risk Aversion preference (CRRA) to Epstein-Zin-Weil preference (Epstein & Zin, 1989; Weil, 1990) and by using an updated—and more severe—damage function (Howard & Sterner, 2017).

According to economic theory of risk, a prudent policymaker should increase precaution when exposed to higher risk. In doing so, uncertainty on the post-tipping regime should add a risk premium element to the social cost of carbon (SCC).[2] When compared to the case of certainty, this uncertainty leads to a higher abatement rate and lower emissions, temperature, carbon stock, and damages.

[1]The climate sensitivity is the long-term equilibrium increase in temperature from doubling the atmospheric CO_2 concentration.

[2]The social cost of carbon is the present value the stream of future welfare changes caused by a unit of emission in the current period.

The results in this chapter are in line with the economic theory. However, the changes induced by uncertainty are surprisingly small: the maximum estimated increase in the SCC under uncertainty is below 0.2%. This result has two main drivers. First, the climate sensitivity effects are only felt in the long term. Hence, for different values of climate sensitivity, the damages of a unit of carbon emitted today remain very similar during the first few decades; temporal discounting then shrinks the damage differences in the distant future. Second, I show that the continuation value has a relatively low curvature on temperatures. Thus, the expectation of the continuation value over a range of next-period temperatures is close to the continuation value on the expected future temperature. Because uncertainty scenarios are mean-preserving, these differences are very small. Section 3.5 discusses which assumptions are responsible for these features of the model.

Most options to deal with climate change involve devising mechanisms to cope with large-scale externalities. These mechanisms require analyses that are familiar to economists, such as incorporating externalities into the decision-making processes, uncovering existing trade-offs, and formulating optimal policies. In doing so, a significant strand of the literature has focused on embedding geophysical dynamics into economic growth models. Such attempts gave rise to a multitude of Integrated Assessment Models (IAMs), with DICE (Nordhaus, 2014), FUND (Anthoff & Tol, 2010) and PAGE (Hope, 2011) being the most widely adopted.[3] In this chapter, I build on an extension of the DICE-2007 model (Nordhaus, 2008) that includes tipping points and uncertain parameters (Lemoine & Traeger, 2014).

[3]Nordhaus (2014) summarizes the main differences between the three models. Stainforth and Calel (2017) draw parallels between the climate dynamics of each model and identify the economic consequences of their different assumptions.

The first dynamic models developed to analyze the economics of climate change did not account for structural uncertainty. Even though stochastic components were present, with temperature noise being the most common, the parameters governing the evolution of the climate system were assumed to be known by the policymaker. At least since Kelly and Kolstad (1999), several studies have analyzed optimal climate policy with uncertain regime parameters. Part of this literature has focused on quantifying the effects of unknown temperature dynamics, especially the climate sensitivity parameter (Kelly & Kolstad, 1999; Leach, 2007; Calel et al., 2015; Kelly & Tan, 2015; Fitzpatrick & Kelly, 2017; Heutel et al., 2018). As precise estimates of the economic damages inflicted by higher temperatures are yet not available, a set of papers has investigated optimal policy under damage uncertainty (Crost & Traeger, 2013, 2014; Rudik, 2020). Future economic growth is another source of uncertainty considered in the literature (Jensen & Traeger, 2014; Cai & Lontzek, 2019).

The introduction of climate tipping points has been shown to significantly affect optimal choices. Given our incomplete understanding of these phenomena, realistic settings to study policy must consider that climate regimes might change after a specific state crosses an unknown threshold, possibly as a result of a stochastic process. The state triggering the tipping mechanism depends on the type of regime change under analysis, with temperature and the carbon stock being the most frequently used. Section 3.2 discusses the types of tipping elements and previous work on the economic consequences of their existence.

This chapter contributes to the literature by bringing into the climate tipping point analysis a vital part of what makes tipping a reason for concern: its unknown consequences. The problem presented here models the uncertainty over

the climate sensitivity, a fundamental warming parameter which has received significant attention in previous studies. With this setting, it is possible to analyze alternative distribution assumptions and other sources of uncertainty in a post-tipping world.

The remainder of this chapter is structured as follows. Section 2 formalizes the definition of tipping points and discusses related modeling aspects. Section 3 summarizes the climate-economy model used in this chapter. Section 4 presents the main results while section 5 discusses the main findings of this chapter and compares them with the original model by Lemoine and Traeger (2014). Finally, section 5 display concluding remarks and future directions for research.

3.2 Climate tipping points in the geosciences and economics

The literature in geosciences and in economics defines tipping events in several different ways. Each definition provides a different set of modeling choices available to researchers. This section presents some of the most relevant definitions of tipping events and what are choices facing climate-economy modelers.

3.2.1 Definitions of tipping points

A widely used definition of abrupt climate change was introduced in a report by the National Research Council (2002, p. 14): "An abrupt climate change occurs when the climate system is forced to cross some threshold, triggering a transition to a new state at a rate determined by the climate system itself and

faster than the cause". These thresholds, or tipping points, are associated with sub-domains of the climate system referred to as tipping elements.

Lenton et al. (2008) propose a broader and more precise definition of tipping elements: a component of the climate system is a tipping element if

$$|F\left(x \geq x_{crit} + \delta x|T\right) - F\left(x_{crit}|T\right)| \geq \hat{F} > 0,$$

that is, if there exists a control x and a critical value x_{crit} (a tipping point) at which a small variation δx leads to a qualitative change $\hat{F}$ in a system feature F after time T. Meaningful values for δx should be within the range of the natural observed variability (e.g., around 0.2°C for temperature). This definition highlights small changes that cause substantial regime shifts, expanding the set of possible tipping phenomena to include, for example, non-climatic variables and slow transitions. Their definition proceeds to establish three conditions that characterize a policy-relevant tipping element: (i) human decisions within the "political time horizon" ($\sim$100 years, the human lifespan) can affect the crossing of x_{crit}; (ii) the qualitative change $\hat{F}$ can be observed within the "ethical time horizon" ($\sim$1,000 years, the lifetime of civilizations); and (iii) the change affects human welfare or aspects that humans care about.

Several tipping elements satisfy the policy-relevance conditions. The corresponding thresholds for most of these elements are estimated between 0.5 and 6°C, with transition times as short as one year or as long as thousands of years. Some of the potential policy-relevant cases are: the melting of the Greenland ice sheet, the weakening of the Atlantic thermohaline circulation, the strengthening of the El Niño-Southern Oscillator, and the dieback of the Amazon rainforest and the Boreal forest. An overview of the tipping elements and their consequences is beyond the scope of this chapter and can be found in Lenton et al.

(2008) and Kopp et al. (2016).

The complexity of tipping dynamics usually requires researchers to rely on simplifying assumptions when analyzing climate policy. Different approaches have been tried recently to capture specific aspects of the matter. Next, I discuss the main modeling choices in this literature.

3.2.2 Modeling choices

With many possible tipping elements, authors may choose to analyze one element separately or to allow multiple elements to coexist. Most papers allow a single tipping element at a time and perform separate runs with different cases (Lemoine & Traeger, 2014, 2016a; Cai et al., 2015; Lontzek et al., 2015, 2016; Belaia et al., 2017). Other papers study the interaction of multiple tipping points and how their effects might reinforce each other (Lemoine & Traeger, 2016b; Cai et al., 2016).

Three main options exist for modeling the crossing of tipping points, with most thresholds being based on average temperatures. The first option assumes that crossing is deterministic and once warming is beyond a threshold, the regime shifts; the planner, however, does not know the precise location of the threshold and has prior beliefs that adjust over time (Lemoine & Traeger, 2014, 2016a, 2016b; Heutel et al., 2016). The second possibility is stochastic crossing with a calibrated hazard function (Cai et al., 2015, 2016; Lontzek et al., 2015, 2016). The third option is to directly model the physics of the tipping element, which has been done for the Atlantic thermohaline circulation in Belaia et al. (2017).

Finally, researchers determine ahead of time the potential effects of the studied tipping elements. Effects can be direct or indirect. Direct effects cause an output or welfare loss without necessarily affecting the climate regime; the direct effects are permanent and either settle in immediately (Cai et al., 2015, 2016; Heutel et al., 2016) or transition to its full impact over time (Lontzek et al., 2015). On the other hand, indirect effects change structural parameters of the regime, which in turn lead to economic impacts. Most indirect effects assume either an increase in the climate sensitivity or a decrease in the carbon sink (transfer of CO_2 to the oceans) (Lemoine & Traeger, 2014, 2016a, 2016b; Heutel et al., 2016; Belaia et al., 2017). Other parameters can be used to study the impacts of tipping elements on smaller systems; Lontzek et al. (2016), for example, consider changes in the size and in the growth rate of forests.

3.3 Optimal plans for an uncertain world

This subsection outlines the specific choices made in this chapter, highlighting the aspects in which these choices depart from the previous literature. Moreover, I explain how tipping points are embedded in a dynamic stochastic climate-economy model.

3.3.1 Modeling the climate sensitivity uncertainty

In this chapter, I follow the tipping point model proposed by Lemoine and Traeger (2014). In this model, the policymaker does not know a priori the location of the threshold. However, she knows that if the current temperature has

not triggered tipping, lower temperatures will not trigger tipping in the future. Prior beliefs on this threshold are assumed to follow a uniform distribution between T_t, the global average surface temperature in period t, and an exogenous upper limit $\bar{T}$.[4] As future temperature increases, the policymaker estimates that the probability of crossing the unknown threshold also increases, resulting in the following hazard function:

$$h\left(T_t, T_{t+1}\right) = \max\left\{0, \frac{\min\left\{T_{t+1}, \bar{T}\right\} - T_t}{\bar{T} - T_t}\right\} \tag{3.1}$$

In this model, the policymaker immediately detects the climate tipping. Once the tipping has occurred and the system is in a new regime, the climate sensitivity parameter, s, is no longer certain to her; she will form beliefs as to what the new value could be. This uncertainty could be caused by changes of unknown extent in one or several feedback processes; for example, the melting of the Greenland ice sheet would lead to a lower albedo, decreasing the Earth's reflectivity index and increasing energy absorption from solar radiation.

I consider three possible cases to model this uncertainty, all of which reproduce the currently observed fat-tailed feature of climate sensitivity estimation in the geosciences. The derivations of each case can be found in Appendix B.1. The three cases are based on the frequently used climate feedback formulation of s, introduced by Roe and Baker (2007):

$$s = \frac{s_0}{1 - f} \tag{3.2}$$

where $s_0 \approx 1.2$°C is the reference climate sensitivity estimated for a system without feedback and $f \in [0, 1)$ is the feedback coefficient. In many climate

[4] In this chapter, $\bar{T}$ is fixed at 4.27°C, implying an expected initial threshold of 2.5°C. In their analyses, Lemoine and Traeger (2014) also consider other values in the range between 3 to 9°C, corresponding to expected thresholds between 1.9 and 4.9°C.

models, feedback processes, such as energy absorption and inwards reflection in the intermediary atmospheric layers, lead to higher equilibrium temperatures. Because the total feedback is the sum of a large number of these processes, Roe and Baker (2007) assume that f follows a normal distribution with parameters $(\bar{f}, \sigma_f)$. This assumption implies a fat-tailed distribution for s that approximates current kernel estimates of the climate sensitivity uncertainty.

One caveat of the distribution in Roe and Baker (2007) is that f has unbounded support, but physical constraints preclude $f > 1$. The first uncertainty case of this chapter thus truncates the normal distribution at 0 and 1. The resulting truncated Roe-Baker (TRB) distribution has support $[s_0, \infty)$ and probability density function (PDF):

$$\theta^{\text{TRB}}(z) = \frac{1}{\Phi\left(1; \bar{f}, \sigma_f\right) - \Phi\left(0; \bar{f}, \sigma_f\right)} \left\{ \frac{1}{\sigma_f \sqrt{2\pi}} \frac{s_0}{z^2} \exp\left[-\frac{1}{2} \left(\frac{1 - \bar{f} - \frac{s_0}{z}}{\sigma_f} \right)^2 \right] \right\} \tag{3.3}$$

where $\Phi\left(z; \bar{f}, \sigma_f\right)$ is the cumulative probability function (CDF) of the normal distribution with parameters $(\bar{f}, \sigma_f)$ evaluated at z.

Truncating the normal distribution fixes an important issue but leaves at least another aspect of potential concern: there is significant density on the upper bound $f = 1$, which corresponds to runaway, unstoppable climate change with infinite warming—a scenario that violates physical constraints and goes against scientific evidence. A possible workaround is to base f on distributions that are naturally defined over finite supports. The second uncertainty case assumes f follows a beta distribution with parameters (α_f, β_f), which has zero density on the support bounds. The supports of f and s are the same as in the truncated Roe-Baker. The resulting distribution, a Reflected-Inverted-Scaled-

Beta (RIS-Beta), has the following PDF:

$$\theta^{\text{RIS-Beta}}(z) = \frac{(s_0)^{\beta_f}}{B(\alpha_f, \beta_f)} \frac{(z - s_0)^{\alpha_f - 1}}{z^{\alpha_f + \beta_f}} \tag{3.4}$$

where $B(\alpha, \beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1} dt$ is the beta function.

The third and final uncertainty case assumes a prior distribution over f represented by a uniform distribution. Such assumption implies that any scenario between and including zero-feedback and runaway warming are equally probable. Although this assumption does not have any support from scientific evidence, it does serve the purpose of testing extreme tail sensitivity. While the upper bound of the uniform is always fixed at 1, the lower bound, $\underline{f}$, changes to accommodate different expectation scenarios. This implies that s has support $\left[\frac{s_0}{1-\underline{f}}, \infty\right)$. The resulting distribution, a Reflected-Inverted-Scaled-Uniform, (RIS-Uniform) has the following PDF:

$$\theta^{\text{RIS-Uniform}}(z) = \frac{s_0}{1-\underline{f}} \frac{1}{z^2} \tag{3.5}$$

The RIS-Uniform, in fact, is a Pareto distribution with scale parameter set to $\frac{s_0}{1-\underline{f}}$ and shape parameter 1. As the PDF decays polynomially, it recalls the fat-tailed uncertainty assumed in the Dismal Theorem (Weitzman, 2009). Note that the RIS-Uniform distribution does not have a finite first moment. Yet, even in this extreme case where expected climate sensitivity is infinity, the policy problem is well defined and yields well-behaved optimal policies. The reasons for these resulting properties are discussed in Section 3.5.

The specific parameters of each post-tipping uncertainty case are calibrated to make next-period temperature expectations match the certainty scenario with climate sensitivity of 6 (the pre-tipping regime has the certain sensitivity at 3). From the temperature transition equation (B.14), imposing a condition that pa-

rameters must satisfy $\mathbb{E}\left[\frac{1}{s}\right] = \frac{1}{6}$ implies that $\mathbb{E}_{\epsilon,s}\left[\tilde{T}_{t+1}\right] = \mathbb{E}_{\epsilon}\left[\tilde{T}_{t+1}|s=6\right]$ (see details in Appendix B.1.2). Hence, an approach directly matching the expectation of s to 6 would be incorrect, as it would render unfair comparisons with the certainty case.

The condition on the expectation of $\frac{1}{s}$ is sufficient to pin down the RIS-Uniform distribution, as it has only one free parameter. With two parameters, the truncated Roe-Baker and the RIS-Beta require a second condition, for which there is no obvious choice. A reasonable option, which I implement in this chapter, is to impose a quantile condition as in Newbold and Daigneault (2009)[5]: the 95th percentile of the s distributions must be at 20. A climate sensitivity of 20 is well beyond the domain of what current climate models estimate (IPCC, 2013; Knutti et al., 2017). Hence, this is an ad-hoc choice, defined with the intention to significantly spread the probability mass compared to the current (pre-threshold) uncertainty. Figures 3.1 and 3.2 show the distribution functions of f and s for the calibrated parameters.

The analysis in this chapter treats the uncertainty over s as fixed: the policymaker does not learn or improve her knowledge of this parameter. Therefore, learning would lead to a decrease in the uncertainty and to a lower SCC due to precautionary reasons (Kelly & Tan, 2015). Thus, the risk premium-related results from a no-learning solution can be understood as an upper bound to the more realistic learning case. Nevertheless, learning the climate sensitivity is in practice a rather slow process; as Knutti et al. (2017) show, the estimated likely range for the climate sensitivity has barely changed since the late 1970s.

[5]Newbold and Daigneault (2009) consider an expected climate sensitivity of 3 and impose the 95th percentile to be at 10.

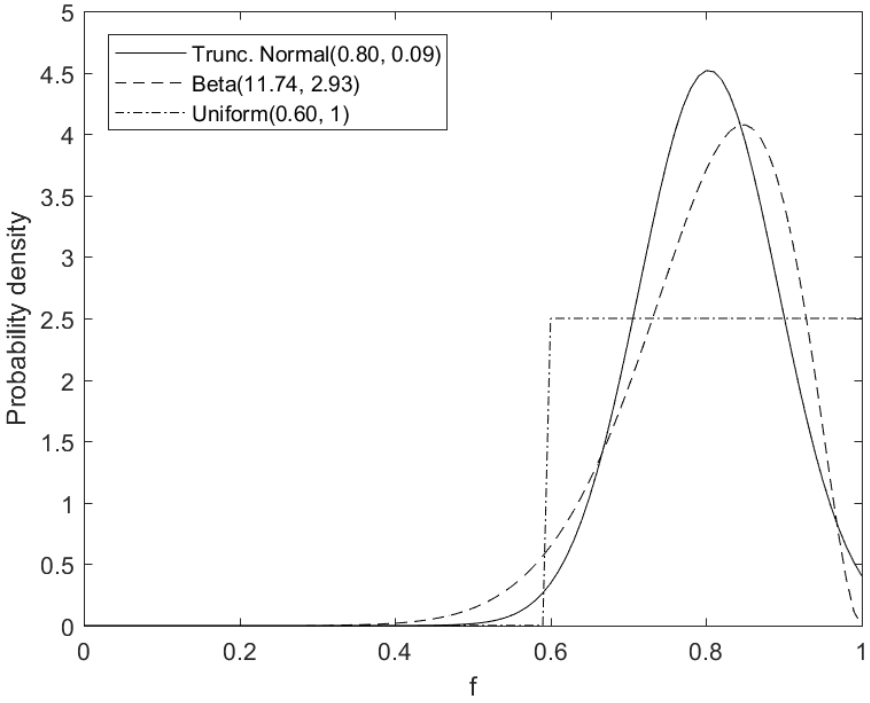

(a) PDFs for f.

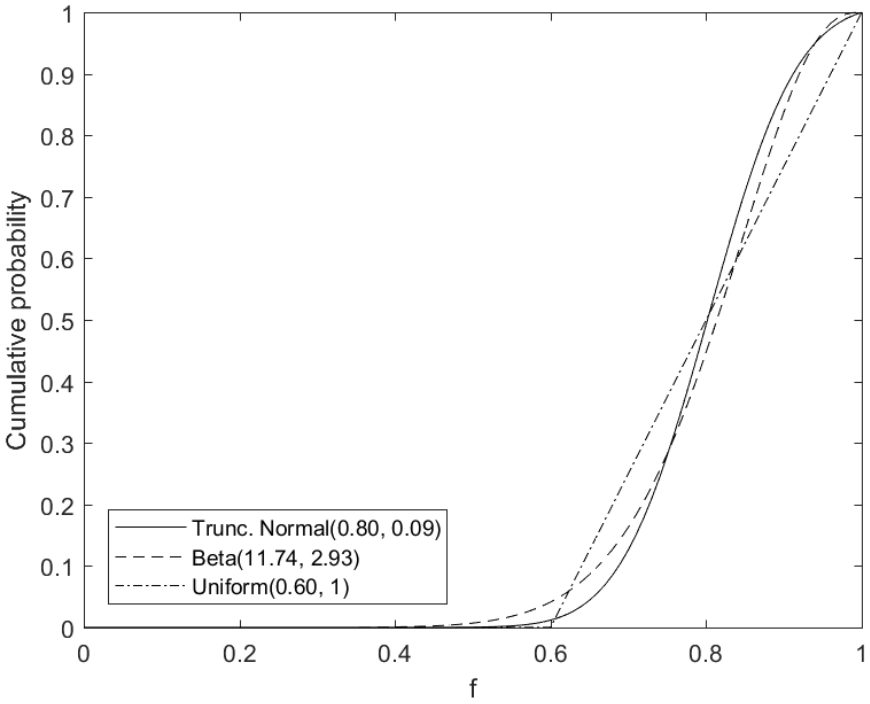

(b) CDFs for f.

Figure 3.1: Calibrated distributions f for each uncertainty case.

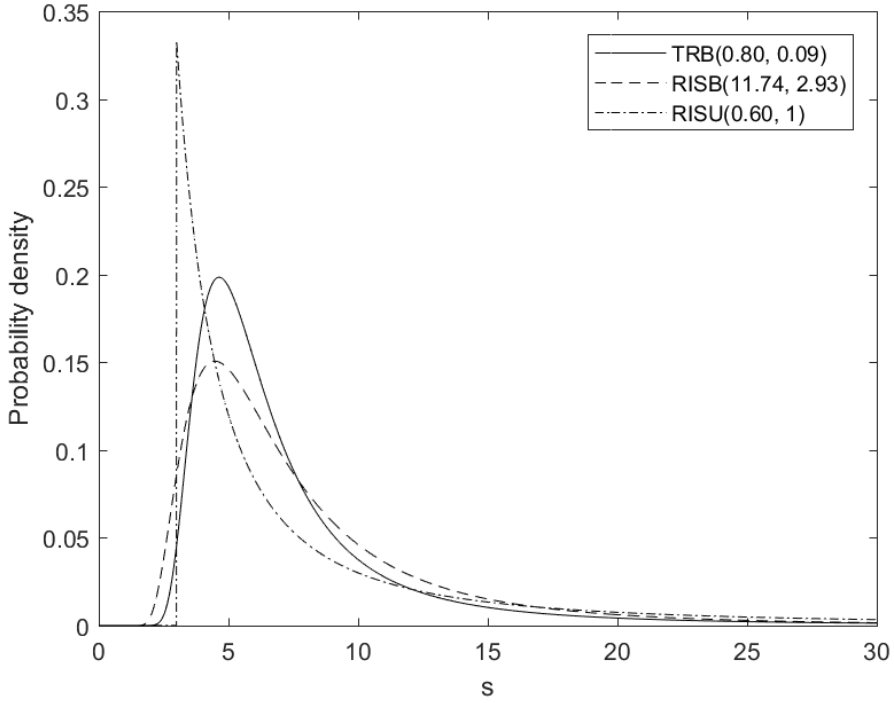

(a) PDFs for s.

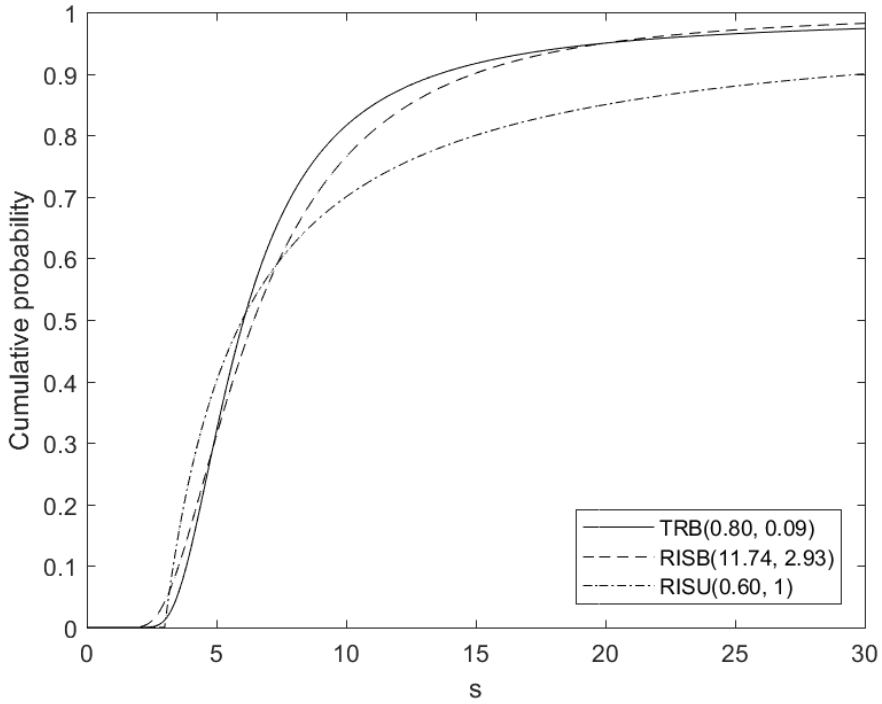

(b) CDFs for s.

Figure 3.2: Calibrated distributions s for each uncertainty case.

3.3.2 The Climate-Economy model

The Dynamic Integrated Climate-Economy (DICE) model, first proposed by Nordhaus (1993), is an infinite-horizon Ramsey-Cass-Koopmans growth model coupled with a stylized climate system. As in the standard growth model, the central planner determines optimal trajectories of capital investment and consumption. In the climate-economy model, these choices determine the level of production, which results in industrial CO_2 emissions that accumulate in the atmosphere; the CO_2 then intensifies the greenhouse effect, increasing the global temperature and causing damages that reduce gross production. In order to avoid climate damages, the planner is given a third resource allocation option: emission abatement. With this extended set of policy choices, the planner dynamically maximizes welfare under a set of economic and climatic constraints. Figure 3.3 presents a diagram of the controls, states, and feedback loops.

Important elements of the model are assumed to be exogenous, such as technological progress, population growth, carbon intensity of industry, nonindustrial emissions, and non-CO_2 forcing. A detailed description of the model is presented in Appendix B.2.1.

The problem can be formulated as a recursive stochastic dynamic program. Let ψ_t represent the state of the climate regime, with $\psi_t = 0$ in the pre-tipping regime and $\psi_t = 1$ in the post-tipping regime. In period t, the joint probability measure of uncertainties with tipping points is then composed of three elements: (i) the temperature shock $\epsilon_{t+1} \in \varepsilon$, (ii) the regime state $\psi_{t+1} \in \{0, 1\}$, and (iii) the post-tipping climate sensitivity $s \in \Sigma$.

Because $\psi = 1$ is an absorbing state, the post-tipping program can be defined

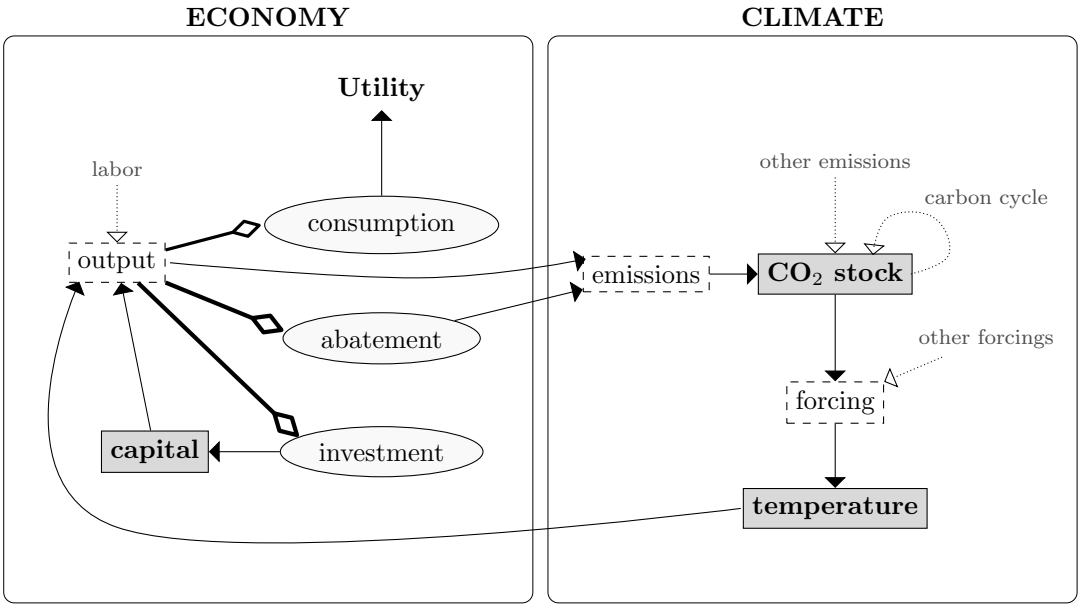

Figure 3.3: Main components and interactions in the DICE model. Solid, shaded boxes represent the states. Planner's decisions are represented by diamond arrows, with shaded ellipses being the variables of choice. Dashed boxes represent endogenous processes. Dotted lines represent exogenous processes.

and solved separately. Furthermore, ϵ_{t+1} and s are uncorrelated, allowing to define separate probability measures $\mathbb{P}_\epsilon$ and $\mathbb{P}_s$. Let V_ψ denote the value function in climate state ψ and period t. Then, the post-tipping problem is given by

$$
\begin{aligned}
& V_1\left(\mathbf{S}_t\right) = \max_{\mathbf{x}_t} U\left(\mathbf{x}_t, \mathbf{S}_t\right) + \beta \int_\Sigma \left[\int_\varepsilon V_1\left(\mathbf{S}_{t+1}\right) d\mathbb{P}_\epsilon \right] d\mathbb{P}_s \\
\text{s.t.} \quad & \mathbf{S}_{t+1} = g_1\left(\mathbf{x_t}, \mathbf{S}_t, \epsilon_{t+1}, s\right) \\
& \mathbf{x_t} \in \Gamma\left(\mathbf{S}_t\right)
\end{aligned} \tag{3.6}
$$

where $\mathbf{x}_t$ and $\mathbf{S}_t$ are the vectors of controls and states (other than ψ) in period t; g_ψ is the state transition function in regime ψ; $\Gamma(\cdot)$ is the set of admissible controls; and β is the social discount rate.

With V_1 obtained by solving (3.6), the pre-tipping problem is given by

$$V_0(\mathbf{S}_t) = \max_{\mathbf{x}_t} U(\mathbf{x}_t, \mathbf{S}_t) + \beta \int_{\varepsilon} \Big\{ [1 - h(\mathbf{S}_t, \mathbf{S}_{t+1})] V_0(\mathbf{S}_{t+1}) + h(\mathbf{S}_t, \mathbf{S}_{t+1}) V_1(\mathbf{S}_{t+1}) \Big\} d\mathbb{P}_\epsilon$$

$$\text{s.t.} \quad \mathbf{S}_{t+1} = g_0(\mathbf{x}_t, \mathbf{S}_t, \epsilon_{t+1})$$

$$\mathbf{x_t} \in \Gamma(\mathbf{S}_t) \tag{3.7}$$

The objective function in (3.7) accounts for the probability that the climate might tip in the next period; this probability is captured by the hazard function, as defined in (3.1). Once the climate tips, $h(\mathbf{S}_t, \mathbf{S}_{t+1}) = 1$ for the remaining periods, and the problem turns into (3.6). An implicit and important assumption in (3.7) is the instantaneous tipping: if S_{t+1} crosses the threshold, the next period starts under the new regime. Note that the initial states in the post-tipping regime are still governed by the pre-tipping transition g_0.

The outer integral in the objective function of (3.6) represents the post-tipping uncertainty. In comparison with the certainty case in Lemoine and Traeger (2014), this additional term implicitly defines a different V_1. Carried into the pre-threshold problem, the redefined V_1 leads to a different V_0 and a different set of optimal policies.

The recursive formulation has the advantage of being solvable by value function iteration. However, the large set of states imposes computational limits to finding numerical solutions. One useful approach is to group similar states, substantially reducing the total number of points in the solution grid. The model used in this chapter follows a version of this approach proposed by Traeger (2014): it represents the ocean temperature as a function of the atmospheric temperature, and both deep and shallow ocean carbon concentrations as functions of the atmospheric concentration. As Traeger (2014) shows, this technique

reduces the total number of states from 7 to 4 without significant loss of precision. Additional adjustments are made to improve the efficiency of the computational solution methods; Appendix B.2.2 presents additional details of the solution methods used in this chapter.

3.4 Results

To assess the effect of post-tipping regime uncertainty, I compare the optimal trajectories for the certainty case and each of the uncertainty cases. The trajectories of greater interest are the optimal carbon tax (which equals the SCC), warming, and atmospheric stock of CO_2. Results are presented for the baseline model with a CRRA utility function, as detailed in Appendix B.2.1, and for two alternative specifications: (i) with Epstein-Zin-Weil (EZW) preferences, to investigate whether the results are driven by the more restrictive assumptions about risk attitude in the CRRA; and (ii) with a steeper damage function estimated by Howard and Sterner (2017), to assess whether the small observed effects are associated with the lower damages implied by the standard damage function.

The main finding of this chapter is that post-tipping uncertainty over the climate sensitivity does not affect optimal policy substantially, at least under the set of assumptions adopted in the literature of DICE with tipping points. As expected, increasing the uncertainty by making tails fatter increases the risk premium. However, the effect is so small that it barely affects the main outcomes: changes in states and controls are well below 1%. For this reason, the trajectory plots display only the certainty case for illustration and discussion, as the stacked curves for uncertainty are visually indistinguishable in low-resolution

figures.

The trajectories presented in Figures 3.4 and 3.5 for the baseline model are similar to those reported by Lemoine and Traeger (2014) for the case of $s = 6$. Peak CO_2 concentration, below 550 ppm, and peak warming, below 3°C, occur between 100 and 150 years from the start of the policy. Panel d summarizes the mean relative changes along the trajectories of each uncertainty case with respect to the certainty case. The increase in the optimal tax for the case with the highest uncertainty (RIS-Uniform) is approximately 80% higher than the case with the lowest uncertainty (truncated Roe-Baker). Hence, the intuitive notion that an increase in uncertainty leads to more precaution finds correspondence in the results. However, the magnitude of the effects is so small that it has little practical impact on the outcomes.

The results above assume CRRA preferences, which do not distinguish risk aversion and intertemporal smoothing of consumption. Such assumption can arguably condition the optimal policy to underestimate the risk aversion. The EZW preferences, on the other hand, disentangles these attitudes by having independent parameters for relative risk aversion (RRA) and intertemporal elasticity of substitution ($1/\eta$). To investigate if the small effects for the baseline model are an artifact of the CRRA risk attitude assumptions, optimal climate policy is calculated for a planner with EZW with parameters frequently used in the literature.[6]

The results with EZW preference are presented in Figures 3.6 and 3.7. In line with Crost and Traeger (2014), an optimal policy under EZW preference is sub-

[6]I adopt the same values as in Traeger (2014): $RRA = 10$ and $\eta = \frac{2}{3}$. Traeger's paper also presents a short discussion on some of the main estimates used in previous papers, especially in the finance literature.

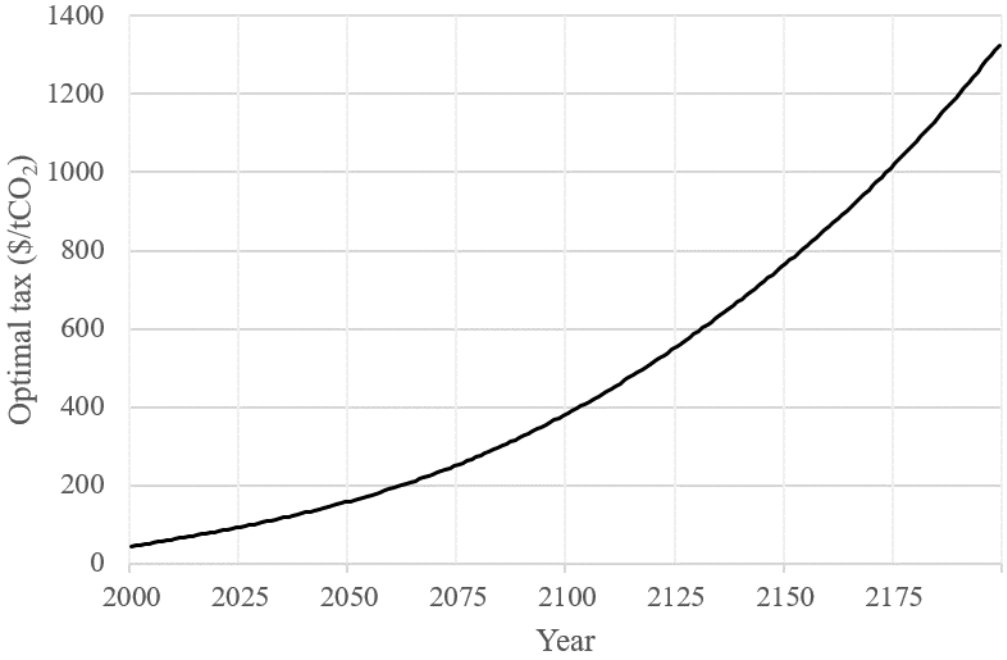

(a) Optimal tax.

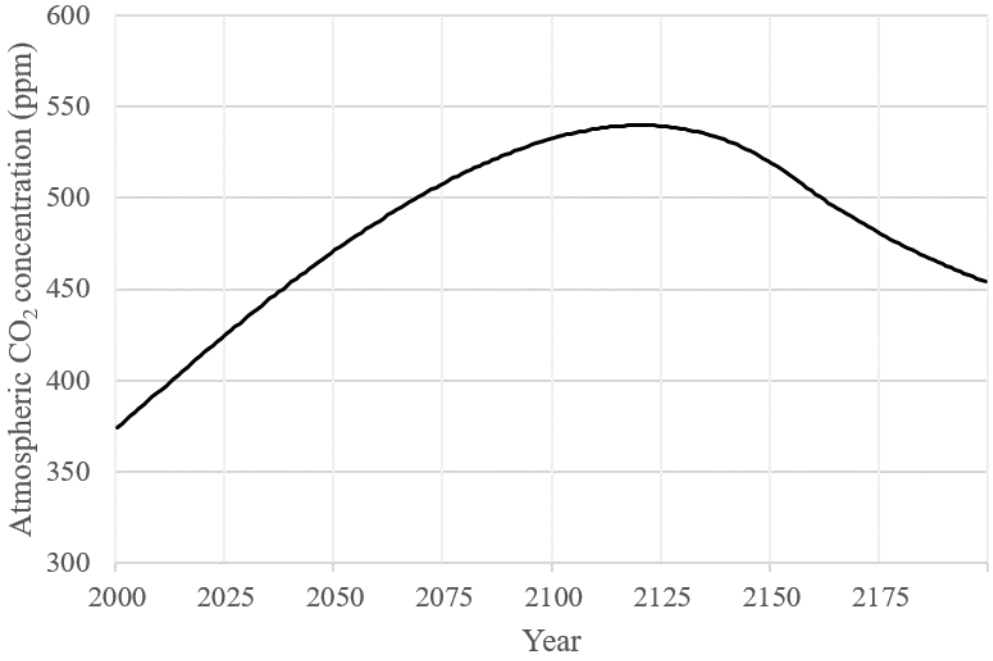

(b) Atmospheric CO_2 stock.

Figure 3.4: Results for the baseline model (with CRRA utility function).

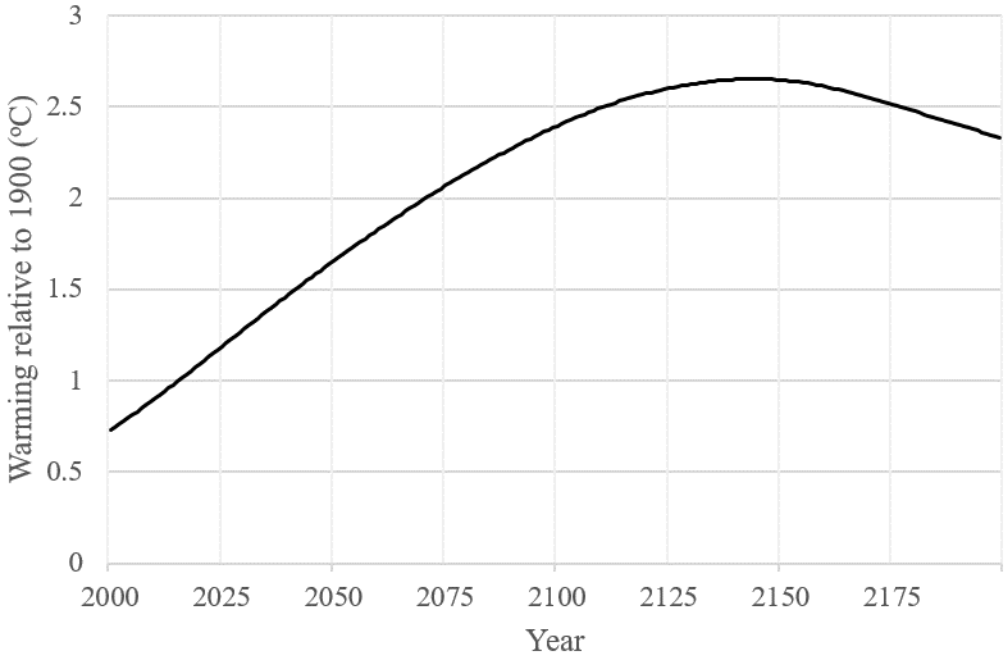

(a) Warming.

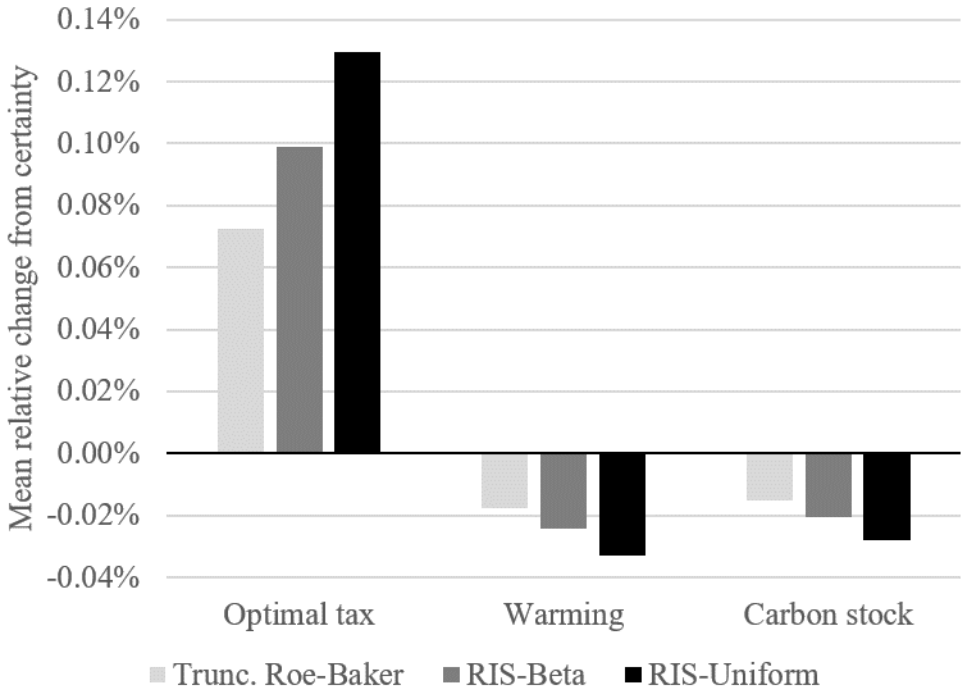

(b) Difference in outcomes from certainty.

Figure 3.5: Other results for the baseline model (with CRRA utility function).

stantially more cautious than under CRRA: optimal tax is higher and increases faster, and peak warming is limited below 2°C; this result is primarily driven by the higher intertemporal elasticity of substitution rather than the increased risk aversion. Nevertheless, uncertainty over the climate sensitivity only introduces a modest increase in the relative changes from the certainty case, as demonstrated in Panel d.

In DICE, the planner cares about higher temperatures only through economic damage. If damages are small enough, the optimal climate policy will prescribe low abatement efforts. Howard and Sterner (2017) empirically estimate a climate damage function and show that the function used in Nordhaus (2008) and subsequent papers underestimate damages by a factor of 3 to 4. A steeper damage curve, as in Howard and Sterner (2017), spreads the effect of different temperature and could in principle intensify the effects of warming uncertainty. The third scenario in this chapter then calculates the optimal climate policy for the baseline model with a damage function consistent with up-to-date evidence.

As Figures 3.8 and 3.9 show, higher damages result in higher optimal taxes and greater abatement; peak CO_2 and warming, consequently, are much lower and occur earlier than in the baseline. The relative changes caused by uncertainty, however, are barely affected; in fact, they decrease slightly.

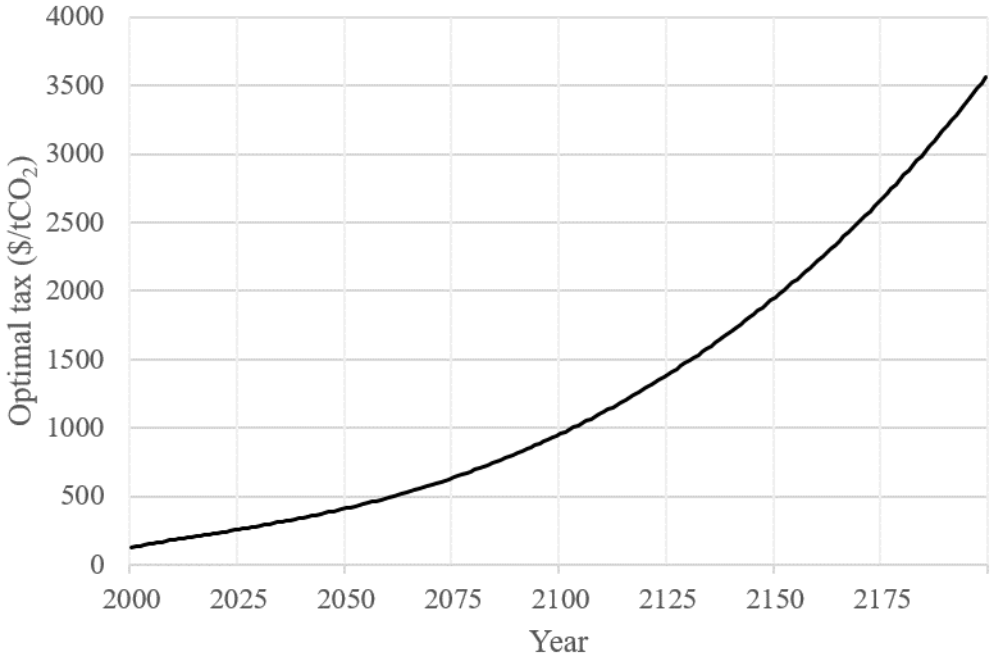

(a) Optimal tax

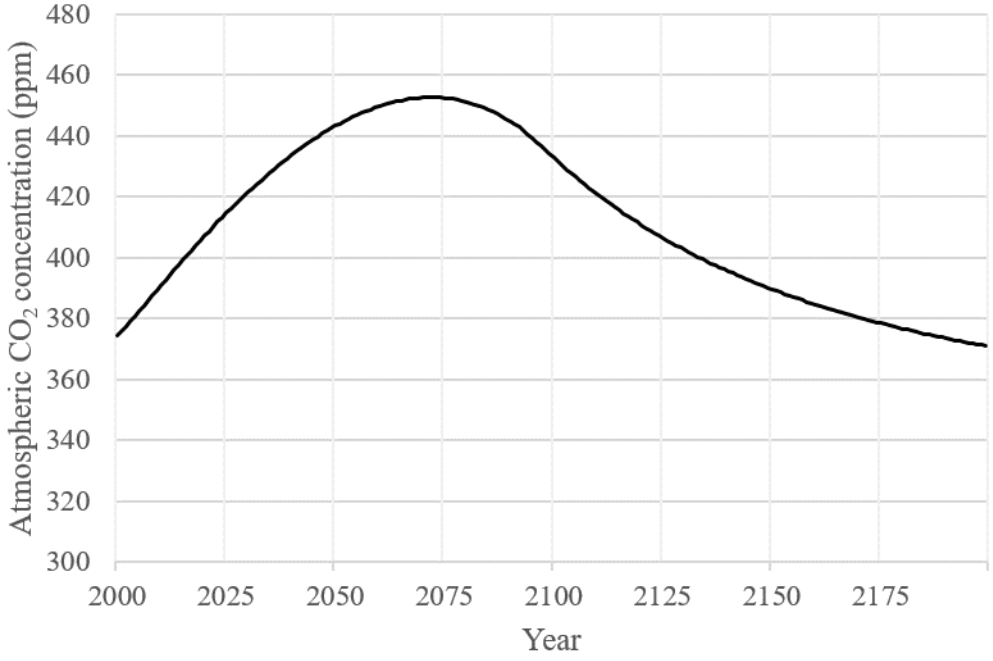

(b) Atmospheric CO_2 stock

Figure 3.6: Results for Epstein-Zin-Weil preference.

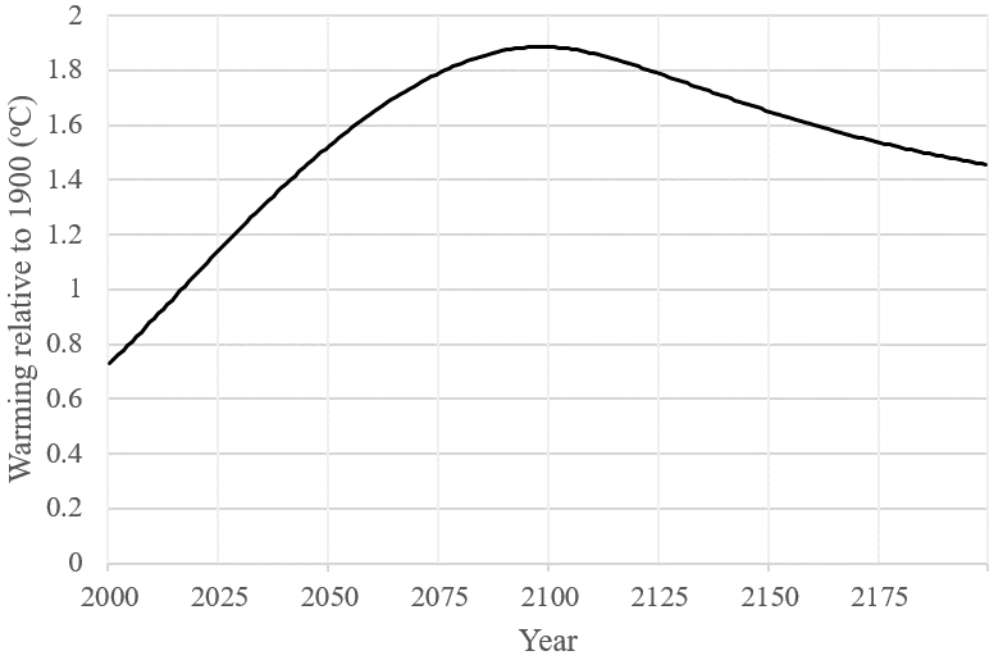

(a) Warming.

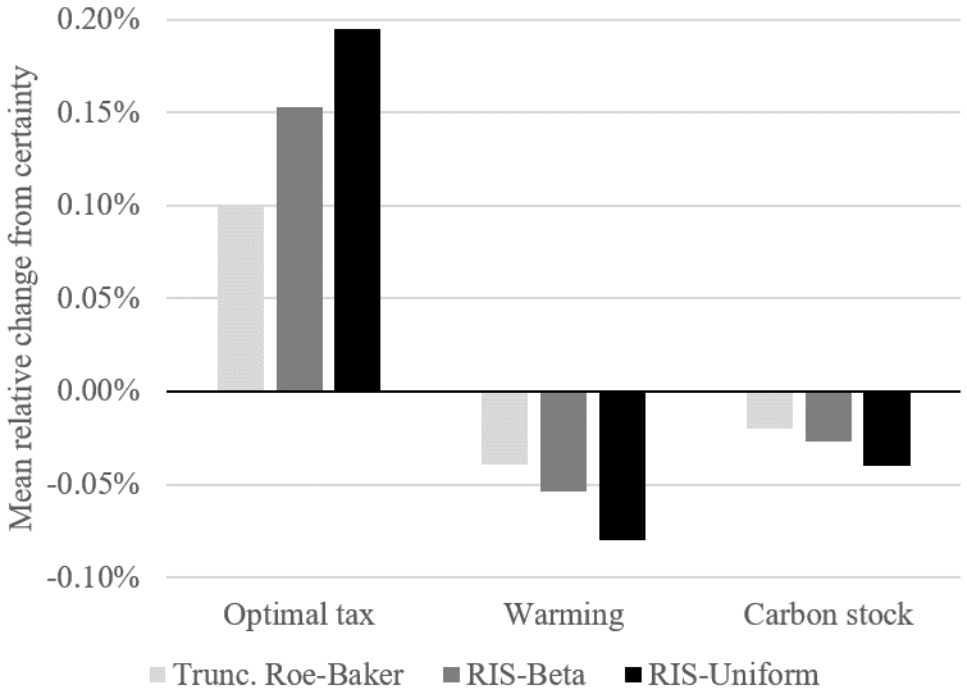

(b) Difference in outcomes from certainty.

Figure 3.7: Other results for Epstein-Zin-Weil preference.

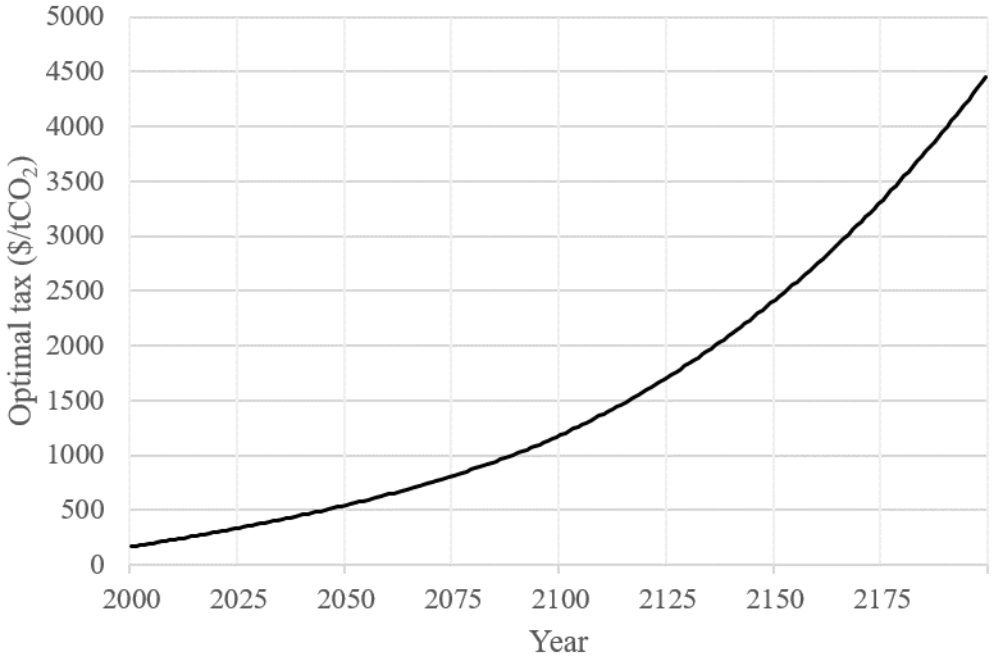

(a) Optimal tax.

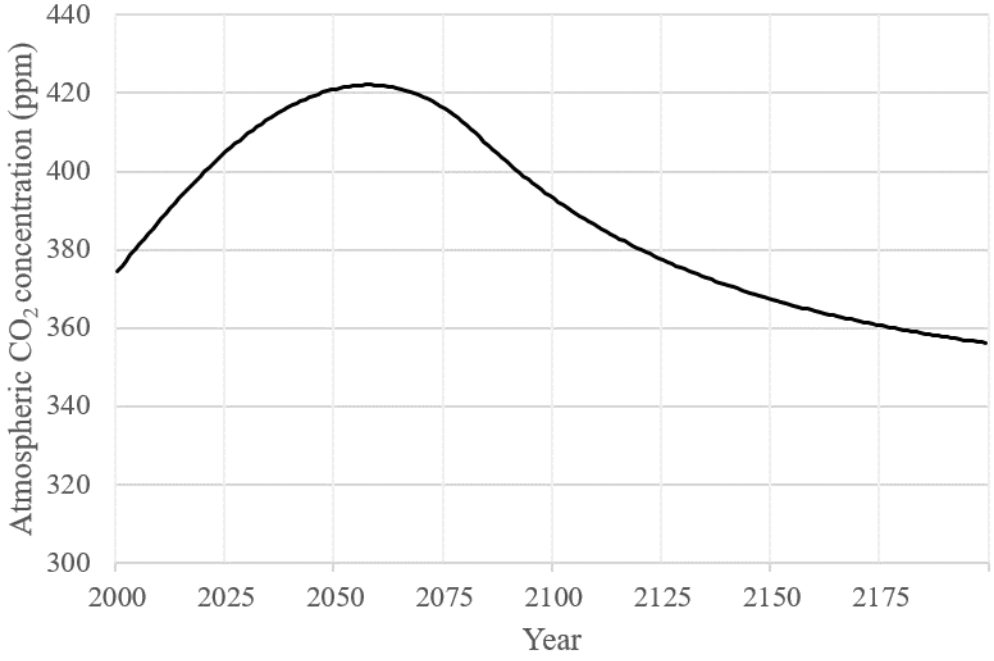

(b) Atmospheric CO_2 stock.

Figure 3.8: Results for the damage function estimated in Howard and Sterner (2017).

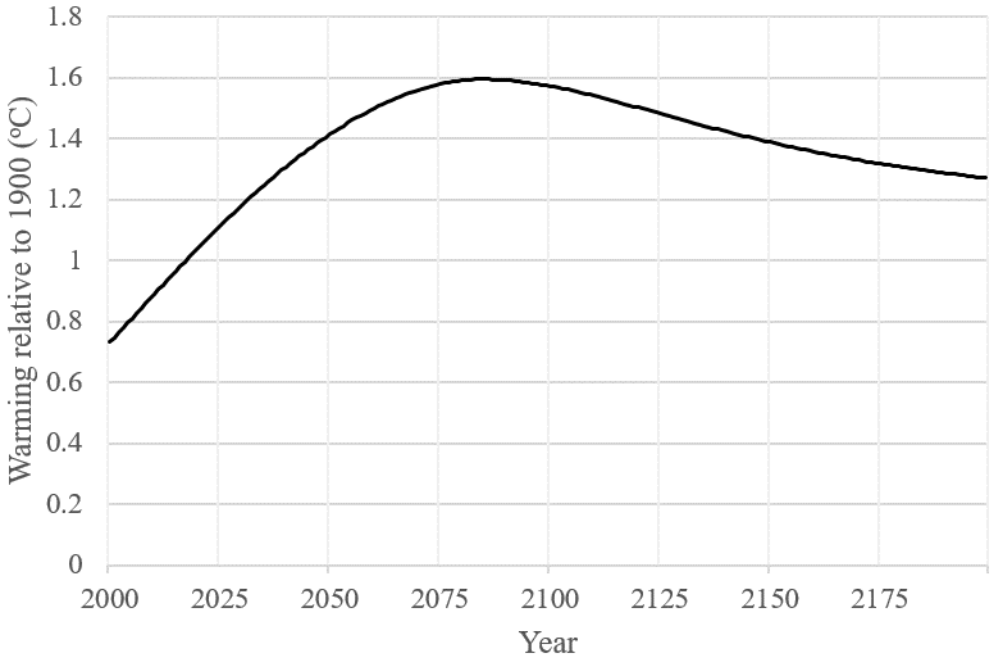

(a) Warming.

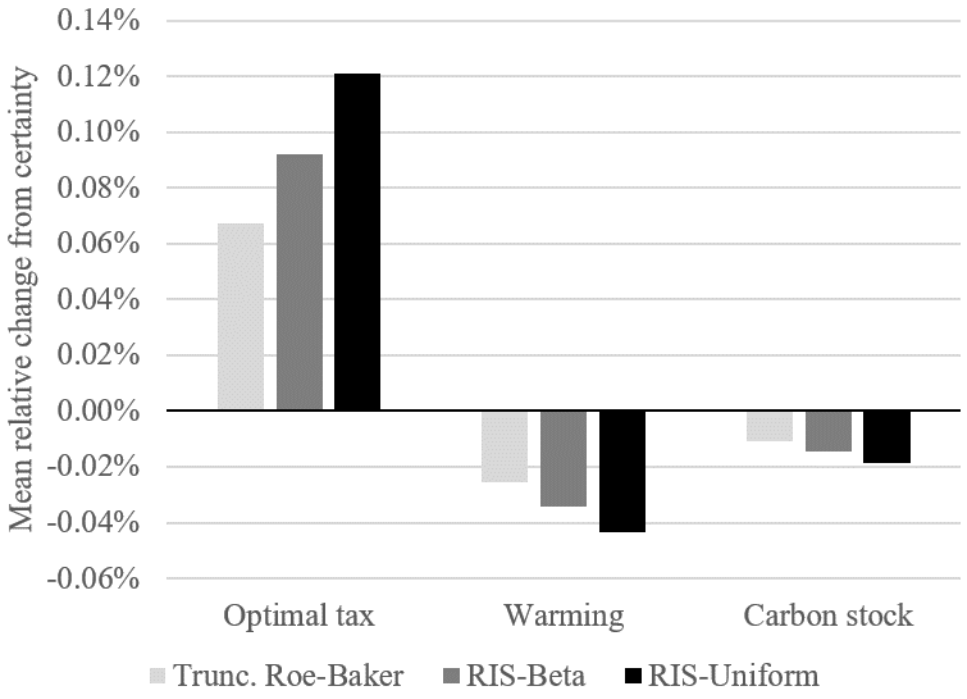

(b) Difference in outcomes from certainty.

Figure 3.9: Other results for the damage function estimated in Howard and Sterner (2017).

3.5 Discussion

The direction of the effects observed in the results is consistent with the economic theory: higher exposure to risks leads to more precaution. The magnitude of these effects, however, raises questions. The climate sensitivity parameter has received significant attention in the economics of climate change. It is understood that a change in this parameter has large impacts on the optimal policy. Nevertheless, the fact that substantial uncertainty over s matters so little prompts further investigation.

The first step is to understand how the climate sensitivity enters the DICE model and what are its consequences to warming. For this matter, we dissociate the economy part of the model and focus only on the stylized climate system in the post-tipping regime. In this exercise, assume a post-tipping regime with emissions fixed at 8 tCO_2 per year, a value close to the initial emission level. Furthermore, capital, output, technological progress, and labor are also fixed at their initial levels. This "fixed economy" holds constant the inputs to the climate system and allows us to turn our attention to temperature and damages.

At each period along the trajectory of temperature, we can calculate the support of the distribution of next-period temperature ($\tilde{T}_{t+1}$) using equation (B.14). For each of the uncertainty cases, we can also obtain confidence intervals. Due to the calibration, the expected value of $\tilde{T}_{t+1}$ matches the actual T_{t+1} in all of the cases. Taking the case of the RIS-Beta distribution, Figure 3.10 illustrates how narrow are the support and the 99% confidence interval (CI) for next-period temperature under climate sensitivity uncertainty. Even though the support for warming after many periods is wider, the one-period case is particularly im-

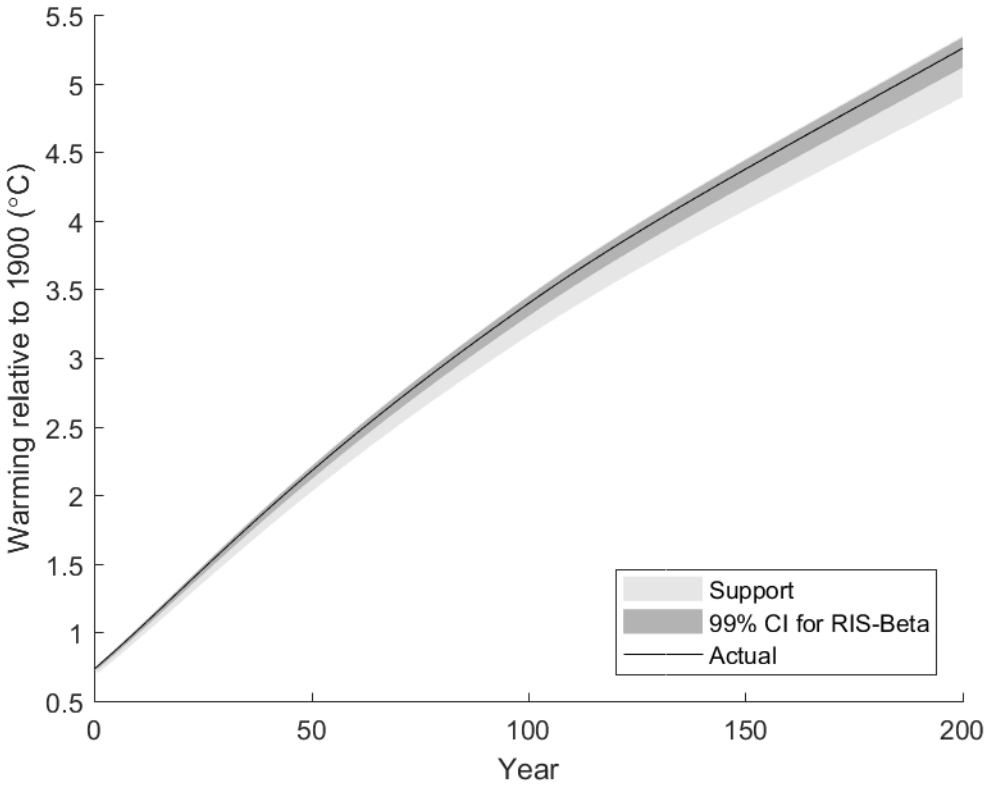

Figure 3.10: Next-period temperature uncertainty for the fixed economy illustration.

portant to evaluate continuation values in a recursive setting; we return to this point below.

Next, we evaluate how the distribution of future temperatures translates into damages. As the damage function is convex, by Jensen's inequality the expected damage will be different from the actual damage. However, as shown in Figure 3.11, the support of $\tilde{T}_{t+1}$ is so narrow that in this range the damage function is approximately linear. Hence, the expected damage and the actual damage practically overlap and cannot be distinguished visually.

Table 3.1 has the numerical values presented in Figures 3.10 (Panel a) and 3.11 (Panel b) at five different periods to illustrate the small differences between

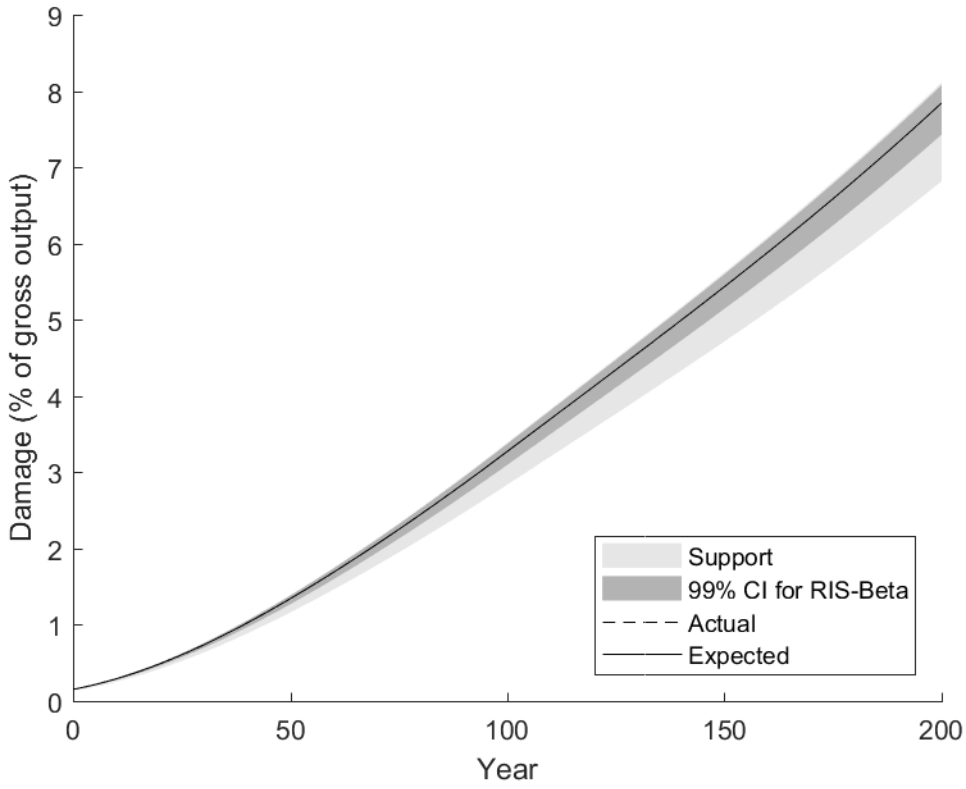

Figure 3.11: Illustration of the next-period damage uncertainty for the fixed economy.

expected and actual temperatures and damages. The second column has the actual trajectory (under certainty), while the third and fourth columns show the expected values for the certain and uncertain climate sensitivity cases, both with temperature shocks; the remaining columns show the support bounds and the 0.5 and 99.5 percentiles for the RIS-Beta distribution.

A closer look on Table 3.1 provides further insights on the role of temperature shocks. Within the stylized climate system of DICE and without learning, a mean-preserving uncertainty over climate change, in fact, acts as a transformation of the i.i.d. temperature shock ϵ_t. This becomes clear by comparing the third and fourth columns in Table 3.1: the sensitivity uncertainty slightly am-

Panel A: Next-period warming under RISB (° C)

Period	Actual	$\mathbb{E}_\epsilon\left[\tilde{T}_{t+1}\right]$	$\mathbb{E}_{\epsilon,s}\left[\tilde{T}_{t+1}\right]$	Lower bound	P0.05	P99.5	Upper bound
1	0.7578	0.7578	0.7578	0.7061	0.7374	0.7692	0.7708
10	0.9849	0.9849	0.9849	0.9173	0.9582	0.9997	1.0018
50	2.1462	2.1462	2.1462	1.9976	2.0875	2.1787	2.1833
100	3.3736	3.3736	3.3736	3.1411	3.2818	3.4245	3.4317
200	5.2390	5.2390	5.2390	4.8846	5.0990	5.3165	5.3275

Panel B: Next-period damages under RISB (% of gross GDP)

Period	Actual	$\mathbb{E}_\epsilon\left[D(\tilde{T}_{t+1})\right]$	$\mathbb{E}_{\epsilon,s}\left[D(\tilde{T}_{t+1})\right]$	Lower bound	P0.05	P99.5	Upper bound
1	0.1630	0.1630	0.1631	0.1415	0.1544	0.1680	0.1687
10	0.2754	0.2754	0.2754	0.2389	0.2606	0.2837	0.2849
50	1.3076	1.3077	1.3078	1.1328	1.2371	1.3475	1.3533
100	3.2309	3.2311	3.2313	2.8010	3.0574	3.3291	3.3432
200	7.7916	7.7919	7.7925	6.7733	7.3809	8.0240	8.0573

Table 3.1: Illustration of the next-period distributions for the fixed economy.

plifies the effects of ϵ_t. Even though different sensitivities lead to substantial temperature differences in the long term, the uncertainty over the sensitivity is never resolved: the policymaker starts each period without knowing the true parameter and only foreseeing a small range of possible warming for the next period.

This illustration with a fixed economy provides an intuition as to why mean-preserving uncertainty over the climate sensitivity would generate small deviations in optimal policy. Having in mind the narrow support of next-period temperature, we can return our attention to the full model. With such a small uncertainty, especially for warming below 3°C, only a continuation value function with high curvature would result in significant differences between expected and actual values; as Figure 3.12 shows, this is not the case. For small intervals over T, V_0 is approximately linear.

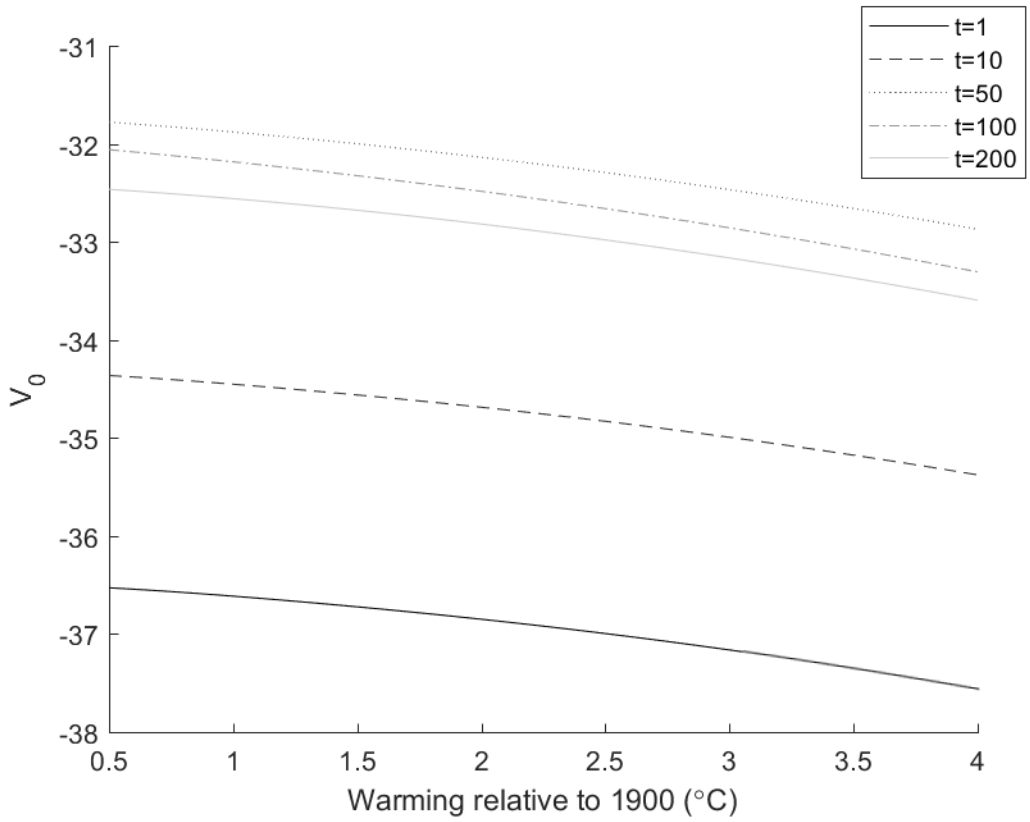

Figure 3.12: V_0 at different warming levels at different periods along the optimal trajectory for the RIS-Beta case.

3.6 Conclusion

This chapter has shown that post-tipping uncertainty over the climate sensitivity in the DICE model does not affect optimal policy substantially. Building on the warming tipping point model proposed by Lemoine and Traeger (2014), I showed that increases in the optimal tax when the post-tipping climate sensitivity is unknown are below 0.2% for a set of different modeling assumptions; warming, CO_2 concentration, and other outcomes also observe small changes. These results are qualitatively in line with economic theory: higher risks translate into more precaution. On the other hand, the meager quantitative effects are a consequence of the small short-term impact of changing the climate sensi-

tivity; with an uncertainty that preserves the mean of next-period temperatures, the range of possible one-period warming is rather narrow.

Do the results mean we should not be worried about post-tipping uncertainty? Absolutely not. First, all the results are conditional on climate sensitivity beliefs that are correct on the mean; even with learning, prior beliefs with incorrect means could substantially increase the realized cost of the policy. Second, this chapter has not explored other important sources of post-tipping uncertainty, such as damages and growth, which could lead to much bigger impacts. Third, the questions posed here most likely push the DICE model to a region beyond its usefulness. In many geophysical climate models, sensitivity is an emergent coefficient rather than a parameter; it is frequently used for model intercomparison and focuses on the very long term equilibrium, which is heavily discounted in DICE. This suggests that in order to properly quantify uncertainties, pre- and post-threshold, we could instead focus on the fundamental regime parameters, such as the ones dictating heat and mass flow that might be affected by a tipping point. As DICE is one of the main models currently used by governments to guide policy, it is important to understand the model's limitations.

In climate systems, changes in structural parameters caused by tipping elements do not occur instantaneously nor are promptly detected; these parameters are in fact hidden states and have to be understood as such in more realistic tipping models. Advancing the analysis of robust climate policy with tipping points requires acknowledging the dynamic formation of beliefs on possibly changing parameters, thus accounting for both the detection and learning of tipping mechanisms.

CHAPTER 4

WATER AFFORDABILITY IN THE UNITED STATES

4.1 Introduction

Water is necessary for human survival. The United Nations identified "equitable access to safe and clean drinking water and sanitation as an integral component of the realization of all human rights" (UN General Assembly, 2010). Water is also an economic good, whose price should reflect its value to society and the long-run costs associated with its treatment and distribution to customers (Olmstead, 2010). Utilities typically price water to recover costs of provision and recent evidence suggests that utilities do not price water to reflect scarcity (Luby, Polasky, & Swackhamer, 2018). To maintain current levels of service in the US, however, water and wastewater infrastructure will require substantial investment over the next several decades, with some estimates totaling more than 1-trillion USD (American Water Works Association, 2012). Compliance with the US Environmental Protection Agency's (EPA) Clean Water Act and Safe Drinking Water Act further adds to water supply costs (National Academy of Public Administration, 2017; Jerch, 2019). The vast majority of those costs will end up on household water and wastewater bills, potentially tripling the current cost of water and sewer service for US households (American Water Works Association, 2012).

Access to affordable, clean water is a significant concern in developing countries (Whittington, Nauges, Fuente, & Wu, 2015), although it is also an issue in developed countries (Martins, Quintal, Cruz, & Barata, 2016). In the US, recent

*This chapter follows from a working paper co-authored with Casey J. Wichman.

trends suggest that the cost of water and wastewater is rising three times faster than other goods and services at a time when economic inequality is increasing (Figure 4.1). States and municipal authorities are beginning to develop policies to reduce the burden of water and sewer bills for low-income households—in 2015, California passed a law to develop a statewide low-income water rate assistance program (California State Assembly, 2015), the City of Philadelphia implemented in July 2017 the nation's first income-based water rates (City of Philadelphia, 2015), and many utilities are adopting low-income water rate assistance programs.

In this chapter, we demonstrate how widespread water affordability issues are in the US, how policies can be designed to reduce burdens on low-income populations, and how underlying economic incentives drive policy effectiveness. We estimate that approximately 14% of households in the US face water affordability concerns by assembling one of most comprehensive data sets of water and sewer rates to date, using rate structures from 1,545 utilities that provide water and sewer service to approximately 45% of the US population. Our data set is compiled from rates surveys conducted by the Environmental Finance Center at the University of North Carolina at Chapel Hill (UNC EFC) and the American Water Works Association (AWWA) and includes water and sewer prices matched with community-level socioeconomic characteristics and typical water use. Using these data, we show that using the full income-distribution at a local level, rather than median household income, is imperative for capturing the water affordability burden of low-income households. We estimate that households in the lowest income decile spend 8.1% of their annual income on water and sewer services on average. Additionally, we find that water affordability concerns are positively correlated with price levels, the proportion of

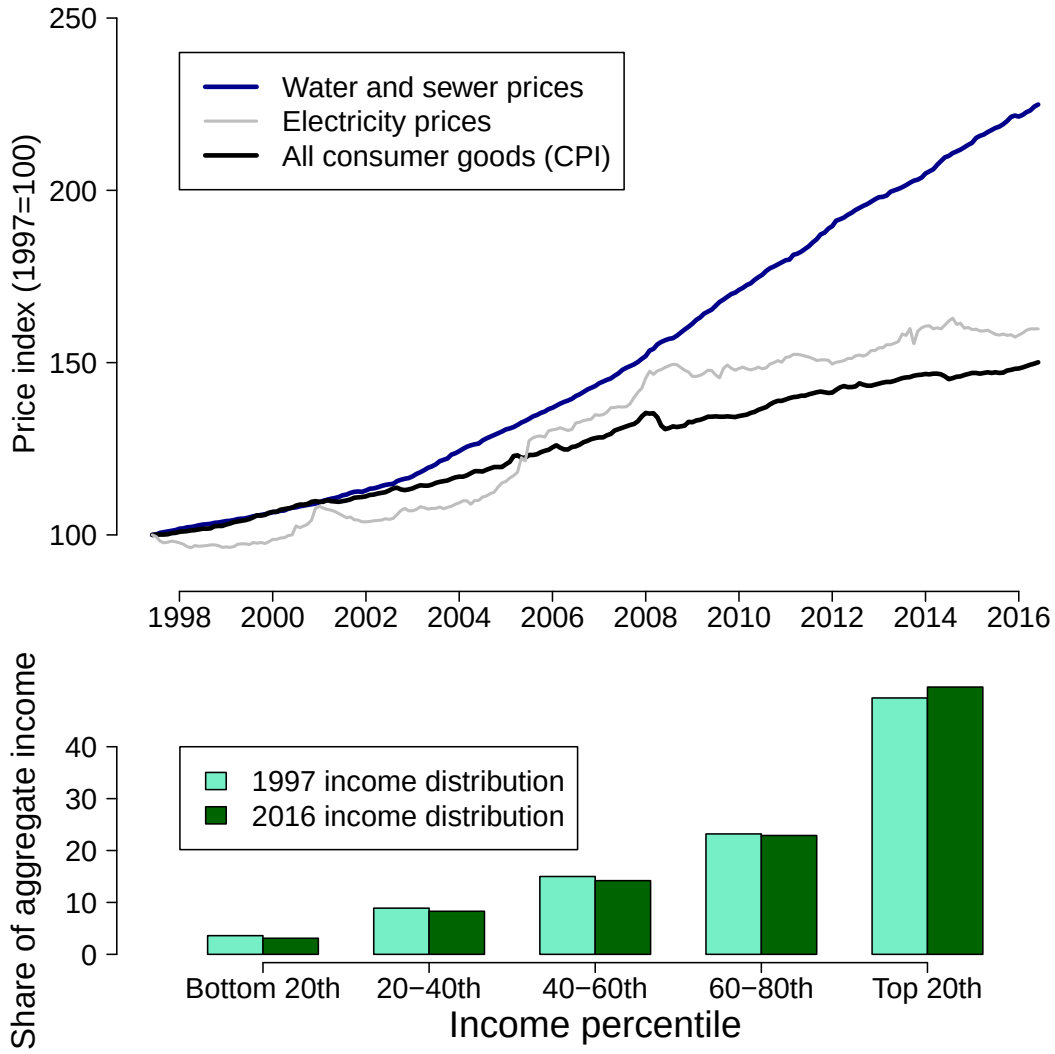

Figure 4.1: U.S. price indexes and income distribution over time.
Top Panel: Monthly price (U.S. city average, all urban consumers, seasonally adjusted) for all consumer goods, electricity, and water and sewer. Series begins in December 1997 (=100). Water and sewer price index includes trash collection. *Source:* U.S. Bureau of Labor Statistics. **Bottom Panel:** Share of aggregate income received by each fifth of households in 1997 and 2016. *Source:* U.S. Census Bureau, Current Population Survey, Annual Social and Economic Supplements.

impoverished residents, and the proportion of black residents within a Census block-group even after conditioning on poverty rates. Finally, we show that affordability policies that provide a lump-sum rebate to low-income households and are paid for by income taxes dominate policies that change marginal incen-

tives for water and sewer consumption.

Surprisingly, there are very few national estimates of water affordability and the geographic distribution of the most vulnerable communities is poorly understood. One exception is a recent paper that provides a large-scale geographic assessment of communities that are at-risk of water poverty, although this analysis relies on several limiting assumptions (Mack & Wrase, 2017). The authors evaluate water bills at a constant level much larger than typical household consumption levels, ignore geographic differences in water prices and consumption patterns, and evaluate all results at the Census tract level, which severely limits the validity and usefulness of those estimates. In the present manuscript, we approach a similar question with much more tenable assumptions. By focusing on behavior at the representative household level, our framework not only addresses the shortcomings of previous analyses but also provides a richer set of policy-relevant results that allow for policy simulations and assessment of distributional impacts.

Further, EPA's oft-used threshold for determining a "high burden" of water and sewer bills—whether combined water and sewer bills (CWSBs) exceed 4.5% of a community's median household income—has received increasing scrutiny as an adequate measure of a household's ability to pay for water and sewer services (National Academy of Public Administration, 2017; Teodoro, 2018; Mumm & Ciaccia, 2017). Common concerns are that using median household income at a community level poorly captures the burdens on the most vulnerable low-income residents and the 4.5% threshold (for combined water and sewer bills) is arbitrary. The origins of this median household income threshold can be traced back to EPA guidance for determining economic impacts of water quality reg-

ulations, but no formal justification for the level of the threshold was provided (U.S. EPA, 1995). Here, we demonstrate why using the full income distribution at a local level is imperative for measuring affordability. We do not focus on the normative question of whether EPA's 4.5% threshold is valid, but our methods align with alternative metrics and are general enough to accommodate alternative income-based affordability measures. Finally, there is no comparative evidence on the efficacy or distributional consequences of water affordability policies, which are beginning to be adopted at scale (California State Assembly, 2015; City of Philadelphia, 2015). Within different affordability policies, we show how underlying economic incentives can influence ultimate program effectiveness.

Overall, we provide evidence of the widescale water-affordability burden on low-income households. By assembling a novel data set covering 45% of the US population, we estimate that approximately 1 in 7 (13.6%) households in our sample have water and sewer expenditures that exceed 4.5% of their annual household income. We also show that redistributive programs that do not distort marginal incentives to consume water can be effective policy options.

4.1.1 Measuring affordability

EPA initially defined a water affordability metric as a tool to measure a community's financial capability for regulatory compliance (U.S. EPA, 1995, 1997). This threshold assigned a high burden of water and sewer expenditures if the typical CWSB exceeded 4.5% of a community's median household income. Critics of this measure often cite the arbitrary nature of the 4.5% threshold, the inabil-

ity of median income to capture the full distributional burden borne by poorest community members, and whether bills should be evaluated at average or necessary consumption levels (Teodoro, 2018). In this analysis, we follow the recently released panel recommendations from the National Academy of Public Administration (NAPA) for defining community affordability criteria for clean water services (National Academy of Public Administration, 2017). These recommendations include development of an improved affordability metric that is: (i) readily available from public data sources; (ii) clearly defined and understood; (iii) simple, direct, and consistent, (iv) valid and reliable according to conventional research standards, and (v) applicable for comparative analyses.

In line with current EPA guidance and forward-looking NAPA recommendations, we contrast the burden of water and sewer expenditures under varying definitions of water affordability. We adopt the threshold that affordable CWSBs should not exceed 4.5% of income, although we apply that standard at different income and geographic resolutions. Our analysis and methods can readily be applied to different thresholds. Second, we consider necessary levels of consumption (50 gallons per person per day, or gppd) as a more appropriate indicator for eligibility in affordability programs. Third, we put forward transparent and readily calculable metrics for ease of communication and decision-making by policymakers, thereby responding directly to NAPA's recommendations. Our preferred affordability measure is defined as the proportion of a community's population that pays more than 4.5% of annual household income on water and sewer service. The 4.5% threshold is readily scalable to different income thresholds. For context, the average US household spends 4.6% of their annual income on health insurance and 4.6% percent of their income on food away from home, according to the 2017 Bureau of Labor Statistics Consumer

Expenditure Survey.

4.2 Data and Methods

4.2.1 Data and calculations

Our primary data set contains water and sewer rates from 1,545 utilities that cover 92,445 census block groups from 521 counties across 42 states. This sample corresponds to approximately 145 million people, which comprises 45% of the total U.S. population as of 2016. This data set combines local water and sewer rates, number of service accounts, average consumption, climate characteristics, and a variety of socioeconomic indicators.

Water and sewer rates are obtained from two sources: rate surveys cataloged by the Environmental Finance Center at the University of North Carolina at Chapel Hill (EFC), current as of July 1, 2017, and the 2016 American Water Works Association (AWWA) Water and Wastewater Rate Survey. We aggregate water and sewer rates to the county level, weighted by the number of accounts in each utility within the county. To account for increasing-block rates, we approximate rate structures as a piecewise linear function of consumption with up to three rate blocks.

Average residential water and sewer usage is reported by the utilities in the AWWA survey but not by those in the EFC survey. We estimate the average consumption per person-day for EFC utilities with following linear model for

AWWA utilities,

$$\begin{aligned}\log(w_c) = \quad & \beta_0 + \beta_1 \log(Pop_c) + \beta_2 \log(y_c) + \beta_3 BaseCharge_c \\ & + \beta_4 \log\left(Rate5_c\right) + \beta_5 \log\left(Rate10_c\right) + \\ & + \textstyle\sum_{z \in Z} \gamma_z 1(CZ_c = z) + \sum_{s \in S} \delta_s 1(State_c = s) + e_i\end{aligned} \tag{4.1}$$

where $\log(w_c)$ is the log of average consumption per person-day in gallons in county c, Pop_c is county population, y is median household income, $BaseCharge_c$ is the fixed minimum service fee, $Rate5_c$ and $Rate10_c$ are the volumetric rates charged between 5 and 10 ccf and above 10 ccf per month, $1(CZ_c = z)$ is a dummy variable equal to 1 if county c is in climate region z and 0 otherwise, and $1(State_c = s)$ is a dummy variable equal to 1 if county c is in state s and 0 otherwise, and e_i is the idiosyncratic error term. The estimated parameters of Equation 4.1 are used to predict $\log(\hat{w}_c)$ for EFC utilities.

We consider four levels of geographic resolution (or aggregation). The unit of observation in the lowest resolution (highest aggregation) is a county, which considers a representative household that has its characteristics set at the county median (or average, depending on the variable). Similarly, in the second resolution level, each block group is represented by a single household with the block-group median/average characteristics. In the third resolution level, block groups are represented by 16 households that share the same socio-demographic characteristics (the block group median/average) but with different incomes corresponding to the center of US Census income brackets. Each of the 16 households has a different weight corresponding to the block-group income distribution. The fourth and highest resolution considers a continuum of households with income varying within each block group. In this latter approach, we interpolate a 16-node income distribution using monotone pre-

serving cubic splines to obtain a continuous cumulative probability distribution function.

County-level average per capita consumption is assigned to households within Census block groups as follows. Let b and c denote the block group and county, and i denote a node of the 16-node income distribution given by Census. For each county, we assume the average per capita water consumption (w_c) corresponds to a household with the median income ($\bar{y}_c$). Using an income elasticity (ϵ_y) of 0.1 (Havranek, Irsova, & Vlach, 2018), we adjust the per capita water consumption for each node in the income distribution (y_i) and multiply by the average household size (h_{bc}) to obtain representative household consumption at the monthly level, given by:

$$W_{ibc} = 30 \times h_{bc} \times w_c \times \left(\frac{y_i}{\bar{y}_c}\right)^{\epsilon_y}. \tag{4.2}$$

Then, for each block group and using the county average rates, we form a 16-node distribution of combined water and sewer bills (CWSBs).

4.2.2 Estimation of socioeconomic and demographic conditional correlations

We investigate whether local affordability of water is correlated with a set of local socioeconomic and demographic factors including:

1. Population density, measured in persons per square mile.
2. The percentage of the block group population that identify their race as Black or African American alone.

3. The percentage of the block group population that identify being of Hispanic or Latino origin.
4. The percentages of households with income below the Census Bureau poverty threshold, and with income between one and two times that threshold.
5. The median age of housing units.
6. The median gross rent as a percentage of the household income.
7. The average household size.
8. The percentage of rented units relative to all occupied units.

Local affordability is calculated using the distribution of income and CWSBs within each block group. In particular, the affordability metric of interest is the percentage of households with CWSBs above 4.5% of their income calculated with estimated consumption levels, which we represent by $\hat{U}_{bc}$.

We estimate conditional correlations by estimating the parameter vector $\mathbf{\Gamma}$ in

$$\hat{U}_{cb} = \mathbf{X}'_{cb}\mathbf{\Gamma} + \sum_{z \in Z} \gamma_z 1\left(CZ_c = z\right) + \sum_{s \in S} \delta_s 1\left(State_c = s\right) + u_{cb}, \tag{4.3}$$

where $\mathbf{X}_{cb}$ is the vector of local socioeconomic and demographic factors defined above. The equation above includes climate zone and state fixed effects and an idiosyncratic error term, u_{cb}.

A challenge in estimating the standard errors of the previous equation is the fact that $\hat{U}_{bc}$ is measured with error, but the nonlinearity in $\hat{U}_{bc}$ restricts our ability to derive asymptotic properties of the estimators analytically. For this reason, we implement a residual bootstrap strategy to obtain the appropriate

standard errors (MacKinnon, 2006). The details of the procedure can be found in Appendix C.2.

4.2.3 Policy simulations

In our policy simulations, households are eligible for assistance if annual household expenditures on water and sewer services at 50 gppd is greater than or equal to 4.5% of their annual income. The assistance programs affect water consumption by either changing water prices or income of a household. In both cases, we recalculate households' water consumption and expenditures, based on estimated consumption levels, as a result of the policy.

Assuming constant price elasticity ϵ_p, a price change d leads to a change in consumption of d^{ϵ_p}; hence, the new consumption is given by $W_{cb} \times d^{\epsilon_p}$. For a change in income, with a constant income elasticity, we use equation (C.3) to recalculate water consumption after an income tax or transfer. Our simulations assume a constant price elasticity $\epsilon_p = -0.3$ (Dalhuisen, Florax, De Groot, & Nijkamp, 2003) and constant income elasticity $\epsilon_y = 0.1$ (Havranek et al., 2018). See Appendix C.4 for sensitivity analyses on these parameter choices. The adjusted consumption values are then used to calculate the new expenditure on water and sewer services.

To determine the size of the programs in each county, we first adjust water consumption for assisted households based on a 50% rate discount (in all rate blocks). Then, we calculate the amount necessary to fund the discounts; this amount determines the size of all four programs. We set uniform lump-sum transfers that match the size of the rate discount program. Similarly, we cal-

culate the uniform income tax rate and the price increase needed to fund the assistance programs in each county.

4.3 Results

4.3.1 High-resolution income data and local prices are critical for measuring household-level water affordability

We calculate the number of households whose annual combined water and sewer bills exceed 4.5% of their annual household income for different definitions of income and consumption in Table 4.1. Comparing average water and sewer consumption at the county level with 4.5% of county-level median household income identifies virtually no households with unaffordable water and sewer service in our sample. But this is clearly misleading as it tells us only about a household with median income. Narrowing the geographic area at which we apply our median-income threshold provides a better approximation of local income distributions. Using median household income at the Census block-group level (block groups typically contain 600–3000 people) induces a modest increase in the proportion of households that exceed the water and sewer affordability threshold—0.8% for 50 gallons per person per day (gppd), 1.9% for estimated average consumption, and 4.8% for 100 gppd. 50 gppd is intended to capture "essential" water consumption, whereas estimated average consumption captures prevailing water use. The sample average of estimated consumption is 78.1 gppd.

By using income-group midpoints of a 16-node income distribution at the block-group level to calculate affordability (see Appendix C.2), we determine that 13.6% (10.0%) of households have CWSB greater than 4.5% of income for estimated (essential) consumption levels. Results are nearly identical when using an interpolated income distribution. These quantities are 4–12 times greater than quantities calculated with coarser income information based on median household income. This result is driven by the fact that we are able to identify more households with unaffordable water by using more granular data on income. Income aggregation reduces our ability to identify households in the very low portion of the income distribution. When calculating the distributional burden of water and sewer expenditures it is imperative to capture the local income distribution in its entirety.

4.3.2 One out of every seven households spends more than 4.5% of annual income on water and sewer bills

As shown in Table 4.1 and Figure 4.2, around 14% of households in our sample, approximately 1 out of every 7 households, have current annual water expenditures that exceed EPA's 4.5% affordability threshold. This estimate reflects the fact that more than 20-million households in our sample have unaffordable water and sewer service; a nationally representative estimate of the number of households with unaffordable water and sewer services would be much larger. We obtain this estimate by applying the 4.5% affordability threshold to representative households within a 16-node income distribution within Census block groups. Although applying an affordability threshold at 4.5% of income is ar-

Table 4.1: Percentage of households who pay more than 4.5 percent of annual household income on combined water and sewer bills by income and consumption data resolution.

			Consumption level		
Unit of analysis	**Income metric**[a]	**Estimated consumption**[b]	**Estimated**	**50 gppd**	**100 gppd**
County	Median household income	Average	0.00%	0.00%	0.09%
Block group	Median household income	Income-adjusted	1.88%	0.79%	4.83%
Block group	Income bracket center	Income-adjusted	13.57%	10.03%	18.54%
Block group	Income distribution	Income-adjusted	13.58%	10.10%	18.54%

[a]Results are presented for three income metrics: (i) "Median household income" represents median incomes at the county or block-group level; (ii) "Income bracket center" measures incomes at the midpoint of income brackets evaluated at the block-group level; and (iii) "Income distribution" represents the 16-node income distribution with interpolation between nodes evaluated within Census block groups.
[b]We use two consumption metrics: (i) "Average" consumption levels are estimated (see Section 4.2, sample average estimated consumption is 78.1 gppd); (ii) "Income-adjusted" consumption levels adjust for household income levels using an income elasticity of 0.1 for the relevant income metric. Consumption levels of 50 and 100 gallons per person day (gppd) are fixed across all households.

bitrary, this threshold provides a useful benchmark to compare the burden of water and sewer expenditures across geographies and our methods can be applied at any income threshold.

We also calculate the proportion of households above this affordability threshold based on fixed levels of consumption per person-day. When calculating water and sewer bills based on 50 gppd—an amount less than typical consumption, but that would cover "essential" daily usage—we find that approximately 10% of households in our sample would exceed the water affordability threshold. This result suggests that affordability is affected by household consumption levels and for essential levels of consumption, one out of every ten households faces affordability issues.

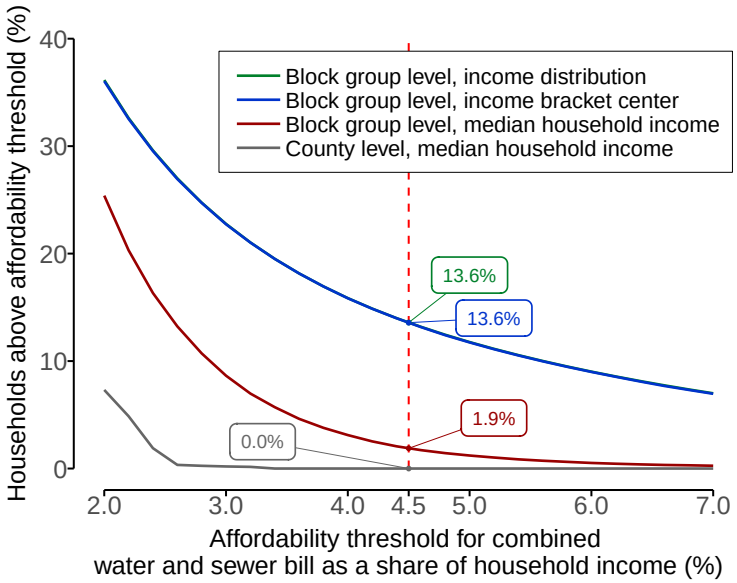

Figure 4.2: Proportion of households above affordability threshold based on water and sewer expenditure-shares of income, based on varying degrees of income data resolution.

4.3.3 Households in the lowest income bracket pay 8.1% of annual income on water and sewer bills

We calculate the burden of water and sewer bills for each income bracket in our data in Table 4.2. As shown, households with annual income less than $15,000 have an average combined water and sewer bill equal to 8.1% of household income. This statistic represents 11.4% of households in our sample. For contrast, households in the $45,000–$59,999 income group, near the US median household income, spend on average 1.7% of their annual household income on water and sewer bills. For the top income group—households earning $200,000 or more—this statistic is only 0.4%. This analysis reveals that the vast majority of households facing unaffordable water service are concentrated in the lowest income deciles.

Table 4.2: Expenditure on water and sewer relative to income by income bracket[a].

Annual income[a]	Frequency (thousands)	Percentage	Percentile	Average CWSB/Income[b]
Under $15,000	5,923	11.4%	11.4	8.1%
$15,000 to $24,999	4,988	9.6%	21.0	4.0%
$25,000 to $34,999	4,899	9.4%	30.5	2.7%
$35,000 to $44,999	4,620	8.9%	39.4	2.1%
$45,000 to $59,999	6,036	11.6%	51.0	1.7%
$60,000 to $74,999	5,092	9.8%	60.8	1.3%
$75,000 to $99,999	6,361	12.3%	73.0	1.1%
$100,000 to $124,999	4,517	8.7%	81.7	0.9%
$125,000 to $199,999	6,013	11.6%	93.3	0.6%
$200,000 and over	3,468	6.7%	100.0	0.4%

[a]Income distribution data are obtained from the U.S. Census 2016 5-Year American Community Survey and defined at the block-group level.
[b]Combined water and sewer bills (CWSB) are estimated.

4.3.4 Water and sewer bills are regressive, but less so, when compared with annual expenditure data

In other contexts, researchers have argued that contemporaneous income provides a poor proxy for the expenditure burdens of energy taxes, for example, because reported income fails to capture government transfers, retirement benefits, unreported income, or "lifetime income" (Hassett, Mathur, & Metcalf, 2009; West & Williams III, 2004). We follow these authors and assume that current expenditures provide a proxy for income throughout the life-cycle as postulated by the permanent-income hypothesis. In Figure 4.3, we plot nationally representative water and sewer expenditures, as reported in the 2016 Consumer Expenditure Survey (CES), as a share of annual income and as a share of total annual household expenditures. The results based on income-shares are consistent with our previous results. However, water expenditures as a share of total expenditures are less regressive than when calculating the water-expenditure share of income. Despite this result, water-expenditure shares of total expen-

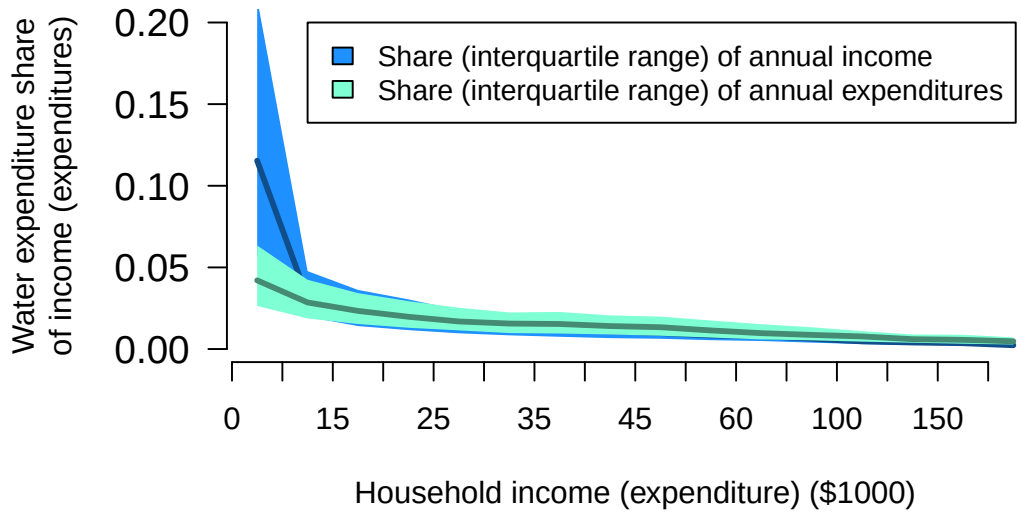

Figure 4.3: Incidence of US water bills using Consumer Expenditure Survey data.

Notes: Median annual water expenditure share (solid lines) and 25th–75th percentile range (shaded regions) as a proportion of annual household income or total annual household expenditures. Annualized water expenditures presented are calculated from survey responses of quarterly 'water and other public services' expenditures at primary place of residence. Statistics include information only for survey respondents that reported positive water expenditures. Medians and percentile ranges are weighted to account for the nationally representative survey sampling design. *Source:* Authors' calculations using Public Use Micro Data from 2016 Consumer Expenditure Survey.

ditures decrease monotonically as income rises, which suggests that water bills are still regressive when using alternative formulations of lifetime income. Using the CES data, we find little evidence of regional or racial heterogeneity in regressivity, as illustrated in Figure 4.4.

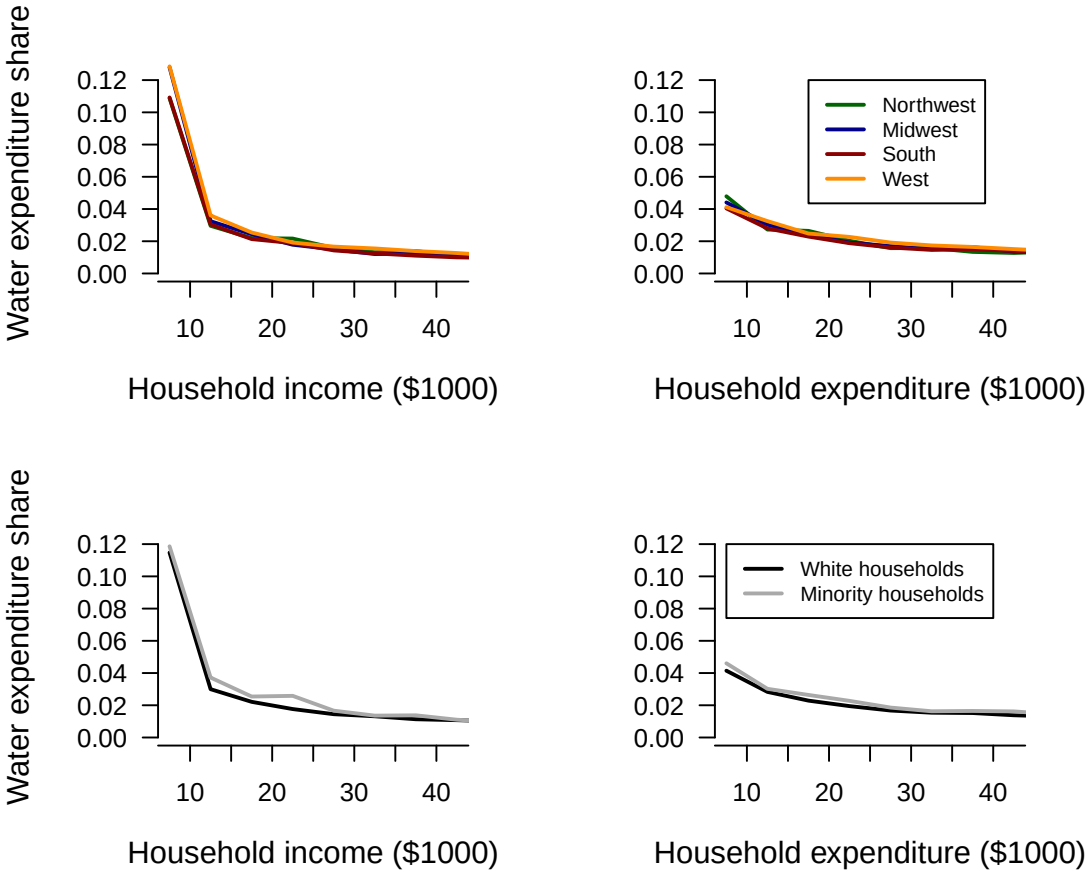

Figure 4.4: Incidence of water bills by region and race using Consumer Expenditure Survey data.

Notes: **Top panels:** Statistics presented are medians weighted to account for the survey sampling design. Geographic definitions are adopted from BLS. **Bottom panels:** Statistics presented are medians weighted to account for the survey sampling design. White households are defined as survey households in which the survey respondent identified as Caucasian. Minority households are defined as survey households in which the survey respondent as anything other than Caucasian. *Source:* Authors' calculations using Public Use Micro Data from 2016 Consumer Expenditure Survey.

4.3.5 Water affordability concerns are pervasive across the US, driven by the local income distribution

Geographically, we find some differences in water affordability across the US. In panel (a) of Figure 4.5, we plot the proportion of households with unaffordable water within each county. We calculate affordability based on estimated consumption levels (see Section 4.2). Some counties in the desert Southwest display high levels of unaffordable CWSBs, with rates of unaffordable water exceeding 25% of households. Several states in the Southeast also possess counties with high rates of unaffordable water and sewer bills.

County-level comparisons, however, mask important heterogeneity at a finer geographic scale. In panels (b)-(d) of Figure 4.5, we plot the same metric evaluated at the Census block-group level. This analysis reveals pockets of water affordability concerns at a more local level. In the Southeast (panel (e)), we observe a patchwork of block groups with high rates of households with unaffordable water bills. Even in the relatively wealthy Northeast (panel (d)), we identify many Census block groups with greater than 25% of households facing unaffordable water and sewer services.

Local maps also illustrate the importance of analyzing water affordability issues at a high resolution. Figure 4.6 shows the percentage of households facing unaffordable water services in block groups of counties corresponding to three large urban areas: Atlanta, GA, Chicago, IL, and Portland, OR. In these cities, we observe clusters of block groups with households facing unaffordable water, which in many cases geographically correlates with low-income areas. The high resolution of the data also allows us to identify several isolated pockets of water

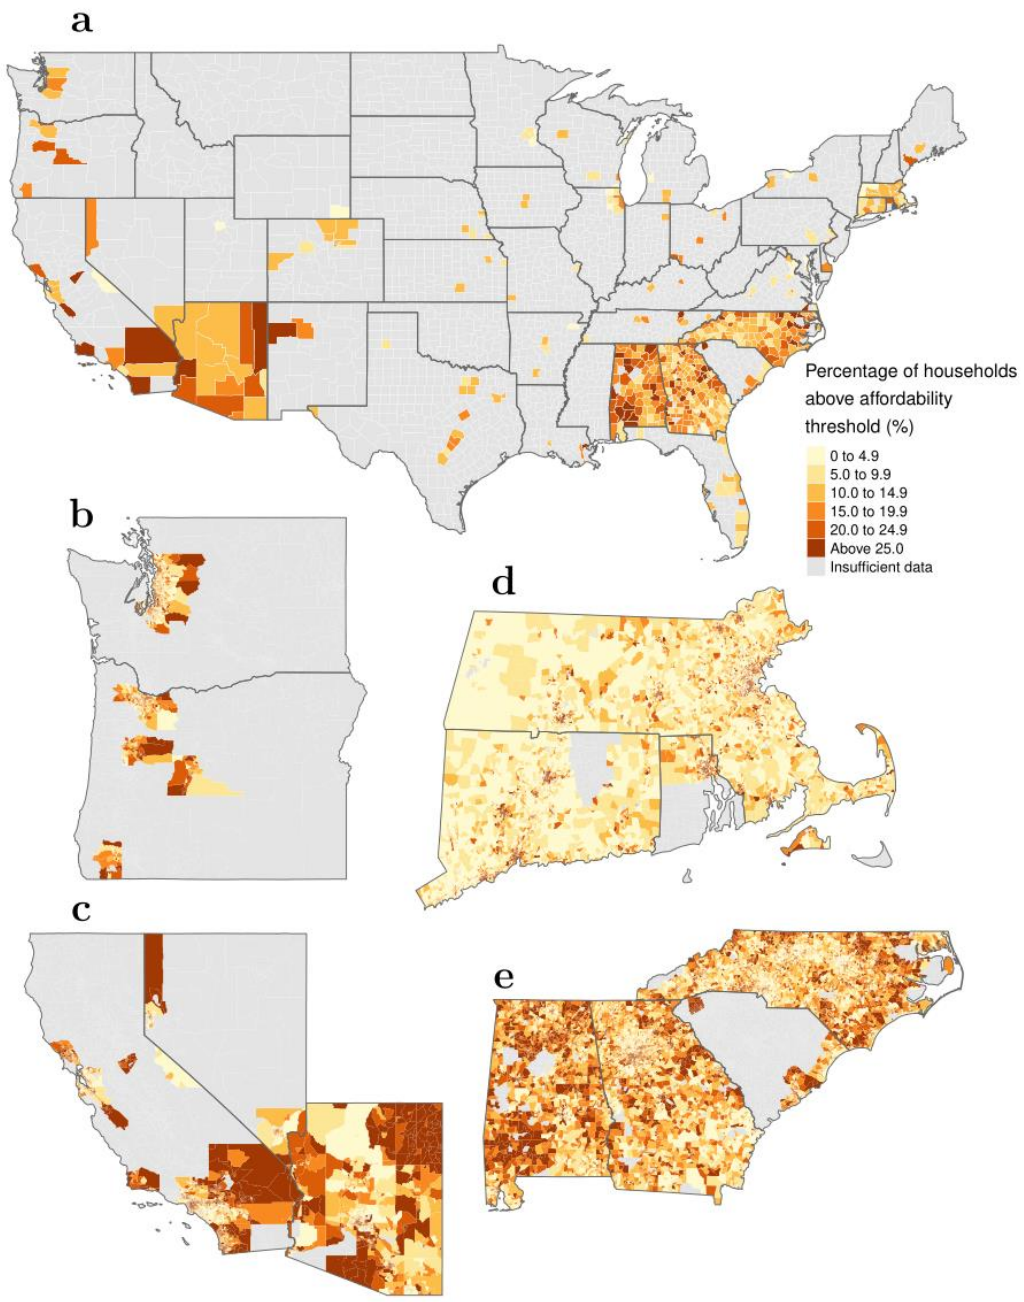

Figure 4.5: Geographic distribution of water affordability using estimated levels of consumption to calculate combined water and sewer bills.

Notes: Shaded colors show the percentage of households within each county (in **a**) and Census block group (in **b**–**e**) that have combined water and sewer bills above 4.5 percent of annual household income.

affordability concerns.

Overall, we find evidence that water affordability concerns are pervasive in the Southwest and Southeast. However, we also uncover serious concerns within states and within urban areas across the US. Because of these findings, we conclude that affordability concerns are inherently a local issue dictated by the distribution of income within a community.

4.3.6 Water affordability concerns are significantly correlated with select community characteristics

We conduct a statistical analysis to test whether any associations between the proportion of households in a block group with unaffordable water and sewer service are statistically significantly correlated with socioeconomic or other community characteristics. To develop conditional statistical tests of correlation, we regress the proportion of households above the affordability threshold on community characteristics and state and climate-zone fixed effects (see Section 4.2 and Appendix C.2 for our detailed statistical methodology).

Conditional correlation coefficients are presented in Table 4.3, along with bootstrapped 95% confidence intervals, between the proportion of households facing unaffordable CWSBs and select socioeconomic variables. The percentage of the population below the federal poverty limit is strongly associated with the prevalence of water affordability concerns. A one percentage point increase in the number of households below the poverty limit is associated with a 0.600 [0.573, 0.626; 95% CI] percentage points increase in the number of households

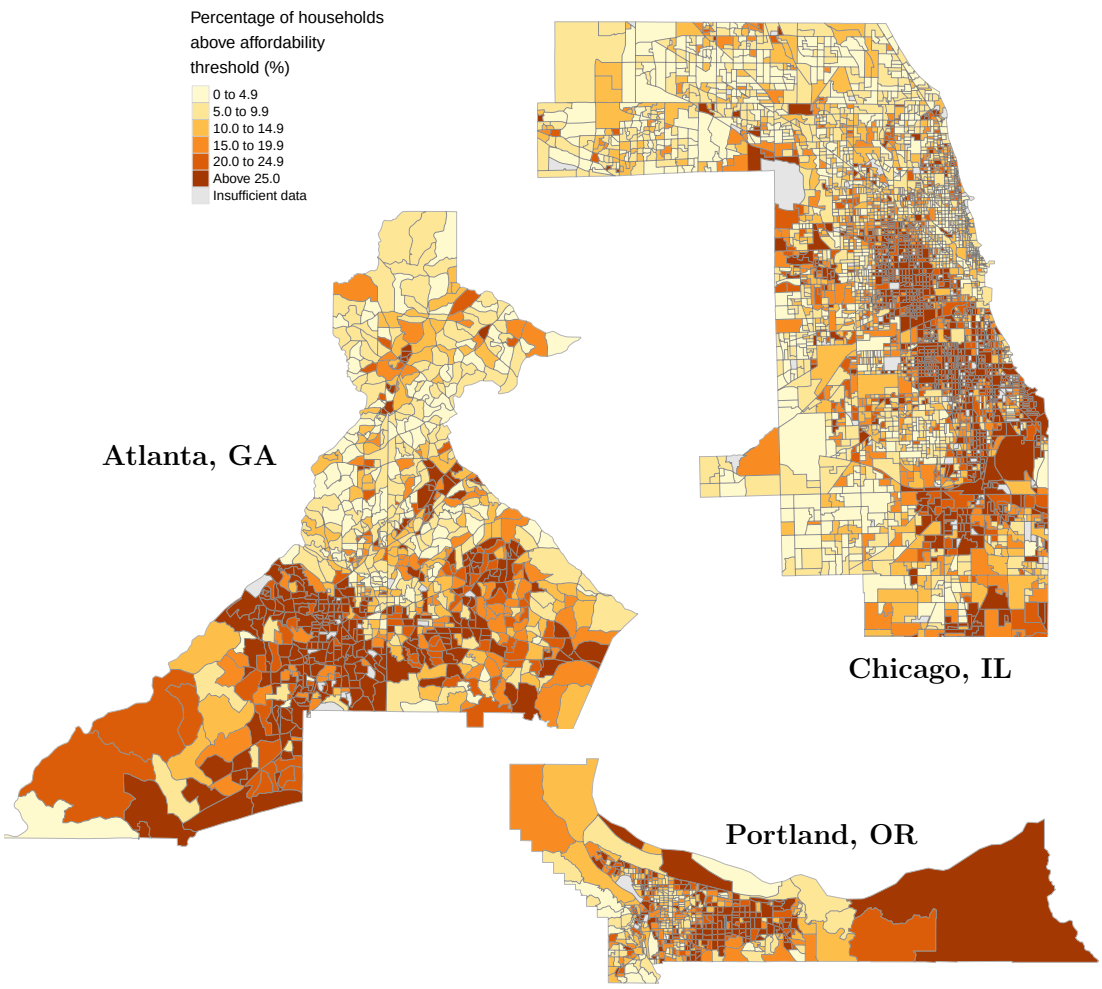

Figure 4.6: Geographic distribution of water affordability in block groups within urban areas.

Notes: Results are presented for Atlanta, GA (DeKalb and Fulton counties), Chicago, IL (Cook county), and Portland, OR (Multnomah county). Shaded colors show the percentage of households within each Census block group that have combined water and sewer bills above 4.5 percent of annual household income. Combined water and sewer bills are based on estimated levels of consumption.

Table 4.3: Conditional correlations between water affordability and select socioeconomic characteristics.

	Coef.	Asymp. SE	Bootstrap SE	Bootstrap 95% CI
log(Population density) (Persons/Sq. mi)	-0.4935	0.0289	0.1129	[-0.7147,-0.2722]
Average household size (Persons)	1.7333	0.0795	0.3939	[0.9612, 2.5054]
log(Volumetric rate at 5–10 ccf) (USD/1,000 gallons)	7.6766	0.1246	1.5523	[4.6342, 10.7191]
Base charge relative to CWSB at essential level (%)	0.0785	0.0026	0.0310	[0.0177, 0.1393]
Households below poverty level (%)	0.5998	0.0042	0.0135	[0.5734, 0.6262]
Households between 1 and 2× poverty level (%)	0.1478	0.0036	0.0123	[0.1236, 0.1719]
Median gross rent relative to income (%)	0.0988	0.0036	0.0076	[0.0839, 0.1137]
Occupied units that are rented (%)	0.0257	0.0019	0.0076	[0.0108, 0.0405]
Median age of housing unit (Years)	0.0129	0.0021	0.0118	[-0.0103, 0.0361]
Population identified as Black/African American (%)	0.0325	0.0020	0.0103	[0.0124, 0.0526]
Population identified as Hispanic/Latino (%)	0.0159	0.0025	0.0168	[-0.0170, 0.0488]
Constant	-32.2983	0.6462	8.9454	[-49.8310,-14.7655]
State fixed efects			Yes	
Climate zone fixed efects			Yes	
Observations			76240	
R^2			0.6466	
F-statistic			2322.96	

[a]Dependent variable is the percentage of households in a block group above the 4.5 percent water affordability threshold calculated with estimated consumption levels. The mean of the dependent variable is 15.78 and its standard deviation is 13.93. Summary statistics for other variables are presented in Appendix C.1. Bootstrapped confidence intervals should be used for inference. See Appendix C for detailed variable definitions and a description of the bootstrap. All variables are defined at the block-group level.

above the affordability threshold. Additionally, we find a significant positive relationship between water affordability concerns and the proportion of black households within a community (0.033 [0.012, 0.053; 95% CI]) even after controlling for poverty levels. However, no significant relationship was found between affordability and the proportion of Hispanic residents (0.016 [-0.017, 0.049; 95% CI]). We also find a small positive correlation between affordability concerns and the proportion of renters within a block group and the median cost of rent relative to income.

Additionally, by using the natural log of population density as a proxy of urbanicity, we find that population density has a negative association with the proportion of households above the affordability threshold. A one log-point increase in population density is associated with a -0.494 [-0.715, -0.272; 95% CI]

percentage point decrease in the number of households above the affordability threshold. The magnitude of this effect, however, is quite small. In other words, a one-percent increase in population density is associated with a -0.0049 percentage point decrease in affordability concerns. Additionally, we believe that we are more likely to falsely assign rural households to utilities when they might in fact not receive public water or sewer service (e.g., rural households are more likely to have septic systems and thus not pay for sewer services directly). As a result, we conclude that water affordability is a concern for both urban and rural areas.

We include two variables that capture the role of water rate-setting practices. One variable captures the mean volumetric price for monthly consumption between 5 and 10 ccf. This variable is positively associated with water affordability concerns. A one log-point increase in average water rates is positively correlated (7.677 [4.634, 10.719; 95% CI]) with the proportion of households above the affordability threshold. Put another way, a one-percent increase in volumetric water rates is associated with a 0.077 percentage point increase in affordability concerns. A second variable captures the proportion of a typical customer's bill (evaluated at 50 gppd) that is composed of the fixed access charge. This variable is also positively correlated with the proportion of households with unaffordable water (0.079 [0.018, 0.393; 95% CI]), which suggests that affordability concerns are not driven entirely by the volumetric price of water and sewer services, but also the fixed service fee.

Overall, the proportion of impoverished households within a block-group and average water prices are strongly associated with unaffordable water. We also find evidence that the proportion of black households is correlated with un-

affordable water after conditioning on poverty levels and other socioeconomic characteristics. We do not, however, find a significant relationship between the prevalence of unaffordable water and the proportion of Hispanic residents within a block group. Additionally, we find positive correlations between water affordability concerns and household size as well as median rents as a proportion of household income.

4.3.7 Affordability policies that provide lump-sum rebates for low-income households and are funded by income taxes are most effective

Affordability policies can reduce the burden of water and sewer bills for low-income customers, although there is virtually no comparative research highlighting the relative effectiveness of different types of programs despite policies being adopted at scale. We consider four illustrative policy options that differ in how programs reduce water and sewer expenditures for low-income customers and in how the programs are funded. In our scenarios, low-income assistance takes the form of a uniform lump-sum transfer or a 50% rate discount for eligible households. These programs are funded either by a uniform water rate increase or a local income tax on relatively wealthy households. All affordability programs are assumed to be administered at the county level. The total size of the programs is equivalent and set equal to the dollar amount needed to cover the 50% rate discount option. The average-income tax increase is 0.1 percentage points and the average lump-sum transfer is $34.6 per month. We assume general equilibrium changes (e.g., changes in labor supply) in response to small

income changes are negligible. Additionally, we abstract from costs associated with policy implementation.

Our policy simulations are simplistic by design, although they possess the key elements inherent in many water affordability policies (California State Assembly, 2015; City of Philadelphia, 2015). In our framework, households above the 4.5% affordability threshold for essential use (50 gppd) are eligible for aid and those below the threshold are not. These options are illustrative and abstract from local regulations that prohibit using water prices for redistributive purposes and any prevailing water affordability programs or rate structures (e.g., "lifeline" rates) that are currently in use.

We compare the effectiveness of different program combinations in Figure 4.7. In the top right panel, we show average expenditure shares for the business-as-usual scenario and for each of our four policy options. Each program reduces the average number of households above the affordability threshold, although aid transfers reduce the number of households above the affordability threshold in the lowest income bracket more than rate reductions. At the 4.5% threshold, all programs considered reduce the 75th percentile of CWSBs to less than 7% of annual income, with transfers reducing it further to approximately 5%. In the top right panel of Figure 4.7 we show changes in the combined water and sewer bill for each policy. All policies substantially reduce the total bill for households in the lowest income bracket, although rebates decrease the net cost to low-income customers by the greatest degree. Raising revenue by price increases for relatively wealthier households increases total bills relative to raising revenue via an income tax.

Additionally, in the bottom left panel of Figure 4.7, we plot the change in

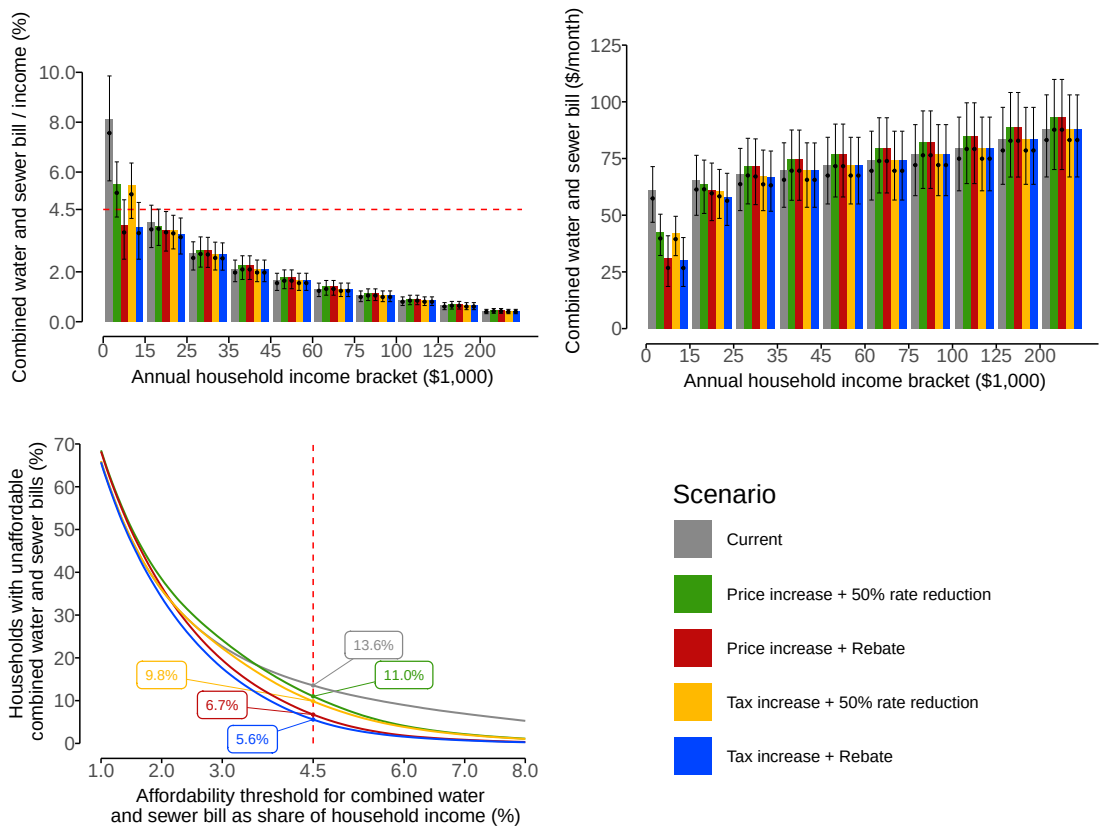

Figure 4.7: Results of program simulations.

Notes: **Top left panel:** Average expenditure shares on water and sewer service by income bracket for business-as-usual and each policy option. **Top right panel:** Combined water and sewer bill by income bracket for business-as-usual and each policy option. **Bottom left panel:** Distribution of sample with unaffordable water and sewer expenditures based on affordability threshold for business-as-usual and each policy option. In all bar charts, whiskers show the interquartile ranges and dots represent median values.

the number of households above the water affordability threshold for each program. Programs designed with income transfers rather than rate reductions can reduce the number of households above the 4.5% affordability threshold from 11.0% to 6.7% if funded by rate increases and from 9.8% to 5.6% if funded by income taxes. Interestingly, rate-reduction programs funded by a price increase evaluated at a 2.0–3.0% affordability threshold can actually worsen outcomes relative to the status quo. For programs of a similar size, structuring water affordability aid as an income transfer funded by income taxes dominates policy

options that alter the unit price of water and sewer consumption. As a practical matter, an income transfer could take the form of individual-specific credits on customer bills (so long as they are not misperceived as a reduction in the price of water (Wichman, 2017, 2014)) or a rate structure in which households pay different fixed access fees for water and sewer services. This finding is a result of the relative sensitivity of water and sewer consumption to changes in price and income. The price elasticity of water and sewer demand is greater in absolute magnitude than the income elasticity of water and sewer demand (Dalhuisen et al., 2003; Olmstead, Hanemann, & Stavins, 2007; Wichman, 2014; Klaiber, Smith, Kaminsky, & Strong, 2014; Wichman, Taylor, & von Haefen, 2016; Havranek et al., 2018). We perform a sensitivity analysis on these parameters in the Supporting Information. Because rate reductions distort marginal incentives for households to consume water more than do income transfers, low-income water customers consume more water as a result of affordability policies directed at making water cheaper. This feedback counteracts the goal of the affordability program. As a result, it is important to understand the demand implications of water affordability policies.

4.4 Discussion and conclusion

Provision of affordable water and sewer service is a growing concern in the United States although the extent of the problem is not known and the effectiveness of corrective policy options are underexplored. In this chapter, we have compiled a database of water and sewer prices for approximately 45% of the United States population to estimate annual expenditures on water and sewer service. We find that nearly one in seven households spend more than 4.5% of

their annual household income on water and sewer service, and that affordability concerns are positively correlated with race after conditioning on poverty levels. Our analysis shows the importance of incorporating geographically resolute information on the local income distribution of residents. Results from policy simulations demonstrate that redistributive water affordability policies designed to provide lump-sum rebates to low-income customers that are funded by income tax increases on relatively wealthier individuals are more effective at reducing the number of households with unaffordable water and sewer services than policies that distort marginal incentives to consume water.

Our analysis, however, does have several limitations. Here, we outline these limitations and discuss detailed robustness checks that explore the implications of these limitations in Appendix C.5. First, our sample covers only 45 percent of the US population, which skews towards urban areas and is not representative of the US. Our population of interest, however, is US residents who receive water and sewer service from public or private water utilities, which mitigates this sample-selection concern. As we show in the Appendix (Figure C.6), the income distribution in our sample is virtually identical to that of the nation as a whole. Moreover, although our sample is not comprehensive, we are aware of no other data set of water and sewer rates matched with socioeconomic characteristics and estimates of consumption as comprehensive as ours.

Second, our results rely on a metric that has received increasing scrutiny as a useful tool for measuring affordability concerns in part because the 4.5% of median-household income threshold is arbitrary and that median income poorly captures the full income distribution. We have shown empirically the substantial difference that using MHI and the full income distribution can have

on measuring affordability. Additionally, two alternative metrics of affordability have been proposed recently and are gaining traction as useful policy tools (Teodoro, 2018). The first is an "affordability ratio" that captures the ratio of essential water and sewer expenditures to a subjective measure of disposable income, evaluated at the 20th percentile of income within a service area. The second is essential water and sewer expenditures in units of hours worked at the minimum wage. Our focus in this analysis is not on contrasting alternative metrics, but we perform a simple comparison in Appendix C.6. For the 20 overlapping cities in our sample and in Teodoro (2018), our preferred metric correlates strongly with these new metrics, which suggests these alternatives may not dominate an income-based threshold affordability metric at face value (see Table S7). This result is important as income-based thresholds are used in the vast majority of other means-tested assistance programs (e.g., the Supplemental Nutrition Assistance Program, the Low Income Home Energy Assistance Program, and California's proposed statewide Low-Income Water Rate Assistance Program). Additionally, our affordability metric is readily scalable and can be used holistically in two ways: (i) to identify communities with a high burden of water and sewer expenditures and (ii) to establish household-level eligibility in low-income water rate assistance programs.

Our affordability framework has several attractive features for policymaking. First, we note that all proposed affordability metrics rely on judgments about what is essential consumption and what defines low-income customers. Our framework is no different in that regard. But, we do not contend that using a threshold of 4.5% of income to determine affordability is inherently wrong. This threshold is an arbirtary judgment of what is deemed affordable, just like all other affordability metrics proposed for water and other goods. What we do

contend is that this threshold should not be applied to the median household in a given geography. If policymakers believe that a household spending more than 4.5% of income is unaffordable, our analysis provides a rigorous framework to estimate the number of households above that threshold based on publicly available data at a broad scale. Second, if a policymaker believes 3% or 6% is a preferable threshold, our framework is easily adaptable to inform and account for that decision, as illustrated in Figure 4.2. Third, we also provide a framework for estimating water consumption (and, thus, water bills) by adjusting average consumption for demographic and economic differences across representative households. Additionally, our framework easily permits, and is robust to, alternative assumptions about what is "essential" consumption (Table S5).

Two other limitations are worth noting. First, we do not know whether the representative customers in our sample are homeowners or renters (who may not pay for water and sewer services directly). If the costs of water and sewer services are passed-through to renters fully in the cost of their rent, affordability is still a concern, but it changes the incentives for efficient water use. We know of no large-scale data set that contains this information useful for the scope of our analysis. To mitigate this concern in our regressions, we control for the proportion of renters within a Census block group and housing rental rates as a proportion of income. Additionally, many utilities have existing rate structures for low-income customers across the US. We know of no large-scale database of these types of rate structures or a synthesis of what the eligibility requirements are. Many of these rate structures lower the marginal price for water consumption for customers with low incomes. In our policy simulations, we show that lowering the marginal price of water counteracts the effectiveness of

low-income water rates, which suggests that "lifeline" rates may be inefficient policies. Exploring ways to address affordability issues for renters and understanding the dynamics of local or more aggregate policies (e.g., state or national) are fruitful areas for future research.

In sum, we provide evidence of the widescale water-affordability burden on low-income households. By assembling a novel data set covering 45% of the US population, we estimate that approximately 1 in 7 (13.6%) households in our sample have water and sewer expenditures that exceed 4.5% of their annual household income. We also show that redistributive programs that do not distort marginal incentives to consume water can be effective policy options. Our analysis provides a consistent framework to evaluate the extent of the water affordability burden and policies ameliorate its worst consequences as municipalities and regulators grapple with policies to fund water infrastructure improvements equitably.

APPENDIX A

APPENDIX OF CHAPTER 2

A.1 Estimation procedures and welfare analyses

The first two sections in this appendix outline the derivation of sufficient statistics and the corresponding expressions for evaluating marginal welfare changes. The third section shows how market equilibria are solved to calculate implied operating costs and counterfactual outcomes with different tax levels. Finally, the fourth section explains how I estimate the structural model using GMM with high-dimensional fixed effects.

A.1.1 Derivation of sufficient statistic ε

Consider a composite flight $Q_{mt} = \sum_{k \in \mathcal{K}_{mt}} q_k$ for market m in period t. Let the attributes of the composite flight be the average of attributes among flights in this market, weighted by the relative market shares $s_{k|g}$ (or shares within nest). That is, we define the average pre-tax price $\tilde{p}_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} \tilde{p}_k$, average price $p_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} p_k$, average fuel use $f_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} f_k$, and average markup $\mu_{mt} = \sum_{k \in \mathcal{K}_{mt}} s_{k|g} \mu_k$. As discussed in section 2.6.1, I assume that, for a small change in tax fuel τ, markups and relative market shares do not change.

With no markup change, any marginal increment in the tax fuel is passed on to consumers

$$\frac{d\tilde{p}_j}{d\tau} = f_k. \tag{A.1}$$

Here, τ is a volumetric tax on jet fuel rather than a tax on emissions. Dividing

τ by the carbon intensity of jet fuel burn ($h = 0.0134$ ton CO_2/gallon) gives the tax per ton of CO_2.

Summing equation (2.20) over all products in the market, it follows that

$$\frac{dQ_{mt}}{d\tau} \equiv \sum_{k \in \mathcal{K}_{mt}} \frac{dq_k}{d\tau} = \sum_{k \in \mathcal{K}_{mt}} \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j} \frac{dp_j}{d\tilde{p}_j} \frac{d\tilde{p}_j}{d\tau} = (1+r) \sum_{k \in \mathcal{K}_{mt}} \sum_{j \in \mathcal{K}_{mt}} \frac{\partial q_k}{\partial p_j} f_j.$$

Thus,

$$\frac{dQ_{mt}}{d\tau} = (1+r) \sum_{j \in \mathcal{K}_{mt}} f_j \frac{\partial Q_{mt}}{\partial p_j}. \tag{A.2}$$

With no change in relative market shares after a marginal change in τ, it follows that price changes are proportional: $dp_j = s_{j|g} dp_{mt}$. Hence,

$$\frac{\partial Q_{mt}}{\partial p_j} = \frac{\partial Q_{mt}}{\partial p_{mt}} \frac{\partial p_{mt}}{\partial p_j} = s_{j|g} \frac{\partial Q_{mt}}{\partial p_{mt}}. \tag{A.3}$$

Combining (A.2) and (A.3) yields

$$\frac{dQ_{mt}}{d\tau} = (1+r) \sum_{j \in \mathcal{K}_{mt}} s_{j|g} f_j \frac{\partial Q_{mt}}{\partial p_{mt}} = (1+r)\, f_{mt} \frac{\partial Q_{mt}}{\partial p_{mt}} \tag{A.4}$$

Then, expressing (A.4) in terms of elasticities, it follows that

$$\frac{dQ_{mt}}{d\tau} = -(1+r) \frac{f_{mt}}{p_{mt}} Q_{mt} \varepsilon = -(1+r) \frac{1}{w_t} \frac{F_{mt}}{\tilde{p}_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon = (1+r) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon.$$

Thus,

$$\frac{dQ_{mt}}{d\tau} = (1+r) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon, \tag{A.5}$$

where $\varepsilon = \frac{\partial Q_{mt}}{\partial P_{mt}}$ is the elasticity of aggregate demand with respect to the market average ticket price—this is the sufficient statistics estimated in section 2.6.1. Note that $\varepsilon < 0$, so an increase in τ leads to a decrease in Q_{mt}.

A.1.2 Marginal welfare calculations with sufficient statistics

Using equation (A.1), rewrite equation (2.6) to evaluate a marginal change in consumer surplus using the composite flight Q_{mt}

$$\frac{dCS_{mt}}{d\tau} = -(1+r)\sum_{k=1}^{K_m} q_k \frac{d\tilde{p}_k}{d\tau} = -(1+r)\,Q_{mt}\sum_{k=1}^{K_m} s_{k|g} f_k.$$

Thus,

$$\frac{dCS_{mt}}{d\tau} = -(1+r)\,Q_{mt} f_{mt}. \tag{A.6}$$

With constant relative market shares under a marginal tax change, it follows that changes in quantities are proportional: $dq_k = s_{k|g} dQ_{mt}$. Then, rewrite equation (2.7) for marginal changes in profits using the composite flight as follows

$$\frac{d\Pi_{mt}}{d\tau} = \sum_{k=1}^{K_m}\left\{\mu_k \frac{dq_k}{d\tau} + \left[\frac{d\tilde{p}_k}{d\tau} - f_k\right] q_k\right\} = \sum_{k=1}^{K_m} \mu_k s_{k|g} \frac{dQ_{mt}}{d\tau} = \frac{d\Pi_m}{d\tau} = \mu_{mt}\frac{dQ_{mt}}{d\tau},$$

where the second step uses (A.1), canceling out the term for changes in markups. Then, using (A.5), it follows that

$$\frac{d\Pi_{mt}}{d\tau} = \mu_{mt}(1+r)\frac{\eta_{mt}}{w_{mt}}\frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt}\varepsilon. \tag{A.7}$$

Following a similar procedure for changes in tax revenue, rewrite equation (2.8) as

$$\begin{aligned}
\frac{dT_{mt}}{d\tau} &= \sum_{k=1}^{K_m}\left[q_k f_k + \tau f_k \frac{dq_k}{d\tau}\right] + r\sum_{k=1}^{K_m}\left[q_k f_k + \tilde{p}_k \frac{dq_k}{d\tau}\right] \\
&= (1+r)\,Q_{mt} f_{mt} + \sum_{k=1}^{K_m}(\tau f_k + r\tilde{p}_k)\, s_{k|g}\frac{dQ_{mt}}{d\tau} \\
&= (1+r)\,Q_{mt} f_{mt} + (\tau f_{mt} + r\tilde{p}_{mt})\frac{dQ_{mt}}{d\tau},
\end{aligned}$$

so that

$$\frac{dT_{mt}}{d\tau} = (1+r)\,Q_{mt}\left\{f_{mt} + (\tau f_{mt} + r\tilde{p}_{mt})\frac{\eta_{mt}}{w_{mt}}\frac{\tilde{p}_{mt}}{p_{mt}}\varepsilon\right\}. \tag{A.8}$$

Lastly, for damages, rewrite equation (2.9) as

$$\frac{d\Phi_{mt}}{d\tau} = \phi \sum_{k=1}^{K_m} f_k \frac{dq_k}{d\tau} = \phi f_{mt} \frac{dQ_{mt}}{d\tau},$$

so that

$$\frac{d\Phi_{mt}}{d\tau} = \phi f_{mt} \left(1 + r\right) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon, \tag{A.9}$$

where ϕ is the marginal climate damage per gallon of jet fuel burned.

Combining equations (A.6)–(A.9), the marginal change in welfare for market m in period t, using sufficient statistics μ_{mt}, η_{mt}, and ε, is given by

$$\frac{dW_{mt}}{d\tau} = \left(\mu_{mt} + r\tilde{p}_{mt} + \left(\tau - \phi\right) f_{mt}\right) \left[\left(1 + r\right) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon\right]. \tag{A.10}$$

The change in aggregate welfare for the sector is the sum of (A.10) evaluated for each market and period,

$$\frac{dW}{d\tau} \equiv \sum_{m \in \mathcal{M}} \sum_{t=1}^{4} \frac{dW_{mt}}{d\tau}, \tag{A.11}$$

where $\mathcal{M}$ represents the set of all markets. The marginal change to short-run private surplus is a similar expression, only eliminating the marginal damage term:

$$\frac{dSRPS_m}{d\tau} = \left(\mu_{mt} + r\tilde{p}_{mt} + \tau f_{mt}\right) \left[\left(1 + r\right) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon\right]. \tag{A.12}$$

The marginal abatement cost (MAC), then, divides (A.12) by the marginal change in emissions $\frac{de_{mt}}{d\tau} = h f_{mt} \frac{dQ_{mt}}{d\tau}$, where $h = 0.0134$ ton CO_2/gallon is the carbon intensity of fuel burn. Hence,

$$MAC_{mt}\left(\tau_0\right) \equiv \frac{\frac{dSRPS_{mt}}{d\tau}}{\frac{de_{mt}}{d\tau}} = \frac{\left(\mu_{mt} + r\tilde{p}_{mt} + \tau f_{mt}\right) \left[\left(1 + r\right) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon\right]}{h f_{mt} \left[\left(1 + r\right) \frac{\eta_{mt}}{w_{mt}} \frac{\tilde{p}_{mt}}{p_{mt}} Q_{mt} \varepsilon\right]},$$

so

$$MAC_{mt}\left(\tau_0\right) = \frac{1}{h} \left[\tau + \frac{\mu_{mt} + r\tilde{p}_{mt}}{f_{mt}}\right], \tag{A.13}$$

where $\tau_0 = \$0.044$/gallon is the baseline jet fuel tax. Aggregating for the entire sector, the marginal abatement cost is given by

$$MAC\left(\tau_0\right) \equiv \frac{\frac{dSRPS}{d\tau}}{\frac{de}{d\tau}} = \frac{1}{h}\left\{\tau + \frac{\sum_{m\in\mathcal{M}}\sum_{t=1}^{4}\left(\mu_{mt} + r\tilde{p}_{mt}\right)\left[\frac{\eta_{mt}}{w_{mt}}\frac{\tilde{p}_{mt}}{p_{mt}}Q_{mt}\right]}{\sum_{m\in\mathcal{M}}\sum_{t=1}^{4} f_{mt}\left[\frac{\eta_{mt}}{w_{mt}}\frac{\tilde{p}_{mt}}{p_{mt}}Q_{mt}\right]}\right\}. \tag{A.14}$$

Note that the expressions for the MAC do not depend on elasticity ε; this reflects the assumption that changes in quantities are proportional, so emissions change linearly with quantities.

A.1.3 Solving market equilibria and operating costs

As described in section 2.4.3, firms choose a vector of pre-tax prices $\tilde{\boldsymbol{P}}_{\boldsymbol{imt}}$ that maximize market profits taking the competitor prices as given

$$\tilde{\boldsymbol{P}}_{\boldsymbol{imt}} = \arg\max \sum_{k\in\mathcal{K}_{imt}} \left(\tilde{p}_k - c_k\right) s_k,$$

where $\mathcal{K}_{imt}$ is the set of flights offered by airline i in market m and period t. The first-order optimality conditions for each product is described in equation (2.17). In a Nash-Bertrand equilibrium, each price choice satisfies (2.17), forming a system of nonlinear equations with dimension equal to the number of products. To simplify notation, drop subscripts mt and index products in a given market from 1 to K. Then, we can represent the stacked first-order conditions as

$$\underbrace{\begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_K \end{bmatrix}}_{S} + \underbrace{\left\{\begin{bmatrix} \frac{\partial s_1}{\partial \tilde{p}_1} & \frac{\partial s_2}{\partial \tilde{p}_1} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_1} \\ \frac{\partial s_1}{\partial \tilde{p}_2} & \frac{\partial s_2}{\partial \tilde{p}_2} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial s_1}{\partial \tilde{p}_K} & \frac{\partial s_2}{\partial \tilde{p}_K} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_K} \end{bmatrix} \circ \begin{bmatrix} O_{11} & O_{21} & \cdots & O_{K1} \\ O_{12} & O_{22} & \cdots & O_{K2} \\ \vdots & \vdots & \ddots & \vdots \\ O_{K1} & O_{K2} & \cdots & O_{KK} \end{bmatrix}\right\}}_{J} \underbrace{\begin{bmatrix} \tilde{p}_1 - c_1 \\ \tilde{p}_2 - c_2 \\ \vdots \\ \tilde{p}_K - c_K \end{bmatrix}}_{\tilde{P}-C} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \tag{A.15}$$

where $\circ$ is the element-wise (Hadamard) product. Cells O_{ij} indicate product ownership and receive 1 if products i and j are offered by the same airline and 0 otherwise. Partial derivatives of shares with respect to prices are obtained from demand equation (2.15). Own and cross-price derivatives are given by

$$\frac{\partial s_k}{\partial \tilde{p}_k} = \frac{\alpha}{\lambda}(1+r)\, s_k \left[(1-\lambda)\, s_{k|g} + \lambda s_k - 1\right] \tag{A.16}$$

$$\frac{\partial s_k}{\partial \tilde{p}_j} = \frac{\alpha}{\lambda}(1+r)\, s_j \left[(1-\lambda)\, s_{k|g} + \lambda s_k\right], \tag{A.17}$$

where the tax term follows from $\frac{\partial p_k}{\partial \tilde{p}_k} = (1+r)$. Note that

$$\frac{\partial s_k}{\partial \tilde{p}_j} = \frac{\alpha}{\lambda}(1+r)\left[(1-\lambda)\frac{s_j s_k}{1-s_0} + \lambda s_j s_k\right] = \frac{\alpha}{\lambda}(1+r)\, s_k \left[(1-\lambda)\, s_{j|g} + \lambda s_j\right] = \frac{\partial s_j}{\partial \tilde{p}_k}.$$

Hence, since all products are in the same nest, the matrix of partial derivatives is symmetric. Because the ownership matrix is symmetric, the resulting matrix J is also symmetric. Define a auxiliary matrix Γ as

$$\Gamma \equiv \begin{bmatrix} \frac{1-\lambda}{1-s_0}s_1 + \lambda s_1 \\ \vdots \\ \frac{1-\lambda}{1-s_0}s_K + \lambda s_K \end{bmatrix} = \left(\frac{1-\lambda}{1-s_0} + \lambda\right) S = \left(\frac{1-\lambda s_0}{1-s_0}\right) S,$$

where s_0 is the share of the outside option (not flying). Then, we can write the matrix of partial derivatives as

$$\begin{bmatrix} \frac{\partial s_1}{\partial \tilde{p}_1} & \frac{\partial s_2}{\partial \tilde{p}_1} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_1} \\ \frac{\partial s_1}{\partial \tilde{p}_2} & \frac{\partial s_2}{\partial \tilde{p}_2} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial s_1}{\partial \tilde{p}_K} & \frac{\partial s_2}{\partial \tilde{p}_K} & \cdots & \frac{\partial s_K}{\partial \tilde{p}_K} \end{bmatrix} = \frac{\alpha}{\lambda}(1+r)\, S\Gamma' - \frac{\alpha}{\lambda}(1+r) \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ 0 & s_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & s_K \end{bmatrix}.$$

Thus, matrix J can be constructed using the following expression

$$J = \frac{\alpha}{\lambda}(1+r) \left\{ \left(\frac{1-\lambda s_0}{1-s_0}\right) SS' - \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ 0 & s_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & s_K \end{bmatrix} \right\} \circ O, \tag{A.18}$$

where O indicates the ownership matrix.

Operating costs. Using (A.15), baseline operating costs are the solution to the following equation

$$\widehat{\boldsymbol{C}}_{\boldsymbol{mt}} = \tilde{\boldsymbol{P}}_{\boldsymbol{mt}} + \widehat{J}_{mt}^{-1}\boldsymbol{S}_{\boldsymbol{mt}},$$

where vectors $\tilde{\boldsymbol{P}}_{\boldsymbol{mt}}$ and $\boldsymbol{S}_{\boldsymbol{mt}}$ are observed in the data; matrix $\widehat{J}_{mt}$ is calculated using (A.18), so it depends on $\boldsymbol{S}_{\boldsymbol{mt}}$ and estimated parameters $\hat{\alpha}$ and $\hat{\lambda}$.

Counterfactual equilibria. As described in section 2.4.3, the baseline marginal operating cost per passenger is given by $c_k = \tilde{c}_k - (w_t + \tau_0) f_k$, where $\tilde{c}_k$ is the constant marginal cost per passenger excluding fuel costs, w_k is the jet fuel cost per gallon, and f_k is the volumetric fuel consumption per passenger. Hence, for a tax change from τ_0 to some tax level τ, operating cost increases by $\Delta c_k = f_k (\tau - \tau_0)$.

Let $\boldsymbol{C}_{\boldsymbol{mt}}(\tau)$ denote the vector of operating costs after a tax change from the baseline. Then, the new Nash-Bertrand equilibrium satisfies

$$\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} = -\widehat{J}_{mt}^{-1}\boldsymbol{S}_{\boldsymbol{mt}}\left(\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}\right) + \widehat{\boldsymbol{C}}_{\boldsymbol{mt}}(\tau), \tag{A.19}$$

where $\widehat{J}_{mt,\tau}^{-1}$ is itself a function of $\boldsymbol{S}_{\boldsymbol{mt}}$, which in turn is a function of the pre-tax price vector $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}$. Since the right-hand side of (A.19) also depends on $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}$, the equilibrium is characterized by a fixed point problem.

I rewrite equation (A.19) as a root-finding problem, $\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}} + \widehat{J}_{mt}^{-1}\boldsymbol{S}_{\boldsymbol{mt}}\left(\tilde{\boldsymbol{P}}_{\boldsymbol{mt,\tau}}\right) - \widehat{\boldsymbol{C}}_{\boldsymbol{mt}}(\tau) = 0$, and solve it numerically using the trust region method with autoscale (Nocedal & Wright, 2006). For each market and period, I solve for the equilibrium pre-tax price vectors over a grid of tax levels. This tax level grid ranges

from $-\tau_0$ (removing the existing tax) to \$100/ton CO_2. Grid intervals are more dense around 0, to calculate more precise marginal changes from the baseline, and in the \$50–60 range, to increase the precision of revenue-neutral tax effects. To speed up calculations on the grid, I use the closest available solution as the initial guess (or the observed price vector for nodes neighboring the baseline). With equilibrium prices, I calculate equilibrium shares and quantities, which then are used to calculate consumer surplus, profits, tax revenues, emissions, and damages. For profits, in particular, I hold load factors (seat occupancy rate) constant when calculating total fuel expenditure per product.

A.1.4 GMM estimation with high-dimensional fixed effects

As discussed in section 2.6.2, estimating GMM with high-dimensional fixed effects and nonlinear parameters can create technical challenges. On the one hand, nonlinear parameters hinder the possibility of performing within transformation to eliminate fixed effects. On the other hand, leaving fixed effects as dummy variables increases memory and CPU requirements and may cause numerical instability when solving the problem with a large number of moment conditions. I address these issues by adapting the method proposed in Conlon and Gortmaker (2020), which I summarize below.

For a shorthand notation, let $\theta = [\alpha, \lambda]$ represent the vector of nonlinear parameters. Also, let $\theta_n = [\alpha_n, \lambda_n]$ represent the value of such parameters in the n-th iteration of the algorithm. We start with an initial guess θ_0. In this chapter, I define the initial guess to be the parameters estimated using two-stage least squares. Nevertheless, varying the initial guess within reasonable values does

not change the final estimates.

Step 1: concentrate out nonlinear parameters. Define variable Y_k^D by rewriting the estimating equation for demand (2.23) as

$$Y_k^D \equiv (\ln s_k - \ln s_0) + \alpha_n p_k - (1 - \lambda_n) \ln s_{k|g} = X_k^D \beta^D + \delta_i + \delta_{ot} + \delta_{dt} + \xi_k, \quad \text{(A.20)}$$

where fixed effects δ_i, δ_{ot}, and δ_{dt} are separated from the vector product characteristics. These fixed effects represent airline, origin location-by-quarter, and destination location-by quarter, respectively. Similarly, define Y_k^S by rewriting the estimating equation for supply (2.24) as

$$Y_k^S \equiv \underbrace{p_k - \mu_k(\alpha_n, \lambda_n)}_{c_k} = \rho F_k + \beta_i^S \text{Ramp-to-ramp}_k + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k$$

$$Y_k^S = X_k^S \beta^S + \gamma_{i,o} + \gamma_{i,c_1} + \gamma_{i,d} + \gamma_{i,c_2} + \gamma_t + \omega_k, \quad \text{(A.21)}$$

where the term $\mu_k(\alpha, \lambda)$ makes it explicit that the markup used to calculate costs is a function of nonlinear parameters α and λ.

Let X^D be the $N \times M_D$ matrix of product characteristics relevant for demand (not including price and share within nest), where N is the number of products (observations) in the data and M_D is the number of product characteristics excluding fixed effects; in this chapter, $M_D = 8$. Let β^D be the $M_D \times 1$ vector of linear coefficients associated with demand characteristics. Similarly, let X^S be the $N \times M_S$ matrix of product characteristics relevant for supply, where $M_S = 8$. Also, β^S is the $M_S \times 1$ vector of linear parameters associated with X^S. Here, β^S includes parameter ρ and airline-specific parameters β_i^S.

Step 2: absorb fixed effects. With the rearrangement of estimating equations outlined above, fixed effect terms in equations (A.20) and (A.21) can be elimi-

nated. Since there are multiple fixed effects in each equation, simple demeaning would not work. Instead, we can absorb these terms using the method of alternating projections, which iteratively demeans variables until convergence (Bauschke et al., 2003). Such procedure is analogous to repeatedly applying the Frisch-Waugh-Lovell theorem in order to partial out dummy variables. The resulting equations can be written as

$$\bar{Y}_k^D = \bar{X}_k^D \beta^D + \bar{\xi}_k$$
$$\bar{Y}_k^S = \bar{X}_k^S \beta^S + \bar{\omega}_k,$$

where bars on top of letters indicate demeaned variables.

Step 3: estimate linear parameters. The moment conditions in equation (2.25) can be rewritten as

$$\begin{bmatrix} E\left(\bar{Z}_k^D \bar{Y}_k^D - \bar{Z}_k^D \bar{X}_k^D \beta^D\right) \\ E\left(\bar{Z}_k^S \bar{Y}_k^S - \bar{Z}_k^S \bar{X}_k^S \beta^S\right) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

with its sample analogue expressed as

$$\underbrace{\frac{1}{N}\begin{bmatrix} \left(\bar{Z}^D\right)' & 0 \\ 0 & \left(\bar{Z}^S\right)' \end{bmatrix}\begin{bmatrix} \bar{Y}^D \\ \bar{Y}^S \end{bmatrix}}_{\tilde{Y}} - \underbrace{\frac{1}{N}\begin{bmatrix} \left(\bar{Z}^D\right)' \bar{X}^D & 0 \\ 0 & \left(\bar{Z}^S\right)' \bar{X}^S \end{bmatrix}}_{\tilde{X}}\begin{bmatrix} \beta^D \\ \beta^S \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

where Z^D and Z^S denote the demand and supply instrument matrices, as outlined in section 2.6.2.

Note that $\bar{Y}^D$ and $\bar{Y}^S$ are functions of θ. Then, for a given value θ_n, we can estimate the linear parameters using GMM with weight matrix W:

$$\begin{bmatrix} \beta^D(\theta_n) \\ \beta^S(\theta_n) \end{bmatrix} = \left[\tilde{X}' W \tilde{X}\right]^{-1} \tilde{X}' W \tilde{Y}(\theta_n). \tag{A.22}$$

Since β parameters are linear, this step can be estimated relatively quickly by using routines optimized for matrix computation.

Step 4: estimate nonlinear parameters. The nonlinear parameter vector θ is estimated using a two-step efficient GMM estimaton.[1] The estimated parameter satisfies

$$\theta^* = \arg\min_{\theta} NG(\theta)' \hat{W} G(\theta),$$

where $G(\theta)$ is the vector of moment conditions evaluated at θ and $\hat{W}$ is the GMM weighting matrix. In the first step, $\hat{W}_1$ is calculated based on the clustered residuals of a two-stage least squares regression. The solution to the first step is used to calculate matrix $\hat{W}_2$, and the second step solves for the nonlinear parameters that minimize the GMM objective function with this updated matrix.

In this chapter, I solve the GMM minimization problem over θ using a trust region method with auto-scale (Nocedal & Wright, 2006). For each iteration of θ (outer loop), the algorithm repeats steps 1 to 3 described above to calculate β^D, β^S, and residuals (inner loop).

A.2 Additional regression results and alternative estimators

This appendix complements the estimation results in section 2.6, showing first-stage regressions and results using alternative estimators and specifications.

[1] For details on this estimation procedure, see Hayashi (2000).

A.2.1 First stage of the sufficient statistics regression

Table A.1 shows the first-stage regression for the sufficient statistic estimates reported in section 2.6.1. In addition to regular F-statistic, this table also includes the conditional F-statistic, which provides a more conservative F-test that adjusts for the possibility of multiple weak instruments (Sanderson & Windmeijer, 2016). Nevertheless, both F-statistics are substantially above commonly adopted critical values, alleviating concerns about weak instruments.

Table A.1: First stage regression of the sufficient statistic estimation.

	ln(Agg. ticket price)
Airlines in market	−0.042
	(0.006)
Products in market	0.0002
	(0.0001)
Share of nonstop flights	0.072
	(0.044)
Potential legacy entrants	0.076
	(0.012)
Potential LCC entrants	0.002
	(0.004)
log(fuel expenditure)	0.562
	(0.011)
Fixed effects	
Origin-by-quarter	Yes
Destination-by-quarter	Yes
Observations	267,967
F-statistic	1422
Conditional F-statistic	632

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The conditional F-statistic is calculated following Sanderson and Windmeijer (2016).

A.2.2 Estimating the structural model with two-stage least squares

In this chapter, the demand and supply sides of the structural model are jointly estimated using the generalized method of moments (GMM). An alternative approach, used in Aguirregabiria and Ho (2012), is to estimate the demand and supply equations separately using two-stage least squares (2SLS). However, these two estimators are not equivalent. As Conlon and Gortmaker (2020) point out, estimating separate equations with 2SLS implicitly assumes that errors across equation are uncorrelated: $E\left[\xi_k \omega_k\right] = 0$. Even though this assumption may not hold in all applications, the 2SLS estimator is sometimes preferred because it is faster to estimate due to the linearity in parameters. In this section, I show the results of estimating the structural model using 2SLS and the first stage estimates when instrumental variables are used.

Table A.2 shows the estimates for the demand equation using 2SLS. This table shows that demand estimates using 2SLS and GMM (table 2.3) are almost identical. In particular for the key parameters of the model, α and λ, estimates from both models differ only at the third significant digits. To provide further evidence of instrument validity, table A.3 shows the first stage regressions for each instrumented variables, as well as the respective F-statistics. Though conditional F-statistics are substantially smaller than regular F-statistics, especially for price, these values are well above 13.95, the critical value at 5% with 3 endogenous variables and 7 excluded instruments (using values from Stock and Yogo (2005), as suggested in Sanderson and Windmeijer (2016)).

Finally, table A.4 shows estimates of supply-side parameters. Here, esti-

mates are slightly different than those from GMM estimation. In particular, the fuel cost parameter ρ is 12% smaller when estimated by 2SLS, though estimates are less precise in this case due to larger standard errors.

Table A.2: Estimates of demand parameters using 2SLS.

	$\ln(s_k/s_0)$
Price ($100) [$-\alpha$]	-0.800
	(0.110)
ln(share within nest) [1-λ]	0.366
	(0.050)
Departures per week	0.036
	(0.002)
Number of stops	-0.848
	(0.073)
Market distance (100 mi.)	0.077
	(0.015)
Market distance squared	0.0001
	(0.0003)
Connection extra distance (100 mi.)	-0.089
	(0.009)
Connection extra distance squared	0.004
	(0.001)
Share of delayed departures	-0.603
	(0.094)
Destinations from origin	0.011
	(0.002)
Fixed effects	
Airline	Yes
Origin airport-by-quarter	Yes
Destination airport-by-quarter	Yes
Observations	267,967

Note: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

A.2.3 Alternative specifications

Section 2.6.2 shows that the predictions using the estimated structural model reproduce key patterns in out-of-sample data. Nevertheless, it is useful to gauge whether alternative choices in the model would substantially affect the main parameters. Table A.5 considers plausible alternatives based on the literature.

As in Aguirregabiria and Ho (2012), specification (1) is estimated via OLS to illustrate endogeneity bias. When instruments are not used, the price coefficient becomes positive and the nesting coefficient is negative—both violate standard assumptions of demand models.

Specification (2) does not include high-dimensional fixed effects combining location and time. A comparison with the GMM estimates shows that α falls by a half, while other key parameters are less drastically affected. Interestingly, this smaller α is closer to the value of 0.45 found in Pagoni and Psaraki-Kalouptsidi (2016), which estimate a model without a rich set of location fixed effects.

In this main specification, prices are included at level because it is assumed that the representative consumer has constant marginal utility of income equal to α—a standard assumption used in nested logit estimation and in other papers in the aviation literature. Relaxing this assumption involves reformulating the welfare framework and requires data on individual income, which, to my knowledge, does not exist. Nevertheless, we can evaluate whether the choice of using prices at level instead of log-prices is influential to other key parameters in the model. Specification (3) in table A.5 shows that his is not the case: when including the natural log of prices, estimates of other parameters are minimally affected.

Specifications (4) and (5) in table A.5 consider alternative nesting choices. Recall that the main specification in the chapter includes all flights in a single nest, whereas the outside option is kept in a separate nest. Specification (4) nest flights by airline, as in Aguirregabiria and Ho (2012). The results show that this specification does not seem to adjust well to the data, as its estimated price coefficient is positive, violating model assumptions. Since the distinction between stop and nonstop flights is central in the present chapter, I also consider an alternative that nests these flight types separately in specification (5). However, as in the previous case, the results show that nesting by flight type violates model assumptions because λ is negative.

Table A.3: First stage regressions of the demand model using 2SLS.

	Dependent variable:		
	Price	ln(share within nest)	Departures per week
Number of stops	0.036	−1.179	−0.030
	(0.022)	(0.022)	(0.206)
Market distance (100 mi.)	0.068	0.016	−0.970
	(0.010)	(0.013)	(0.084)
Market distance squared	0.001	0.004	0.031
	(0.0002)	(0.0004)	(0.002)
Connection extra distance	−0.024	−0.152	−0.920
	(0.005)	(0.005)	(0.053)
Connection extra distance squared	0.001	0.010	0.076
	(0.0003)	(0.0004)	(0.004)
Share of delayed departures	−0.383	−0.416	−6.059
	(0.094)	(0.107)	(0.798)
Destinations from origin	0.018	−0.010	−0.086
	(0.001)	(0.001)	(0.006)
Airlines in market	−0.060	−0.129	−0.629
	(0.015)	(0.022)	(0.144)
Rivals' products in market	−0.0005	−0.002	0.008
	(0.0002)	(0.0004)	(0.002)
Rivals' % of nonstop flights	−0.602	−0.222	7.351
	(0.077)	(0.225)	(1.266)
Potential legacy entrants	0.074	−0.049	0.006
	(0.031)	(0.062)	(0.428)
Potential LCC entrants	−0.029	0.058	−0.857
	(0.012)	(0.017)	(0.113)
Fuel expenditure	0.003	−0.015	−0.066
	(0.001)	(0.001)	(0.010)
Compl. segment density	0.007	0.007	0.320
	(0.0004)	(0.0003)	(0.005)
Fixed effects			
Airline	Yes	Yes	Yes
Origin airport-by-quarter	Yes	Yes	Yes
Destination airport-by-quarter	Yes	Yes	Yes
Observations	267,967	267,967	267,967
F-statistic	317	957	353
Conditional F-statistic	23	62	62

Notes: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together). The conditional F-statistic is calculated following Sanderson and Windmeijer (2016).

Table A.4: Estimates of supply parameters using 2SLS.

	$\widehat{c}_k$
Fuel expenditure/avail. seat [ρ]	0.669
	(0.161)
Total ramp-to-ramp time (h)	0.156
	(0.017)
$\times$ *American*	-0.038
	(0.007)
$\times$ *Delta*	0.049
	(0.008)
$\times$ *United*	-0.014
	(0.008)
$\times$ *Alaska*	-0.020
	(0.033)
$\times$ *JetBlue*	-0.013
	(0.010)
$\times$ *Other low-cost*	-0.124
	(0.007)
Fixed effects	
Quarter	Yes
Each route airport-by-airline	Yes
Observations	267,967

Note: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

Table A.5: Estimates of demand parameters under alternative specifications.

	Dependent variable:				
	$\ln(s_k/s_0)$				
	(1) OLS	(2) 2SLS No FE	(3) 2SLS Log prices	(4) 2SLS Airline nests	(5) 2SLS Flight type nests
Price ($100) [$-\alpha$]	0.008	−0.411		0.101	−0.152
	(0.0004)	(0.085)		(0.082)	(0.025)
ln(Price)			−2.920		
			(0.453)		
ln(share within nest) [1-λ]	−0.040	0.443	0.381	0.411	1.075
	(0.002)	(0.021)	(0.041)	(0.034)	(0.017)
Departures per week	0.790	0.021	0.032	0.028	−0.001
	(0.008)	(0.001)	(0.002)	(0.002)	(0.001)
Number of stops	−0.346	−0.703	−0.803	−1.187	−0.310
	(0.014)	(0.031)	(0.061)	(0.018)	(0.016)
Market distance (100 mi.)	−0.007	0.037	0.067	0.008	0.003
	(0.009)	(0.013)	(0.014)	(0.009)	(0.003)
Market distance squared	−0.002	−0.001	−0.0003	−0.001	0.001
	(0.0004)	(0.0003)	(0.0003)	(0.0002)	(0.0001)
Conn. extra distance (100 mi.)	−0.027	−0.070	−0.093	−0.072	−0.009
	(0.002)	(0.005)	(0.007)	(0.007)	(0.002)
Conn. extra distance squared	0.001	0.003	0.004	0.001	0.001
	(0.0001)	(0.0003)	(0.0005)	(0.001)	(0.0001)
Share of delayed departures	−0.379	−0.799	−0.522	−0.682	0.176
	(0.054)	(0.079)	(0.088)	(0.075)	(0.027)
Destinations from origin	−0.001	−0.002	0.007	0.009	0.0001
	(0.0002)	(0.001)	(0.002)	(0.002)	(0.0005)
Fixed effects					
Airline	Yes	Yes	Yes	Yes	Yes
Origin airport-by-quarter	Yes	No	Yes	Yes	Yes
Destination airport-by-quarter	Yes	No	Yes	Yes	Yes
Nesting	All flights	All flights	All flights	By airline	Stop vs. nonstop
Observations	267,967	267,967	267,967	267,967	267,967

Note: standard errors, shown in parentheses, are clustered by non-directional city pairs (markets in opposite directions are clustered together).

A.3 Additional details on the data set

This appendix complements section 2.5 and presents additional information on definitions and variables that are calculated using multiple sources.

Airline groups. When assigning ownership of flights to specific airlines, I group firms based on the controlling airline. In doing so, regional carriers that are owned by or have exclusive service to a major airline are identified as being part of that major airline. This procedure is similar to Aguirregabiria and Ho (2012). Table A.6 shows how individual carriers are organized into airline groups. The last group contains five regional carriers that operate under the name of at least two major carriers, so their service contracts are not exclusive at the national level. In these cases, the major carrier is the ticketing carrier I observe in the data; I use this information to assign a flight to its respective major airline group.

Jet fuel use. There are four steps in calculating the average fuel use per available seat in a route. First, I use Table T-100 in the Form 41 Traffic Database (BTS, 2019b) to obtain the total departures performed and available seats by aircraft model, airline, and segment. This information allows me to calculate capacity-based use shares of each aircraft model.

Second, I calculate the total jet fuel burn for each aircraft model and stage length (segment distance) using the ICAO Fuel Consumption Table in Appendix C of the ICAO Carbon Emissions Calculator Methodology v.11 (ICAO, 2018). Matching these data requires a few conversions. Table T-100 identifies aircraft models with DOT codes instead of the standard IATA Designator Codes used

Table A.6: Airline groups and types in the dataset.

Carrier name (code)	Airline group	Type of airline
American Airlines Inc. (AA) PSA Airlines Inc. (16, OH) Envoy Air (MQ)	American	Legacy
Delta Air Lines Inc. (DL) Endeavor Air Inc. (9E)	Delta	Legacy
United Air Lines Inc. (UA) ExpressJet Airlines LLC (EV) Island Air Hawaii (WP) Air Wisconsin Airlines Corp (ZW)	United	Legacy
Alaska Airlines Inc. (AS) Horizon Air (QX)	Alaska	Mixed
Southwest Airlines Co. (SW)	Southwest	Low-cost
JetBlue Airways (B6)	JetBlue	Low-cost
Frontier Airlines Inc. (F9)	Frontier	Low-cost
Allegiant Air (G4)	Allegiant	Low-cost
Spirit Air Lines (NK)	Spirit	Low-cost
Sun Country Airlines (SY)	Sun Country	Low-cost
GoJet Airlines (G7) Compass Airlines (CP) SkyWest Airlines Inc. (OO) Mesa Airlines Inc. (YV) Republic Airline (YX)	American, Delta, or United (varies by market)	Regional partners

in ICAO's table; I manually match the 44 aircraft models in (the subset of) T-100 to IATA Designator Codes. Moreover, ICAO's table provides total fuel burn in kilograms for stage lengths in nautical miles (nm). I convert jet fuel mass to volume using the reference density of 820 kg/m^3 at 15 °C—approximately 0.217 kg/gallon at 60 °F—for Jet Fuel A, the standard used in US aviation (NREL, 2001). Also, one nautical mile corresponds to 1.15078 miles. Total fuel burn is tabulated at values of 125, 250, 500, 750, 1000, and subsequent increments of 500 nm up to 8500 nm, depending on the range of the model. I use linear interpolation to calculate fuel burn between tabulated stage lengths.

Figure A.1 illustrates the importance of factoring stage lengths instead of relying on average fuel efficiency by aircraft model. This figures shows how

Figure A.1: Fuel efficiency by stage length for the aircraft models most used.

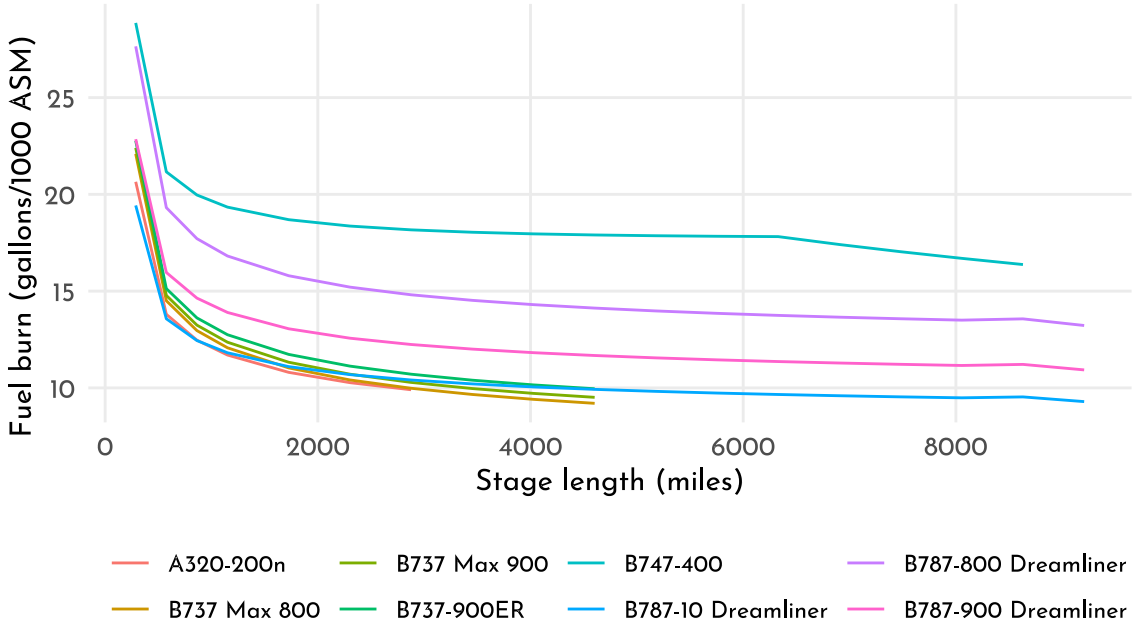

Notes: Fuel efficiency is measured in gallons of jet fuel per 1000 available seat-miles. The eight models shown in this graph correspond to those that offered the highest aggregate capacity (in available seat-miles) in 2018.

fuel burn per available seat-mile decreases with increased distance flown for the eight most used models.[2] The "hockey stick" shape of efficiency curves represent the fact that taking off and climbing consume a substantial fraction of fuel allocated to a flight. Since the cruising stage is more fuel efficient, longer flights burn less fuel per mile on average. ICAO's fuel burn data are representative of the average conditions in which these aircraft models are used in industry. Actual fuel consumption in a given trip—an information not publicly available—varies by weight load, flight plan, weather conditions, and other factors. Nevertheless, passenger occupation rate, or load factor, has a small effect on fuel burn, as passengers and baggage typically account for only 15% of the

[2]The number of available seats of an aircraft model varies by plane, with some airlines using more dense configurations. The calculation in this graph uses the average number of seats across all planes.

take-off weight (Borenstein & Rose, 2014).

Third, I calculate average fuel use for each segment, airline, and quarter. The fuel use of an aircraft model in a segment is the total fuel burn, calculated in the previous step, divided by the average number of seats available for that model and airline. I allow the number of seats to vary by airline in order to accommodate different seat configurations. Then, for each segment and airline in a quarter, I calculate average fuel efficiency across aircraft models used (if the airline used more than one), weighting each model by the capacity-based use share in that quarter.

Fourth and finally, I calculate average fuel use per available seat in a product (i.e., by route, airline, and quarter) by summing average the fuel use of each segment in a route for that airline and quarter.

Jet fuel expenditure. Jet fuel prices are gathered from the US Energy Information Administration (EIA, 2018). I use average end user sales prices by quarter and region. Regions are based on the Petroleum Administration for Defense Districts (PADD). The data cover all five districts (and sub-districts): West Coast, Rocky Mountain, Gulf Coast, Midwest, and East Coast. The East Coast PADD has three sub-districts: New England, Central Atlantic, and Lower Atlantic. End user sales prices include delivery costs, which captures relevant variations in fuel prices across the US. Despite this spatial variation, end user sales prices track closely changes in spot prices, as illustrated in Figure A.2, panel (a). Panel (b) in this Figure shows the variation of end user sales prices across regions and compares it with spot prices.

To calculate jet fuel expenditure per available seat, I multiply average jet fuel

Figure A.2: Jet fuel and crude oil prices.

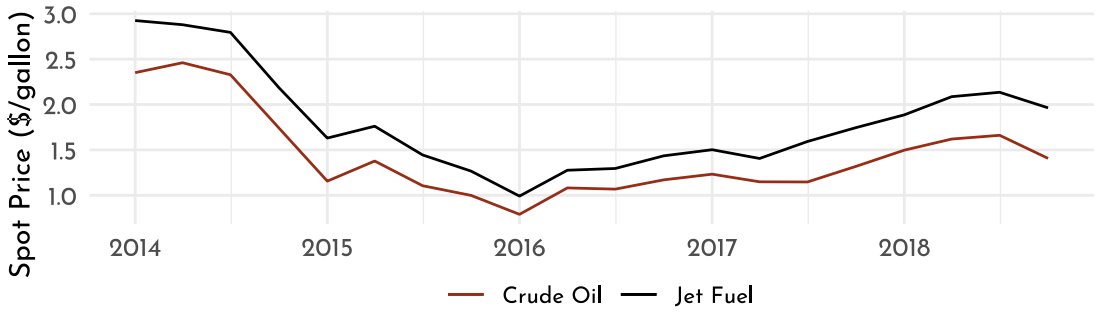

(a) Jet fuel and crude oil spot prices.

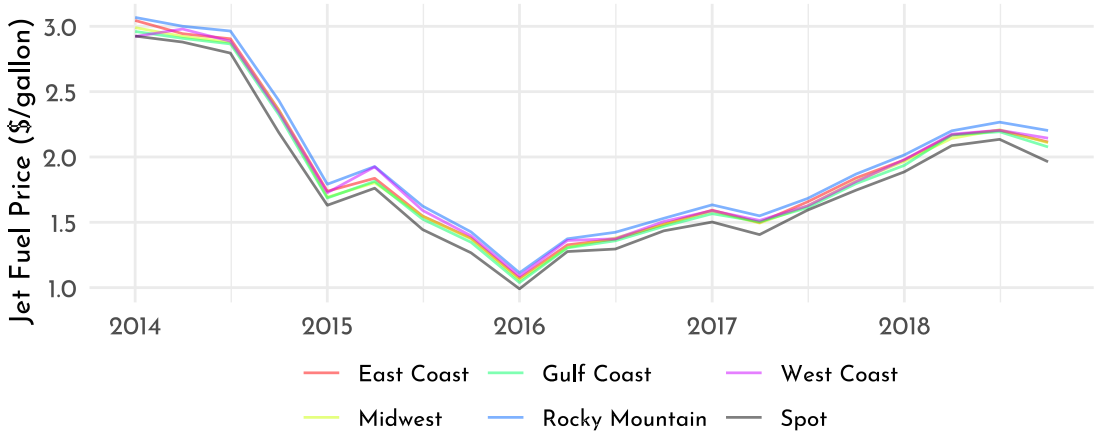

(b) Jet fuel spot and regional end user sales prices.

Notes: Spot prices are for US Gulf Coast Kerosene-Type Jet Fuel and Crude Oil WTI. Regional (PADD) prices are average end user sales prices by quarter. *Source: US Energy Information Administration*.

use in each segment by its respective jet fuel price. The price assigned to each segment corresponds to the PADD where the departing airport is located. Then, for each product, I sum fuel expenditures of all the segments of a product's route.

A.4 Market-specific marginal abatement costs

This appendix complements the discussion about the distribution of marginal abatement costs (MAC) across markets. Figure A.3 shows a histogram of MACs for each market size quartile, with the first quartile containing the smallest markets. This figure indicates that the dispersion of market MACs increases with market size: the proportion of markets with MAC below \$200/ton CO_2 is higher in the first and second quartiles.

Figure A.3: Histograms of baseline market-specific marginal abatement costs (MAC) for each market size quartile.

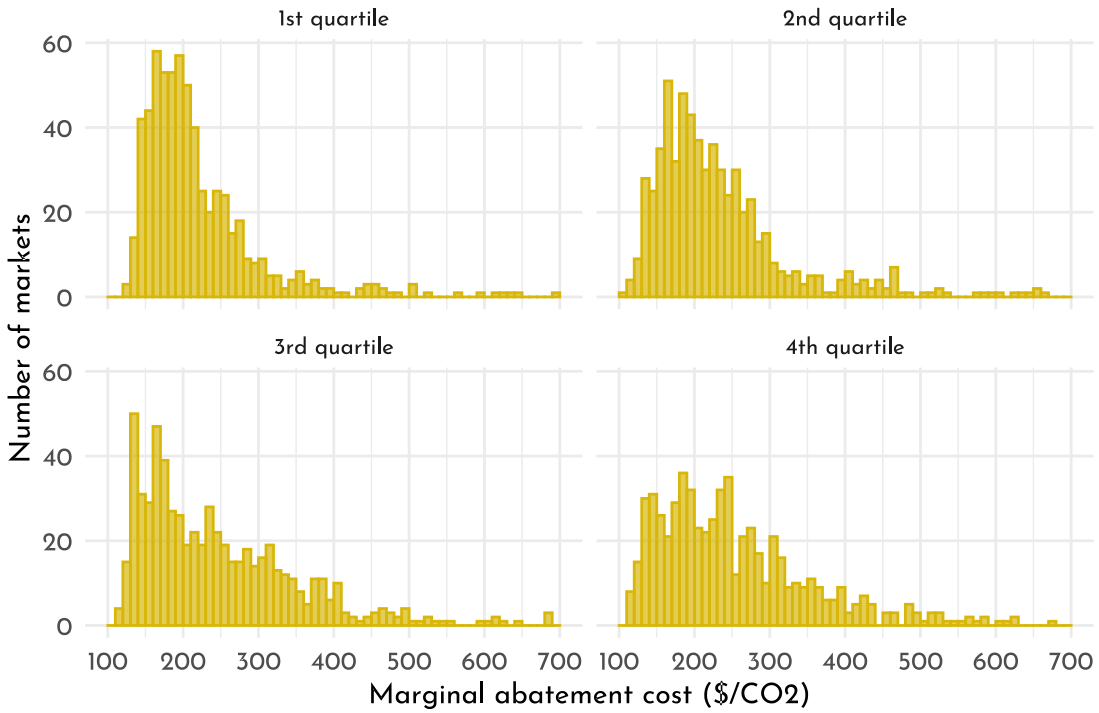

Notes: marginal abatement costs are defined in section 2.3.4. Smallest market sizes are in the first quartile.

Though the first quartile has a higher count of markets with low MAC, these markets generate a small fraction of the emissions in the sector. Since carbon damages are proportional to total emissions, it is also important to evaluate the

distribution of MACs over emissions. To do so, I weight each market by its emissions. Figure A.4 compares the unweighted and emission-weighted distributions of each market size quartile. Panel (a) in this figure reproduces the patterns observed in figure A.3, with the lowest quartiles having a larger proportion of markets with MACs below $200. However, panel (b) in A.4 shows that this order is inverted in the lower range when we consider emission totals: about 18% of the emissions from the fourth quartile are below $150, whereas 9% of emission in the first quartile are in this range. This reflects the fact that more competition in markets serving large urban areas tends to drive markups and prices down, thus decreasing the distortions that contribute to a higher MAC. Since the fourth quartile corresponds to three-quarters of total sector emissions, the vast majority of low-MAC emissions come from large markets: 82% of low-MAC emissions are from the fourth quartile, comparted to less than 5% from the two lowest quartiles combined.

Figure A.4: Cumulative distributions of baseline market-specific marginal abatement costs (MAC) for each market size quartile.

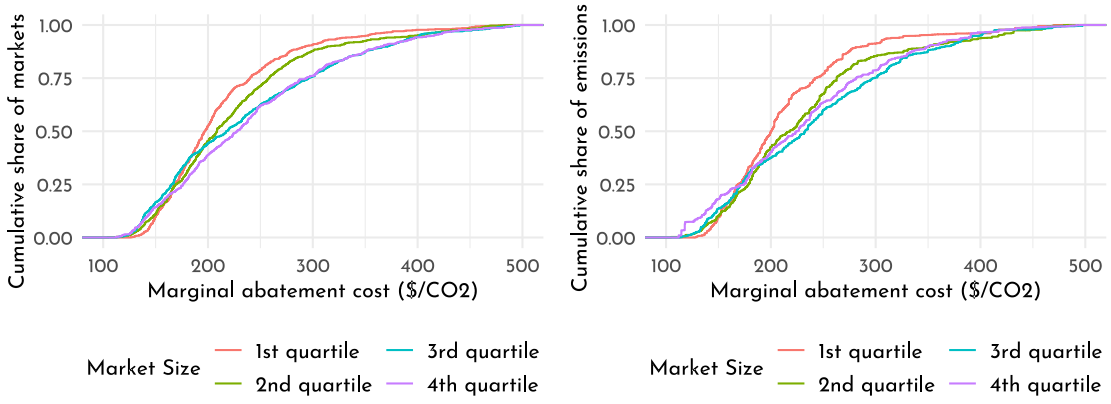

(a) Unweighted distribution

(b) Emission-weighted distribution

Notes: marginal abatement costs are defined in section 2.3.4. Smallest market sizes are in the first quartile. In panel (a), *unweighted* shares are relative to the total number of markets. In panel (b), *emission-weighted* distributions weight each market by its total emissions. These can be interpreted as the distributions of MACs with respect to total emissions, so that each point along the line indicates the share of emissions with a MAC at or below that level.

APPENDIX B

APPENDIX FOR CHAPTER 3

B.1 The feedback-implied distributions

B.1.1 Derivation of the used distributions

Truncated Roe-Baker (TRB) distribution. Physical constraints require the feedback parameter f to be smaller than 1. The Roe-Baker distribution for the climate sensitivity is obtained from assuming that f follows a normal distribution. In the literature, this assumption has led to positive probability on negative climate sensitivities. A first approach to bound the climate sensitivity to a meaningful support is to truncate the normal distribution for f between 0 and 1. Although some climate processes generate a negative feedback, the sum of all feedback processes is positive. Hence, a truncation on the lower end is a valid choice and it has the advantage of resulting in a support for s that is bounded below in s_0 and comparable with the RIS-Beta distribution.

Roe and Baker (2007) derive the PDF for a distribution of s based on a normal distribution for f with mean $\bar{f}$ and standard deviation σ_f. The transformation applied to the normal is monotonic between 0 and 1, so this truncation leads to a density reweighing by $\Pr(0 \leq f < 1) = \Phi\left(1; \bar{f}, \sigma_f\right) - \Phi\left(0; \bar{f}, \sigma_f\right)$. Hence, for $z \in [s_0, \infty)$, the truncated Roe-Baker PDFs is:

$$\theta^{\text{TRB}}(z) = \frac{1}{\Phi\left(1; \bar{f}, \sigma_f\right) - \Phi\left(0; \bar{f}, \sigma_f\right)} \left\{ \frac{1}{\sigma_f \sqrt{2\pi}} \frac{s_0}{z^2} \exp\left[-\frac{1}{2} \left(\frac{1 - \bar{f} - \frac{s_0}{z}}{\sigma_f} \right)^2 \right] \right\} \tag{B.1}$$

where the first term on the right-hand side accounts for the conditional proba-

bility and the second term is the distribution proposed in Roe and Baker (2007).

Reflected-Inverted-Scaled-Beta (RIS-Beta) distribution. The truncation procedure above leaves positive density on the upper bound, which translates into a fixed positive density for $s \to \infty$. Another possible approach is to assume a distribution for f that is built on the $[0, 1]$ support. Due to the reciprocal nature of the feedback parameter, it is reasonable to assume that the density of f tends to zero as f approaches 1; with the peak density above 0.5, this assumption implies some skewness must exist. The beta distribution is flexible enough to accommodate these requirements and is used to derive an alternative distribution for the climate sensitivity: the Reflected-Inverted-Scaled-Beta (RIS-Beta).

Let $f \sim Beta(\alpha_f, \beta_f)$. Then, by property of the beta distribution, the reflected term $(1 - f) \equiv j \sim Beta(\beta_f, \alpha_f)$. The PDF of j is given by:

$$\phi(j) = \frac{1}{B(\alpha_f, \beta_f)} j^{\beta_f - 1} (1 - j)^{\alpha_f - 1} \tag{B.2}$$

where B is the beta function, defined as $B(\alpha, \beta) = \int_0^1 t^{\alpha-1} (1 - t)^{\beta-1} dt$. Note that the beta function is symmetric.

Following the temperature transition equation, the climate sensitivity is given by $z = s_0 j^{-1}$, where s_0 is the reference climate sensitivity (i.e. without feedback). Inverting the expression, define $\omega(z) \equiv s_0 z^{-1} = j$. Letting $\theta(z)$ denote the PDF, for any interval $[z_1, z_2]$ it follows

$$\Pr(z_1 \leq z \leq z_2) = \int_{z_1}^{z_2} \theta(z) \, dz \tag{B.3}$$

Because ω is a monotonic transformation, it is possible to map each interval defined with respect to z to an interval for j. Hence, rewrite (B.3) as

$$\Pr(z_1 \leq z \leq z_2) = \Pr(\omega(z_1) \leq j \leq \omega(z_2)) = \int_{\omega(z_1)}^{\omega(z_2)} \phi(j) \, dj$$

Performing a change of variable in the previous expression, it follows

$$\Pr\left(z_1 \leq z \leq z_2\right) = \int_{z_1}^{z_2} \phi\left(\omega\left(z\right)\right) \left|\frac{d\omega\left(z\right)}{dz}\right| dz \tag{B.4}$$

where it was applied $dj = \left|\frac{d\omega(z)}{dz}\right| dz$. Comparing (B.4) and (B.3), it follows that

$$\theta\left(z\right) = \phi\left(\omega\left(z\right)\right) \left|\frac{d\omega\left(z\right)}{dz}\right|$$

Plugging in (B.2) yields

$$\theta^{\text{RIS-Beta}}\left(z\right) = \frac{s_0^{\beta_f}}{B\left(\alpha_f, \beta_f\right)} \frac{\left(z - s_0\right)^{\alpha_f - 1}}{z^{\alpha_f + \beta_f}} \tag{B.5}$$

which defines the PDF of the RIS-Beta distribution with parameters (α_f, β_f, s_0). In the calibrations and simulations, the CDF for the RIS-Beta distribution was obtained numerically; no closed-form could be derived analytically.

Reflected-Inverted-Scaled-Uniform (RIS-Uniform) distribution. If we consider that the current knowledge of post-tipping regimes is mostly speculative, the normal and beta distributions assumed in the truncated Roe-Baker and RIS-Beta might impose too much structure for priors. An alternative choice is to think of prior beliefs in f as following a uniform distribution, thus assuming that all scenarios between zero-feedback and runaway warming are equally probable.

Let f be uniformly distributed between $\left[\underline{f}, 1\right]$. Then, the PDF for $j = 1 - f$ is simply

$$\phi\left(j\right) = \frac{1}{1 - \underline{f}}$$

The transformation $\omega(z) \equiv s_0 z^{-1} = j$ is the same as in the RIS-Beta case. Hence, plugging in $\left|\frac{d\omega(z)}{dz}\right|$ yields

$$\theta^{RIS-Uniform}\left(z\right) = \frac{s_0}{1 - \underline{f}} \frac{1}{z^2} \tag{B.6}$$

The resulting distribution is effectively a Pareto with scale parameter $\frac{s_0}{1-\underline{f}}$ and shape parameter 1. This distribution, thus, illustrates an uncertainty case with very fat tails.

B.1.2 Distribution calibration condition

Claim: Let S be the climate sensitivity in a certainty scenario. The choice of distribution parameters that sets $\mathbb{E}\left[\frac{1}{\tilde{s}}\right] = \frac{1}{S}$ for the distribution of $\tilde{s}$ is sufficient for preserving the mean of the temperature transition, i.e. $\mathbb{E}_{\epsilon,s}\left[\tilde{T}_{t+1}\right] = \mathbb{E}_{\epsilon}\left[\tilde{T}_{t+1}|s=S\right]$.

Proof: From equation (B.14), define the next-period temperature before the temperature shock as a function of s

$$\Theta(s) = T_t + C_T\left[F(M_{t+1}, t+1) - \frac{f}{s}T_t - [1-\alpha_T(T_t,t)]\,C_O T_t\right]$$

Hence, $\tilde{T}_{t+1} = \tilde{\epsilon}\Theta(s)$. Abusing notation, let $\Theta(\tilde{s})$ be the random variable representing next-period temperature before the shock but with uncertain climate sensitivity. Then, it follows

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{\tilde{s}}\right] &= \frac{1}{S} \\
\mathbb{E}_s\left[\Theta(\tilde{s})\right] &= \Theta(S) \\
\mathbb{E}_\epsilon\left[\tilde{\epsilon}\mathbb{E}_s\left[\Theta(\tilde{s})\right]\right] &= \mathbb{E}_\epsilon\left[\tilde{\epsilon}\Theta(S)\right] \\
\mathbb{E}_\epsilon\left[\mathbb{E}_s\left[\tilde{T}_{t+1}|\epsilon\right]\right] &= \mathbb{E}_\epsilon\left[\tilde{T}_{t+1}(S)\,|\epsilon, s=S\right]
\end{aligned}$$

where the second line acknowledges the fact that $\Theta(s)$ can be obtained after a series of linear operations over $\frac{1}{s}$. Then, by the Law of Iterated Expectations

$$\mathbb{E}_{\epsilon,s}\left[\tilde{T}_{t+1}\right] = \mathbb{E}_\epsilon\left[\tilde{T}_{t+1}|s=S\right]$$

■

B.2 A DICE model with uncertain post-threshold regime

The model in this paper extends Lemoine and Traeger (2014) to allow post-threshold uncertainty. It includes 4 of the 6 original states in DICE-2007 (Traeger, 2014). This is done by replacing ocean-related states—upper and lower carbon stock and temperature—with calibrated, time-varying parameters in the atmospheric dynamics. In doing so, the transitions for carbon stock and surface temperature approximate the trajectories in the full model.

B.2.1 Baseline model

Production and investment. Let K_t, L_t, and A_t be capital, labor, and labor augmenting technology at time t. Define controls in terms of effective labor: $k_t \equiv \frac{K_t}{A_t L_t}$ and $c_t \equiv \frac{C_t}{A_t L_t}$. Gross output, also in effective terms, is then given by

$$y_t^g = \frac{(A_t L_t)^{1-\kappa} K_t^{\kappa}}{A_t L_t} = k_t^{\kappa} \tag{B.7}$$

where: κ is the elasticity of capital in a Cobb-Douglas production function.

Labor grows exogenously, asymptotically approaching L_∞ as

$$L_t = L_0 + (L_\infty - L_0)\left(1 - e^{-t\delta_L}\right) \tag{B.8}$$

where: δ_L determines the speed of convergence for L_t. Hence, the exogenous growth rate for labor (approximated in continuous time) is

$$g_{L,t} = \frac{\frac{\partial L_t}{\partial t}}{L_t} = \frac{\delta_L (L_\infty - L_0) e^{-t\delta_L}}{L_\infty - (L_\infty - L_0) e^{-t\delta_L}} = \delta_L \left[\frac{L_\infty}{L_\infty - L_0} e^{t\delta_L} - 1\right]^{-1} \tag{B.9}$$

Likewise, technological progress is assumed to follow an exponentially declining growth rate

$$g_{A,t} = g_{A,0} \exp\left[-t\delta_A\right] \tag{B.10}$$

which implies the continuous time solution

$$A_t = A_0 \exp\left[\frac{g_{A,0}}{\delta_A}\left(1 - e^{-t\delta_A}\right)\right] \tag{B.11}$$

The impact of climate change on gross output is a function of the total temperature change (T_t), resulting in the post-damage output

$$y_t^d = \frac{y_t^g}{1 + D_t} = \frac{y_t^g}{1 + d_1 T_t^{d_2}} \tag{B.12}$$

where: d_1 and d_2 are calibrated parameters.

Capital investment obeys the following law of motion (in effective terms)

$$k_{t+1} = e^{-\left(g_{L,t} + g_{A,t}\right)}\left[\left(1 - \delta_k\right) k_t + \left(1 - \Lambda_t\right) y_t^d - c_t\right] \tag{B.13}$$

where: δ_k is the capital depreciation rate; and Λ_t is the abatement cost as a share of post-damage output.

Climate. Next period temperature is a random variable $\tilde{T}_{t+1}$. In the pre-threshold regime, the only source of uncertainty is the random multiplicative shock $\tilde{\epsilon}_t$ with mean 1. The post-threshold transition for temperature also has uncertainty over the climate sensitivity $\tilde{s}$.

$$\tilde{T}_{t+1} = \tilde{\epsilon}_t \left\{ T_t + C_T \left[F\left(M_{t+1}, t+1\right) - \frac{F_{2\times C}}{\tilde{s}} T_t - \left[1 - \alpha_T\left(T_t, t\right)\right] C_O T_t \right] \right\} \tag{B.14}$$

where: C_T translates forcing into temperature change; C_O translates surface-ocean temperature into forcing; $F\left(M_{t+1}, t+1\right)$ is the next-period forcing; $F_{2\times C}$ is the forcing from doubling CO_2 (in W/m^2); α_T is the ocean cooling parameter calibrated based on the full model with 2 temperature states.

Forcing is defined as

$$F\left(M_t, t\right) = f \ln\left(\frac{M_t}{M_{pre}}\right) + EF_t \tag{B.15}$$

where: M_t and M_pre are the atmospheric CO_2 stock at period t and at preindustrial level; EF_t is the exogenous (non-CO_2) component of forcing, increasing linearly in the first 100 years as

$$EF_t = EF_0 + 0.01\left(EF_{100} - EF_0\right)\min\left\{t, 100\right\} \tag{B.16}$$

Carbon cycle. Atmospheric carbon stock follows

$$M_{t+1} = e_t + M_t\left[b_{11} + b_{21}\left[b_{12} + \left(b_{22} + b_{32}b_{23}\right)\alpha_B\left(M_t, t\right) + b_{32}b_{33}\alpha_O\left(M_t, t\right)\right]\right] \tag{B.17}$$

where: b_{ij} are transition parameters of the carbon cycle in the model with 3 reservoirs; $\alpha_B\left(M_t, t\right)$ and $\alpha_O\left(M_t, t\right)$ are the shallow and deep ocean stock of carbon measured as a fraction of the atmospheric stock; e_t is time t emissions defined as

$$e_t = \sigma_t\left(1 - \mu_t\right)Y_t + B_t \tag{B.18}$$

where: μ_t is the abatement rate; σ_t is the emission intensity parameter; B_t is the exogenous CO_2 emissions, unrelated to industrial production. The other two parameters follow exogenous paths

$$\sigma_t = \sigma_0 \exp\left[\frac{g_{\sigma,0}}{\delta_\sigma}\left(1 - e^{-t\delta_\sigma}\right)\right] \tag{B.19}$$

$$B_t = B_0 e^{tg_B} \tag{B.20}$$

where: σ_0 and B_0 are the initial values of the parameters; g_B is the annual growth rate of B; $g_{\sigma,0}$ is the initial growth rate of σ; and δ_σ is the annual change in the growth rate of σ.

Abatement. Industrial emissions can be abated at cost

$$\Lambda_t = \Psi_t \mu_t^{a_2} \tag{B.21}$$

with

$$\Psi_t = \frac{a_0 \sigma_t}{a_2} \left(1 - \frac{1 - e^{t g_\Psi}}{a_1} \right) \tag{B.22}$$

where: a_0 is the cost of a backstop technology; a_1 is the ratio of initial to final backstop cost; a_2 defines the cost function curvature; and g_Ψ is the annual growth rate of backstop cost.

Tipping threshold. The climate regime changes when temperature exceeds an unknown threshold $\tilde{T}$. Prior beliefs on the location of the threshold follow a uniform distribution with upper bound $\bar{T}$. The probability of crossing the threshold in the next period is given by the hazard function

$$h\left(T_t, T_{t+1}\right) = \max \left\{ 0, \frac{\min\left\{T_{t+1}, \bar{T}\right\} - T_t}{\bar{T} - T_t} \right\} \tag{B.23}$$

Utility. Agents derive utility following a CRRA with intertemporal elasticity of substitution equals to $\frac{1}{\eta}$:

$$u\left(C_t, L_t\right) = \frac{\left(C_t / L_t\right)^{1-\eta}}{1-\eta} \tag{B.24}$$

Hence, total welfare weighted by population in period t can be written as:

$$W\left(c_t, L_t, A_t\right) = L_t \frac{\left(A_t c_t\right)^{1-\eta}}{1-\eta} = L_t \frac{c_t^{1-\eta}}{1-\eta} A_t^{1-\eta} \tag{B.25}$$

Problem. The planner's problem can be expressed in a dynamic recursive fashion. Let V_0 and V_1 be the pre- and post-threshold value functions. Also, let $\mathbf{S}_t = (k_t, M_t, T_t, t)$ be the set of states, and $\mathbb{P}_\epsilon$ be the probability measure of the temperature uncertainty. Then, the Bellman equation can be written as

$$V_0\left(\mathbf{S}_t\right) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \beta_t \int_\varepsilon \left\{ \left[1 - h\left(T_t, T_{t+1}\right)\right] V_0\left(\mathbf{S}_{t+1}\right) + h\left(T_t, T_{t+1}\right) V_1\left(\mathbf{S}_{t+1}\right) \right\} d\mathbb{P}_\epsilon \tag{B.26}$$

subject to (B.13), (B.14), (B.17), the output constraint

$$c_t + \Lambda_t y_t^d \leq y_t^d, \tag{B.27}$$

and control constraints

$$c_t \geq 0 \tag{B.28}$$

$$\mu_t \in [0, 1] \tag{B.29}$$

The effective discount factor, β_t, varies with time; this is a consequence of time-dependent growth rates which must be accounted for when representing the Bellman equation in effective terms. Letting ρ represent the social discount factor, β_t follows[1]

$$\beta_t = \exp\left(-\rho + (1-\eta)\, g_{A,t} + g_{L,t}\right) \tag{B.30}$$

In analogy to the problem defined in (3.6), the Bellman equation for the post-tipping problem is given by

$$V_1\left(\mathbf{S}_t\right) = \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} + \beta_t \int_{\Sigma} \left[\int_{\varepsilon} V_1\left(\mathbf{S}_{t+1}\right) d\mathbb{P}_{\epsilon} \right] d\mathbb{P}_s \tag{B.31}$$

subject to the same set of constraints as in the pre-tipping problem.

Alternative utility function. A drawback from using the CRRA form is that risk aversion and intertemporal elasticity of substitution parameters are necessarily equal. We can disentangle these effects by implementing the Epstein-Zin-Weil preference. Keeping η as the intertemporal substitution parameter, and adding RRA as the relative risk aversion, rewrite the Bellman equation as

$$\begin{aligned} V_0\left(\mathbf{S}_t\right) = & \max_{c_t, \mu_t} \frac{c_t^{1-\eta}}{1-\eta} \\ & + \frac{\beta_t}{1-\eta} \left\{ \int_{\varepsilon} \left\{ (1-\eta) \left[\left[1 - h\left(T_t, T_{t+1}\right)\right] V_0\left(\mathbf{S}_{t+1}\right) + h\left(T_t, T_{t+1}\right) V_1\left(\mathbf{S}_{t+1}\right) \right] \right\}^{\frac{1-RRA}{1-\eta}} d\mathbb{P}_{\epsilon} \right\}^{\frac{1-\eta}{1-RRA}} \end{aligned} \tag{B.32}$$

[1] The proof can be found in the appendix of Traeger (2014).

The changes in the post-tipping problem are analogous.

Tables B.1 to B.4 display the values of the parameters used in the numerical analysis.

B.2.2 Solving the model

The dynamic stochastic program defined in B.2.1 is solved by value function iteration with modified policy iteration as proposed by Puterman and Shin (1978); each iteration in the value function is followed by one of more policy function iterations. The value functions are approximated using Chebyshev polynomials as basis functions and 7 Chebyshev nodes in each of the four states; the resulting state space is a tensor product with 7^4 nodes. The first run of each scenario starts with 5 nodes per dimension, and subsequent runs increment the number of nodes until 7, with the previous run used as initial guess; this is done to improve the quality of the first guess and speed up the convergence of the solution in the final state space. The tolerance for convergence is 10^{-3} for the sum of absolute changes in the value function basis coefficients.

For each scenario, the post-tipping program was first solved to obtain an approximation of V_1. Then, the basis coefficients of V_1 and solution policy functions were fed into the pre-tipping problem to obtain V_0. Finally, the code simulates trajectories starting from the initial state using the model transitions and the solution policy functions.

The distributions for the temperature shock and the climate sensitivity were generated using Gauss-Legendre quadrature with 8 nodes each. For the post-

Parameter	Value	Description
A_0	0.027	Initial level of technology
$g_{A,0}$	0.009	Initial growth rate of technology
δ_A	0.001	Rate of decline in technology growth
L_0	6514	Initial population [millions]
L_∞	8600	Asymptotic population
δ_L	0.035	Rate of population convergence
σ_0	0.13	Initial carbon intensity [Gt C / unit of output]
$g_{\sigma,0}$	-0.0073	Initial growth rate of carbon intensity
δ_σ	0.003	Rate of change in growth rate of carbon intensity
a_0	1.17	Initial cost of backstop technology [$1,000 / tC]
a_1	2	Ration of initial to final backstop technology cost
a_2	2.8	Abatement cost exponent
g_Ψ	-0.005	Growth rate of backstop cost
k_0	$137/(A_0 L_0)$	Initial effective capital [US $trillions]
ρ	0.015	Pure rate of time preference
$\eta = RRA$	2	Intertemporal elasticity of substitution = relative risk aversion

Table B.1: Parameters of the baseline economic model.

Parameter	Value	Description
B_0	1.1	Initial nonindustrial CO_2 emissions [Gt C]
g_B	-0.01	Growth rate of nonindustrial emissions
EF_0	-0.06	Initial exogenous forcing [W/m^2]
EF_100	0.3	Exogenous forcing in period 100 [W/m^2]
κ	0.3	Elasticity of capital in production
δ_κ	0.06	Capital depreciation rate
d_1	0.0028	Damage function linear coefficient
d_2	2	Damage function exponent
$F_{2\times C}$	3.8	Forcing from doubling CO_2 [W/m^2]
M_{pre}	5.96.4	Preindustrial (1900) atmospheric CO_2 [Gt C]
C_T	0.03	Translation of forcing into temperature change
C_O	0.3	Translation of surface-ocean temperature gradient into forcing
M_0	808.9	Initial atmospheric CO_2 [Gt C]
T_0	0.73	Initial warming relative to 1900 [Gt C]
s_{pre}	3	True pre-tipping climate sensitivity
s_{post}	6	True post-tipping climate sensitivity
b	$\begin{bmatrix} 0.978 & 0.023 & 0 \\ 0.011 & 0.983 & 0.005 \\ 0 & 0.0003 & 0.9997 \end{bmatrix}$	Carbon transfer matrix

Table B.2: Parameters of the baseline climate model.

tipping regime, the uncertainties were combined in a joint probability space with 64 nodes.

The code used for this paper is available under request. I extended the original Matlab code for Lemoine and Traeger (2014), adding support for the parallel processing toolbox and newer features of Matlab version 2017a. The code uses the CompEcon package by Miranda and Fackler (2004). The quadrature rules for the truncated normal use a function written by John Burkardt. Results for this paper were executed in a machine running Linux, but an option for Windows platform is available in the code.

B.3 Additional results

This appendix presents the trajectories of four variables not analyzed in the main text: emissions, damages, abatement rate, and net output. As for the other variables, changes from adding uncertainty are tiny and cannot be visually distinguished in the graphs. The Figures below illustrate the differences caused by changing certain aspects of the model. Note that while the shape of the trajectories is primarily determined by the underlying assumption in DICE, their scale is determined by the alternative specifications. For instance, the EZW case sets higher intertemporal elasticity of substitution and risk aversion, which translates in a more aggressive abatement with lower damages.

Figure B.1 shows the additional results for the baseline model.

Next, Figure B.2 shows the additional results for the model with Epstein-Zin-Weil preferences. Higher intertemporal elasticity of substitution and risk aver-

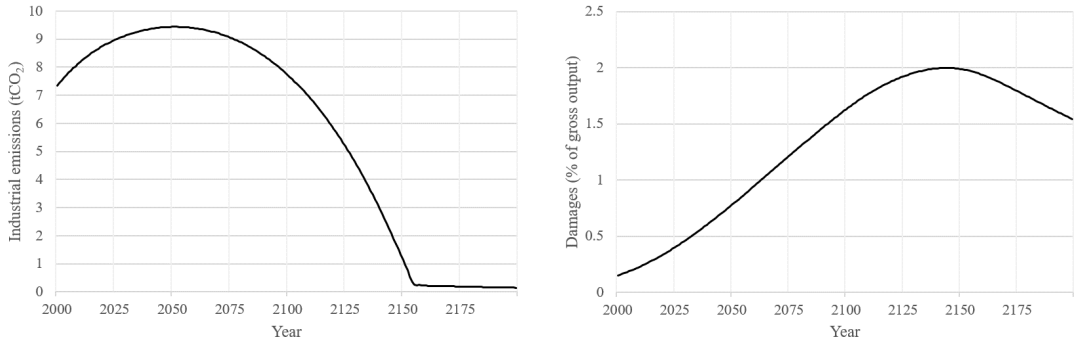

(a) CO_2 emissions

(b) Damages

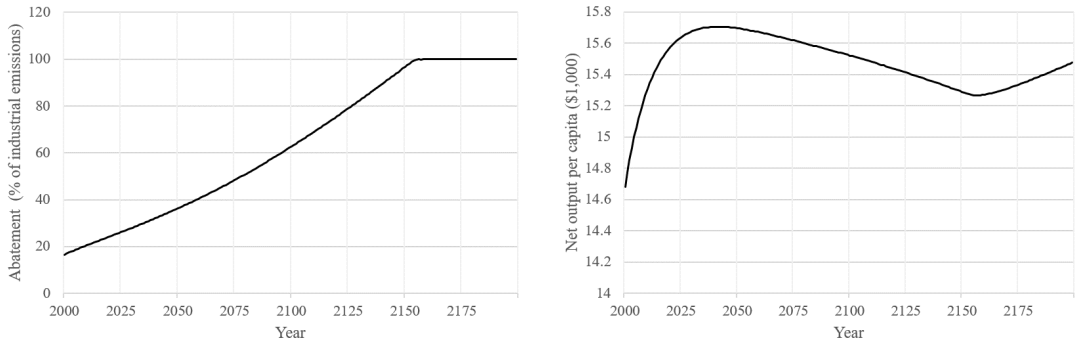

(c) Abatement rate

(d) Net output per capita

Figure B.1: Additional results for the baseline model (with CRRA utility function).

sion lead to a faster total decarbonization of the economy; as a consequences, peak damages are half of the baseline case and net output is higher.

Lastly, Figure B.3 displays the additional results for the model with the damage function proposed by Howard and Sterner (2017). With damages being almost four times higher, the abatement policy is even more aggressive than in the EZW case. The additional abatement effort limits peak damage to a value that is only 50% higher than the baseline case, at the cost of more intense loss of net output.

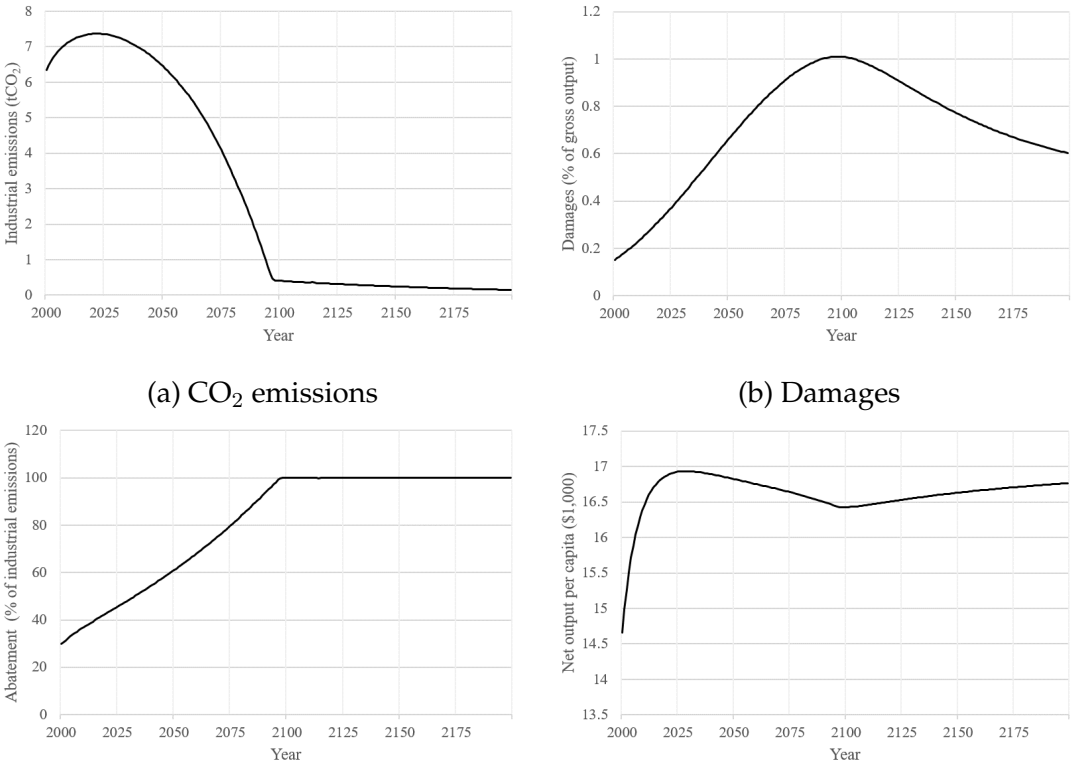

(a) CO_2 emissions

(b) Damages

(c) Abatement rate

(d) Net output per capita

Figure B.2: Additional results for a model with EZW utility function.

Parameter	Value	Description
$\bar{T}$	4.27	Upper bound of threshold beliefs
s_0	1.2	Reference climate sensitivity (with zero feedback)
$\bar{f}$	0.8	Mean feedback in TRB
σf	0.09	Standard deviation of feedbacks in TRB
α_f	11.74	First parameter for RIS-Beta
β_f	2.93	Second parameter for RIS-Beta
$\underline{f}$	0.6	Lower bound for feedback in RISU

Table B.3: Parameters for the climate sensitivity uncertainty cases.

Parameter	Value
Epstein-Zin-Weil	
η	0.6667
RRA	10
Higher damages	
d_1	0.01145

Table B.4: Parameters for alternative specifications.

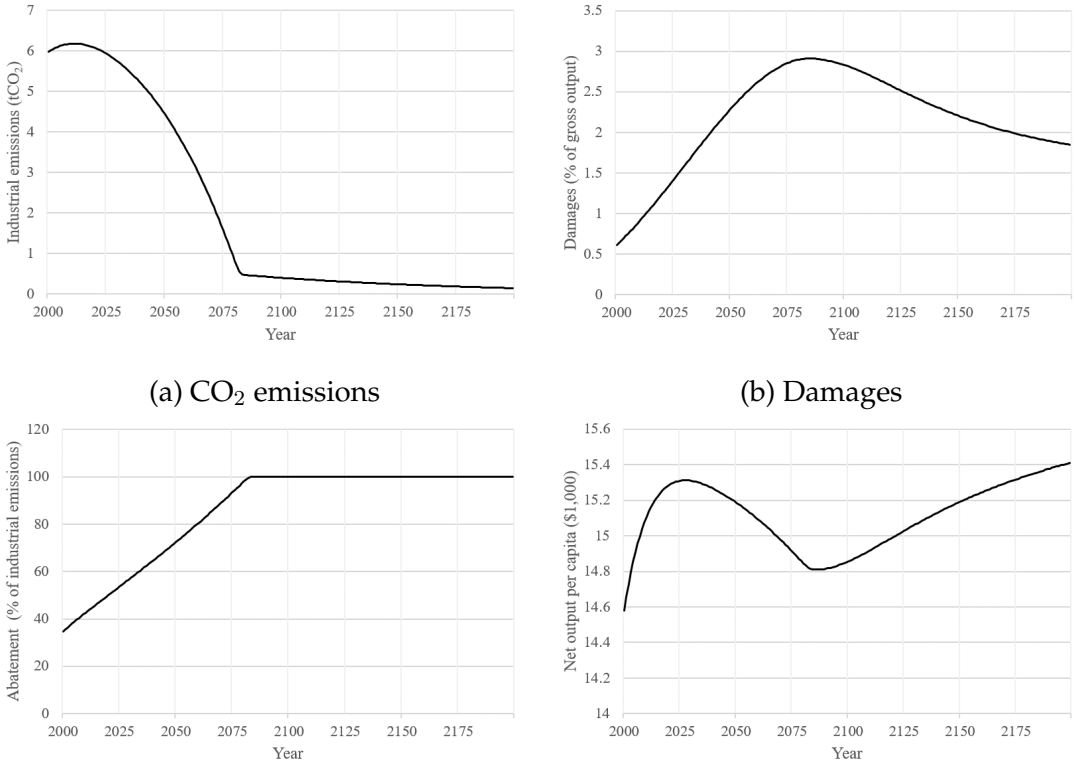

(a) CO_2 emissions

(b) Damages

(c) Abatement rate

(d) Net output per capita

Figure B.3: Additional results for a model with Howard and Sterner (2017) damage function.

APPENDIX C

APPENDIX FOR CHAPTER 4

This Appendix is structured as follows. Section C.1 describes the sources of data and outlines the steps in assembling our data set. Section C.2 formalizes the regression analysis and policy simulations. Section C.3 presents additional results for socioeconomic and demographic factors at select US region. Section C.4 examines how sensitive our results are to key parameters in the model. Section C.5 performs robustness checks by considering alternative specifications for our estimations. Finally, C.6 compares our approach to alternative metrics in the literature.

C.1 Data

Our data set includes 92,445 census block groups from 521 counties across 42 states. The data set covers approximately 145 million people—around 45% of the total US population as of 2016. This data set combines local water and sewer rates, number of service accounts, average consumption, climate characteristics, and a multitude of socio-economic-demographic indicators. Below we cover each domain of the data set and its respective sources.

C.1.1 Water rates and water use

To estimate local rates for water and sewer services, we combine data from two sources. The first source is the Water and Wastewater Rates Dashboard, provided by the Environmental Finance Center (EFC) at the University of North

Carolina. The EFC offers a free online dashboard tool, with detailed information on water and sewer bills at different consumption levels from several utilities in 13 states. We gathered the data available in those dashboards either through published tables or via a custom software code that extracts relevant information from the dashboards. From the dashboards, we include data on 1,356 utilities, from 7 states, that showed the necessary information for our analysis. The second source is the 2016 Water and Wastewater Rate Survey, conducted by the American Water Works Association (AWWA); it includes 264 water and 183 wastewater utilities from 42 states, out of which we consider 189 utilities that have the necessary information on rates and consumer base.[1] The AWWA data are proprietary but can be purchased from AWWA by anyone.

Merging data from AWWA and EFC poses three main challenges: (i) water bills are reported at distinct consumption levels in each source; (ii) the exact geographical boundaries of the area served by each utility are not provided and are difficult to establish in practice; (iii) data from EFC do not contain information on average consumption. Below, we explain how we address each of these issues.

The AWWA survey asks participants to report the year-round (non-seasonal) total monthly bill of residential water and wastewater at six different levels of consumption: 0, 5, 10, 15, and 30 ccf (hundred cubic feet), and at average residential consumption levels. However, the EFC dashboards report total monthly bills at levels from 0 to 15,000 gallons, in increments of 500 gallons. In order to combine these data, we linearly interpolate values from EFC to obtain estimated bills at levels that match those reported by AWWA: 0, 5, 10, and 15 ccf (0, 3740,

[1]For the 26 utilities with at least 200,000 accounts and missing data on sewer (but not water) rates, we imputed the missing values with data obtained directly from these utilities.

7480, and 11220 gallons). Having the data at these four levels, we calculate the total monthly bills at other consumption levels using linear interpolation. Hence, we approximate each utility's rate schedule to block pricing with up to three blocks with different rates: below 5 ccf, between 5 and 10 ccf, and above 10 ccf per month.

Determining the precise area covered by each water utility proved to be a difficult task. In most cases, this information is not available or formatted for software use. To overcome this limitation, our analysis estimates local water rates using county-level averages weighted by the number of accounts in each utility within a county. The data show that in most cases rates within counties are similar. This procedure, however, does not come without loss: in counties with very different rate structures, the weighted average approximates more closely the estimation of water affordability for utilities with more customers.

The EFC dashboards do not report average residential water consumption by utilities. However, this information is needed to estimate affordability based on current consumption. To address this issue, we predict average consumption using ordinary least squares with parameters estimated from the AWWA data set. The details of the estimation procedure are described in the Section C.2.

Figure C.1 illustrates the heterogeneity of water rates across the country. Each horizontal bar represents a county with a population above 500,000; the leftmost, light gray segment is the base rate, and each subsequent segment accounts for the incremental charge of consuming an additional 5 ccf. The white dots show the estimated average bill.

Additionally, in Figure C.2 we plot the unconditional correlation between

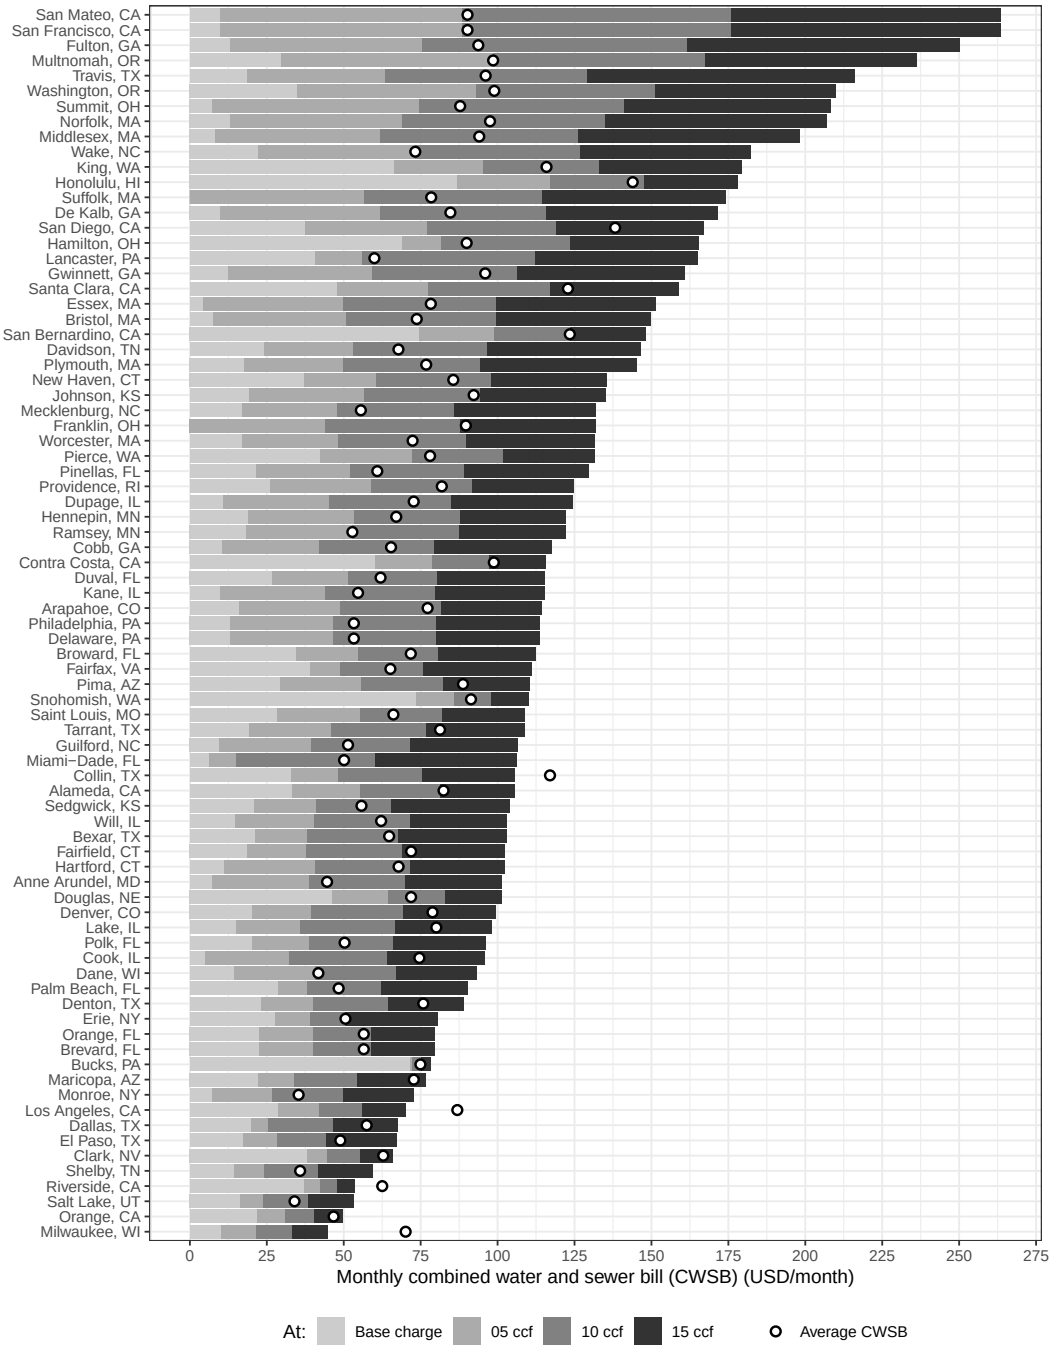

Figure C.1: Combined water and sewer bills in counties with populations greater than 500,000.

Source: Authors' calculations from AWWA and EFC data.

the proportion of the CWSB that is the fixed service fee (for CWSBs evaluated at 50 gppd) and the proportion of households who have unaffordable water

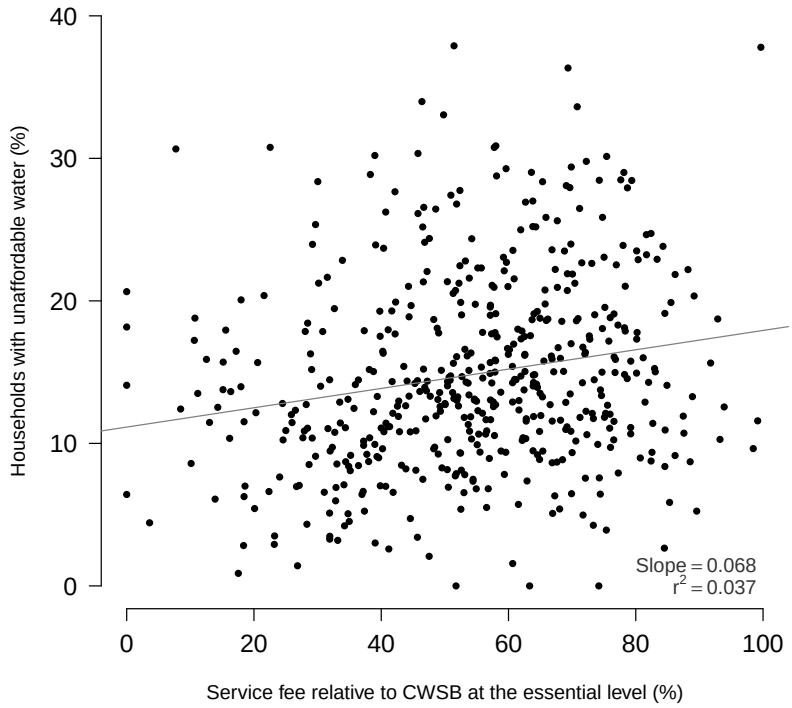

Figure C.2: Correlation between the proportion of the combined water and sewer bill (CWSB) that is the fixed access fee (evaluated at 50 gppd) relative to the proportion of households with unaffordable water and sewer service within a county.

and sewer service within the county. Interestingly, we see a positive correlation between the proportion of the bill that is fixed and affordability concerns. We can infer from this figure that the unconditional correlation between the proportion of CWSBs that is volumetric at essential levels is negatively correlated with affordability concerns. This analysis suggests that affordability concerns are associated with bills that have high fixed service fees.

C.1.2 Climate zones

The definition of climate zones follows the 2004 International Energy Conservation Code (IECC). Each zone is characterized along two dimensions: average temperatures, categorized by a number from 1 (hottest) to 8 (coldest), and humidity, categorized by letters A (humid), B (dry), or C (marine). County-level climate zone data is obtained from the County Characteristics 2000-2007 data set, compiled by the Inter-University Consortium for Political and Social Research (ICPSR 20660).

C.1.3 Socioeconomic-demographic factors

Population, income, and other socioeconomic data come from the 2016 5-Year American Community Survey (ACS), conducted by the US Census Bureau. The publicly available data are reported at the block group level; each census block group typically includes from 600 to 3,000 people. We map block groups to counties using the 12-digit FIPS code identifier.

Of particular interest for our analysis are the income data reported by the ACS. For each block group, there are data on the median annual household income and an estimate of the number of households in each of the 16 income brackets. These values allow us to construct block-group level income distributions and to estimate more precisely the number of households facing affordability issues. Using the information on the population in each block group, we can also aggregate local distributions to obtain a county-level summary of income. Having an income distribution instead of only local medians makes it possible to estimate a distribution of water consumption, which is central to our

analysis.

In addition to data on income, the ACS provides a rich characterization of the social, economic, and demographic factors within a block group. Among the extensive list of indicators available in ACS, we study the correlations between water affordability and a selected set of factors aggregated to the block level group:

1. Population density, measured in persons per square mile.
2. The percentage of the block group population that identify their race as Black or African American alone.
3. The percentage of the block group population that identify being of Hispanic or Latino origin.
4. The percentages of households with income below the Census Bureau poverty threshold[2] and with income between one and two times that threshold.
5. The median age of housing units.
6. The median gross rent as a percentage of the household income.
7. The average household size.
8. The percentage of rented units relative to all occupied units.

[2]This poverty threshold is adjusted for inflation and takes into account the size of the householder's family. More details can be found at the ACS documentation.

C.2 Methods

C.2.1 Preliminaries

This section establishes common definitions and assumptions of our analysis. As described in the paper, we examine metrics using four levels of aggregation (or resolution): county, Census block group, household income brackets within a block group, and a continuum of households within a block group.

The regressions and policy simulations are based on the third resolution level, i.e, a representative household of an income bracket within a block group. For each block group, the ACS provides an estimate of the number of households in each of the sixteen annual income brackets. Using this information, we construct a sixteen-node income distribution per block group, using the center of each bracket as a node. There are two exceptions: (i) for the lowest bracket (up to \$10,000) we set the node at \$7,500, which follows from assuming \$5,000 is the lowest possible income; (ii) for the upper bracket, we set the node at \$250,000, assuming \$300,000 as the highest possible income.

A household is indexed by its income node i, its block group b, and its county c. We denote the income level corresponding to an income node as y_i, and the set of all income nodes as I. Each income node has a probability mass (or share) $f_b(y_i)$ within its block group. Each county c is formed by a set B_c of block groups. Each block group has a probability mass $f_c(b)$ within a county.

A variable x observed at the household level is represented as x_{ibc}. Block-

group averages are then defined as

$$\bar{x}_{bc} = \sum_{i \in I} x_{ibc} f_{bc}\left(y_i\right).$$

County averages are defined as

$$\bar{x}_c = \sum_{b \in B_c} \left[\sum_{i \in I} x_{ibc} f_{bc}\left(y_i\right) \right] f_c\left(b\right).$$

We assume the following concerning water demand functions:

1. County average consumption per capita, $\bar{w}_c$, corresponds to a household with the county median income
2. Household consumption per capita follows:

$$w_{ibc} = \bar{w}_c \left(\frac{y_i}{\bar{y}_c} \right)^{\epsilon_y}, \tag{C.1}$$

where $\epsilon_y = 0.1$ is the constant income elasticity (Havranek et al., 2018). Letting $\bar{h}_{bc}$ denote the average household size in block group b, the household's monthly consumption of water is given by

$$W_{ibc} = 30\bar{h}_{bc} w_{ibc} = 30\bar{h}_{bc}\bar{w}_c \left(\frac{y_i}{\bar{y}_c} \right)^{\epsilon_y}. \tag{C.2}$$

Let $\Phi_c\left(W_{ibc}\right)$ denote the monthly combined water and sewer bill (CWSB) for county c, where W_{ibc} is given in gallons per month. In this paper, we assume Φ_c is piecewise linear to approximate block pricing schedules. This approximation allows for up to three blocks following the AWWA reported data, with break-points at 3,740 and 7,480 gallons per month (or 5 and 10 hundred cubic feet). Then, the estimated CWSB is given by:

$$\hat{\phi}_{ibc} = \Phi_c\left(\hat{W}_{ibc}\right) = \Phi_c\left(\bar{h}_{cb} y_i^{\epsilon_y} \bar{y}_c^{-\epsilon_y} \hat{\bar{w}}_c \hat{\nu}_c\right). \tag{C.3}$$

Hence, the annual share of income corresponding to the CWSB is

$$\hat{s}_{ibc} = \frac{12\hat{\phi}_{ibc}}{y_i}. \tag{C.4}$$

C.2.2 Regression analysis and bootstrapping

This section describes a two-stage procedure to estimate the conditional correlations between water service affordability and socio-economic-demographic characteristics.

First stage: predicting average water and sewer expenditure.

The AWWA data set provides the average household water consumption by utility, W_u, whereas the EFC does not. The first step, then, is to use W_u to (i) calculate the county average water consumption per capita, $\bar{w}_c$, and (ii) use $\bar{w}_c$ to predict $\hat{\bar{w}}_c$ for EFC counties.

If a county has more than one utility, we calculate $\bar{W}_c$ by averaging all values of W_u weighted by the number of accounts. Then, $\bar{w}_c = \frac{\bar{W}_c}{\bar{h}_c}$, where $\bar{h}_c$ is the county average household size.

Next, we estimate the parameter vectors β, δ, γ in the following linear model for county average per capita consumption:

$$\begin{aligned}\log(\bar{w}_c) = \quad & \beta_0 + \beta_1 \log(Pop_c) + \beta_2 \log(\bar{y}_c) + \beta_3 BaseCharge_c + \beta_4 \log(Rate5_c) \\ & + \beta_5 \log(Rate10_c) + \textstyle\sum_{z \in Z} \gamma_z 1(CZ_c = z) + \sum_{s \in S} \delta_s 1(State_c = s) + e_c,\end{aligned} \tag{C.5}$$

where Pop_c is the county population, $\bar{y}_c$ is the county median household income, $BaseCharge_c$ is the minimum service fee charged to households, $Rate5_c$ and

$Rate10_c$ are the marginal combined water and sewer volumetric rates charged between 5 and 10 ccf (3,740 and 7,480 gallons) and above 10 ccf per month, $1(CZ_c = z)$ is a dummy variable equal to 1 if county c is in climate region z and 0 otherwise, $1(State_c = s)$ is a dummy variable equal to 1 if county c is in state s and 0 otherwise, and e_c is the idiosyncratic error term.

Using the estimated parameters in Equation C.5, we predict $\hat{\bar{w}}_c$ for EFC counties. Consequently, we can only estimate $\hat{e}_c$ for counties in the AWWA set, but not for the ones in EFC. For the next stage, it will be helpful to express the exponential of the error term, $\nu = \exp(e_c)$, which allows us to write:

$$\log(\bar{w}_c) = \log\left(\hat{\bar{w}}_c\right) + \hat{e}_c \Rightarrow \bar{w}_c = \hat{\bar{w}}_c \hat{\nu}_c.$$

Summary statistics for the variables used in the first-stage regression are presented in Table C.1. The estimation results of this stage are reported in Table C.2. Despite our aggregation to the county level, point estimates for the elasticities relative to median income (0.0999) and rate at 5–10 ccf (-0.2952) are representative of central values of the ranges estimated in the literature (Dalhuisen et al., 2003; Havranek et al., 2018) and close to the values assumed in our analyses. Nevertheless, we reiterate that the goal of this first stage is the prediction of local average per capita consumption levels rather than the identification of elasticity parameters of aggregate water demand. As such, state and climate zone fixed effects play an important role and explain approximately 30% of the variance in water consumption.

Table C.1: Summary statistics

Panel A: First-stage estimation[a]	Mean	Median	Min	Max	SD
Average consumption (gallons per person per day)	78.06	69.39	43.65	266.14	32.67
Median household income (1,000 USD)	64.01	60.57	31.98	129.74	17.65
Population (1,000 persons)	631.74	329.97	8.15	10,030.21	1,037.18
Base charge (USD)	27.62	22.43	0.00	142.94	19.90
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	8.53	8.01	0.80	23.42	4.23
Volumetric rate above 10 ccf (USD/1,000 gallons)	9.19	8.39	0.80	23.89	4.91
Observations	175				

Panel B: Second-stage estimation[b]	Mean	Median	Min	Max	SD
Households above affordability threshold at estimated level (%)	15.78	11.97	0.00	100.00	13.93
Households above affordability threshold at essential level (%)	11.49	8.06	0.00	100.00	11.67
Households above affordability threshold at 100 gppd (%)	21.07	16.95	0.00	100.00	16.64
Population density (Persons/Sq. mi)	7,141.02	4,461.77	0.13	656,711.81	9,952.46
Average household size (Persons)	2.71	2.64	1.02	9.34	0.67
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	8.39	8.16	0.80	23.42	4.22
Base charge relative to CWSB at essential level (%)	48.49	50.84	0.00	100.00	22.06
Households below poverty level (%)	17.09	12.84	0.00	100.00	14.63
Households between 1 and 2× poverty level (%)	19.75	18.01	0.00	90.19	12.49
Median gross rent relative to income (%)	32.34	31.00	10.00	50.00	9.71
Occupied units that are rented (%)	43.63	39.27	1.19	100.00	25.24
Median age of housing unit (Years)	45.24	43.00	2.00	77.00	18.66
Population identified as Black/African American (%)	16.24	5.37	0.00	100.00	24.24
Population identified as Hispanic/Latino (%)	21.40	10.73	0.00	100.00	25.21
Observations	76,240				

[a] All variables are aggregated to the county level.

[b] Marginal prices are aggregated to the county level. All other variables are aggregated to the block group level.

Intermediate step: calculating local (un)affordability.

Let t be the affordability threshold, usually defined as 4.5% of the annual income. We define the local unaffordability metric, i.e., the percentage of households with CWSBs above the affordability threshold within a block group, as follows:

$$\hat{U}_{bc} = \sum_{i \in I} 1(\hat{s}_{ibc} \geq t) f_{cb}\left(y_i\right),$$

where $1(\hat{s}_{ibc} \geq t)$ is an indicator function receiving 1 if $\hat{s}_{ibc} \geq t$ or 0 otherwise. Thus, $\hat{U}_{bc}$ is effectively a block-group average of the unaffordability indicator. We can also denote $\hat{U}_{bc}$ as a nonlinear function $U\left(\bar{h}_{bc}, \bar{y}_c, \hat{\bar{w}}_c, \hat{\nu}_c, t, \Phi_c, \mathbf{f}_{bc}\right)$, where $\mathbf{f}_{bc}$ is a vector with $f_{bc}\left(y_i\right)$ evaluated at all $i \in I$.

Table C.2: Results for the first-stage regression[a]

	Coef.	SE	95% CI
log(Median household income) (1,000 USD)	0.0999	0.1769	[-0.2468, 0.4467]
log(Population) (Persons)	0.0453	0.0252	[-0.0041, 0.0948]
Base charge (USD)	-0.0027	0.0018	[-0.0063, 0.0008]
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	-0.2952	0.1335	[-0.5568,-0.0336]
Volumetric rate above 10 ccf (USD/1,000 gallons)	0.0094	0.1342	[-0.2535, 0.2724]
Constant	3.3609	1.7617	[-0.0920, 6.8139]
State fixed effects		Yes	
Climate zone fixed effects		Yes	
Observations		175	
R^2		0.5756	
F-statistic		3.10	

[a] Dependent variable is the log of county average per capita water consumption. All variables are defined at the county level.

Similarly, at the county level, we have

$$\hat{U}_c = \sum_{b \in B_c} \left[\sum_{i \in I} 1(\hat{s}_{ibc} \geq t) f_{bc}\left(y_i\right) \right] f_c\left(b\right).$$

Stage two: affordability and socioeconomic-demographic factors.

In this stage, we estimate the linear association between local unaffordability and various factors. In doing so, we note that this estimation does not intend to uncover causal relationships. Instead, the goal of this exercise is to make sense of which socioeconomic and demographic factors are correlated with local water affordability after controlling for prices and other variables.

Let $\mathbf{X}_{bc}$ be a $n \times 1$ vector of n socio-economic-demographic factors of a block group. We want to estimate the $n \times 1$ parameter vector Γ in the following equation:

$$U\left(\bar{h}_{bc}, \bar{y}_c, \hat{\bar{w}}_c, \hat{\nu}_c, t, \Phi_c, \mathbf{f}_{bc}\right) = \hat{U}_{bc} = \mathbf{X}'_{bc}\Gamma + \sum_{z \in Z} \gamma_z 1(CZ_c = z) + \sum_{s \in S} \delta_s 1(State_c = s) + u_{bc}, \tag{C.6}$$

where $1(CZ_c = z)$ and $1(State_c = s)$ are defined as in the first stage, and u_{bc} is the idiosyncratic error term.

A challenge in estimating the standard errors of the previous equation is the fact that $\hat{U}_{bc}$ is measured with error, $\hat{\nu}_{bc}$; the nonlinearity in $\hat{U}_{bc}$ restricts our ability to derive asymptotic properties of the estimators analytically. For this reason, we obtain standard errors by bootstrap. Resampling the variables in Equation C.5 could lead to iterations in which certain states or climate zones would not have estimated coefficients; we implement residual bootstrap to avoid this problem. The procedure can be outlined in six steps:

1. Estimate Equation C.5 to obtain the parameter vectors β, δ, γ and the predicted residuals. Rescale residuals to correct their variance: $\ddot{e}_{bc} = \left(\frac{n}{n-k}\right)^{1/2} \hat{e}_c$, where n is the number of observations and k the number of estimable parameters (MacKinnon, 2006). Let Θ denote the set containing all $\ddot{e}_{bc}$.
2. Calculate $\hat{\bar{w}}_c$ for all counties using the parameters estimated in Equation C.5.
3. For each county, randomly draw $\tilde{e}_c$ from Θ and obtain $\tilde{\nu}_c = \exp(\tilde{e}_c)$.
4. For each block group, calculate $\tilde{U}_{bc} = U\left(\bar{h}_{bc}, \bar{y}_c, \hat{\bar{w}}_c, \tilde{\nu}_c, t, \Phi_c, \mathbf{f}_{bc}\right)$.
5. Estimate Equation C.6 using sample $\left\{\tilde{U}_{bc}, X_{bc}\right\}$.
6. Obtain the empirical distribution of Γ by repeating steps 3-5 many times. In our application, we perform 1,000 bootstrap replications.

The generated distribution of Γ is used to calculate the standard errors of parameter estimates and form their respective confidence intervals.

Summary statistics for the variables used in the second stage are presented in Table C.1. Section C.5 discusses alternative one-stage estimations when the average water consumption is either observed or held at a fixed level.

C.2.3 Policy simulations

We investigate the effects of four hypothetical assistance programs with different benefits and funding options. The benefits can be either a uniform lump-sum transfer or a 50% rate discount for assisted households. The programs can be funded either by uniform water rate increase or local income tax, both considering only non-assisted households.

In our simulations, programs are implemented at the county level. To make fair comparisons, all programs have equal size, as measured by the annual aggregated transfers. To determine the size of the programs in each county, we first adjust water consumption for assisted households based on a rate discount of 50% in all price blocks. Then, we calculate the dollar amount necessary to fund the discounts in a scenario of increased consumption; this amount sets the size of all four programs. Once the size is fixed, we calculate the uniform lump-sum transfers for the programs that offer such benefit. Similarly, we calculate the income tax rate and the price increase necessary to raise funds for their corresponding benefits.

Households are eligible for assistance if their annual expenditure in water services at 50 gppd is greater than or equal to 4.5% of their annual income. Using this threshold, we obtain an estimate of the number of assisted households per county. The assistance programs affect water consumption by either changing

water prices or income of a household. For a price change d and a constant price elasticity ϵ_p, the ratio of water consumption after and before the change is d^{ϵ_p}. Hence, following the notation defined in Section C.2, the expenditure on water services after a rate change d is given by:

$$\hat{\phi}_{ibc} = \Phi_c \left(\hat{W}_{ibc} d^{\epsilon_p} \right) = \Phi_c \left(\bar{h}_{bc} y_i^{\epsilon_y} \bar{y}_c^{-\epsilon_y} \hat{\bar{w}}_c \hat{\nu}_c d^{\epsilon_p} \right). \tag{C.7}$$

The simulations assume a constant elasticity $\epsilon_p = -0.3$ (Dalhuisen et al., 2003). Thus, a price discount of 50% ($d = -0.5$) increases water consumption by about 23% and decreases water expenditure by approximately 38%. Similarly, Equation C.3 is used to adjust consumption after changes in income due to transfers.

We evaluate the results of simulations based on two metrics that emphasize the distribution of expenditure shares rather than point values. The first metric is the interquartile range of expenditure shares per income bracket; this metric allows us to observe not only the effect of a policy in reducing the average burden on poorer households but also how it alleviates those under the most stress by lowering the upper bound. The second metric is the tail (or complementary cumulative) distribution function, which indicates the percentage of households with expenditure share greater than or equal to a specific value; using this curve, we assess how each program promotes redistribution by shifting relative mass from higher to lower expenditure shares.

C.3 Water affordability and socio-economic-demographic factors in US regions

Our main regressions focus on national averages of correlations between water affordability and socio-economic-demographic factors. Nevertheless, heterogeneity is likely to exist across different regions of the US. In this section, we estimate conditional correlations for three regions with higher data availability in our sample: New England, Southeast, and Southwest. The definition of regions follows the Bureau of Economic Analysis (BEA). The New England region includes Connecticut, Massachusetts, Maine, New Hampshire, Rhode Island, and Vermont. The Southeast region includes Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Virginia, and West Virginia. The Southwest region includes Arizona, New Mexico, Oklahoma, and Texas.

Table C.3 presents the estimates and bootstrapped standard errors for regional regressions, with summary statistics presented in Table C.4. With fewer observations than the national case and less variation in the covariates that are correlated regionally, the regional regressions have less statistical power and larger standard errors; this result is particularly salient for price covariates in New England and the Southwest due to higher correlation across block groups in fewer counties.

Although point estimates change across regions, the 95% confidence intervals of regional regressions contain the point estimate of the national regression in almost all cases; the exceptions are discussed below. Overall, poverty and water rates remain important determinants of water affordability in regional

Table C.3: Conditional correlations between water affordability and select socioeconomic characteristics for subsamples containing states in US regions[a]

	All regions Coef. (SE)	New England Coef. (SE)	Southeast Coef. (SE)	Southwest Coef. (SE)
log(Population density) (Persons/Sq. mi)	-0.4935*	-0.1005*	-0.4775*	-0.8586*
	(0.1129)	(0.3536)	(0.2295)	(0.3014)
Average household size (Persons)	1.7333*	2.1567*	3.2451*	1.1892
	(0.3939)	(0.9936)	(0.6257)	(3.088)
log(Volumetric rate at 5–10 ccf) (USD/1,000 gallons)	7.6766*	12.4559*	8.8374*	4.4855
	(1.5523)	(5.8052)	(1.9489)	(8.3736)
Base charge relative to CWSB at essential level (%)	0.0785*	0.0427	0.0758	-0.0357
	(0.031)	(0.0785)	(0.0454)	(0.2752)
Households below poverty level (%)	0.5998*	0.6412*	0.5381*†	0.6346*
	(0.0135)	(0.0454)	(0.0194)	(0.0294)
Households between 1 and 2× poverty level (%)	0.1478*	0.1159*	0.1013*†	0.1304*
	(0.0123)	(0.0395)	(0.0159)	(0.0307)
Median gross rent relative to income (%)	0.0988*	0.0665*	0.0797*	0.0961*
	(0.0076)	(0.0222)	(0.0099)	(0.0189)
Occupied units that are rented (%)	0.0257*	0.0242	0.032*	0.0101
	(0.0076)	(0.0161)	(0.0069)	(0.018)
Median age of housing unit (Years)	0.0129	-0.0318†	0.0324*	0.0187
	(0.0118)	(0.0183)	(0.0144)	(0.022)
Population identified as Black/African American (%)	0.0325*	0.0638	0.0175	-0.0126
	(0.0103)	(0.0476)	(0.0134)	(0.0332)
Population identified as Hispanic/Latino (%)	0.0159	0.0444	-0.1088*†	0.0122
	(0.0168)	(0.0413)	(0.0458)	(0.0349)
Constant	-32.2983*	-33.7254*	-33.6289*	-3.1786
	(8.9454)	(15.8409)	(8.4868)	(41.6757)
State fixed efects	Yes	Yes	Yes	Yes
Climate zone fixed efects	Yes	Yes	Yes	Yes
Observations	76240	6304	20411	9066
R^2	0.6466	0.7221	0.5962	0.6748
F-statistic	2322.96	1167.36	1368.27	1042.89

* denotes statistical significance at 5%. † indicates the 95% confidence interval does not contain the point estimates of the "All regions" regression.

[a]Dependent variable is the proportion of households above the 4.5% water affordability threshold within a block group. Bootstrapped standard errors in parentheses. Water affordability is calculated using reported consumption only. Marginal prices are aggregated to the county level. All other variables are aggregated to the block group level.

regressions, albeit with small variations. In New England, the results suggest that prices and poverty are relatively more important drivers of unaffordability; this is in line with higher water prices and lower poverty rates in this region. In the Southeast, where poverty rates are higher on average, the coefficients on poverty are smaller, and their respective 95% confidence intervals do not contain the point estimates for the national regression; however, poverty coefficients in the Southeast remain relatively large and statistically significant. For

Table C.4: Summary statistics for US regions[a]

(a) New England

	Mean	Median	Min	Max	SD
Households above affordability threshold at estimated level (%)	14.58	10.81	0.00	81.76	13.00
Households above affordability threshold at essential level (%)	10.74	7.41	0.00	77.94	10.70
Households above affordability threshold at 100 gppd (%)	21.77	17.75	0.00	90.64	15.91
Population density (Persons/Sq. mi)	8099.53	4177.46	16.77	179414.23	11419.34
Average household size (Persons)	2.51	2.49	1.16	4.87	0.48
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	12.12	11.49	2.00	17.58	3.46
Base charge relative to CWSB at essential level (%)	29.02	27.93	0.00	91.88	19.77
Households below poverty level (%)	13.71	9.03	0.00	88.61	13.45
Households between 1 and 2× poverty level (%)	15.47	12.97	0.00	78.50	11.27
Median gross rent relative to income (%)	31.45	30.00	10.00	50.00	9.47
Occupied units that are rented (%)	44.13	40.85	2.86	100.00	25.86
Median age of housing unit (Years)	59.33	61.00	4.00	77.00	16.61
Population identified as Black/African American (%)	9.64	2.96	0.00	100.00	16.01
Population identified as Hispanic/Latino (%)	14.50	6.13	0.00	100.00	19.11
Observations	6304				

(b) Southeast

	Mean	Median	Min	Max	SD
Households above affordability threshold at estimated level (%)	15.29	12.19	0.00	96.00	12.93
Households above affordability threshold at essential level (%)	13.39	10.28	0.00	96.00	12.14
Households above affordability threshold at 100 gppd (%)	25.93	22.46	0.00	100.00	17.15
Population density (Persons/Sq. mi)	3230.89	1856.58	0.91	136167.10	5116.18
Average household size (Persons)	2.60	2.57	1.09	5.99	0.54
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	9.37	9.47	1.50	23.03	3.46
Base charge relative to CWSB at essential level (%)	48.48	49.73	6.53	99.85	18.31
Households below poverty level (%)	19.83	16.28	0.00	100.00	14.91
Households between 1 and 2× poverty level (%)	22.27	21.18	0.00	90.19	11.88
Median gross rent relative to income (%)	32.66	31.40	10.00	50.00	10.03
Occupied units that are rented (%)	40.36	35.31	2.00	100.00	23.42
Median age of housing unit (Years)	37.46	36.00	5.00	77.00	15.10
Population identified as Black/African American (%)	26.99	15.62	0.00	100.00	29.03
Population identified as Hispanic/Latino (%)	12.97	5.36	0.00	100.00	19.57
Observations	20411				

(c) Southwest

	Mean	Median	Min	Max	SD
Households above affordability threshold at estimated level (%)	15.90	11.82	0.00	100.00	13.96
Households above affordability threshold at essential level (%)	9.68	6.36	0.00	100.00	10.67
Households above affordability threshold at 100 gppd (%)	18.69	14.23	0.00	100.00	16.07
Population density (Persons/Sq. mi)	5128.33	4291.43	0.14	75048.68	4896.00
Average household size (Persons)	2.80	2.77	1.14	9.34	0.70
Volumetric rate at 5–10 ccf (USD/1,000 gallons)	7.17	5.61	1.86	17.59	3.49
Base charge relative to CWSB at essential level (%)	56.49	56.51	14.05	92.79	16.36
Households below poverty level (%)	18.79	14.46	0.00	100.00	15.63
Households between 1 and 2× poverty level (%)	22.09	20.51	0.00	87.84	13.47
Median gross rent relative to income (%)	30.95	29.30	10.00	50.00	9.43
Occupied units that are rented (%)	43.96	38.26	1.99	100.00	26.34
Median age of housing unit (Years)	35.43	34.00	3.00	77.00	16.33
Population identified as Black/African American (%)	9.56	3.77	0.00	100.00	14.98
Population identified as Hispanic/Latino (%)	37.26	28.81	0.00	100.00	28.44
Observations	9066				

[a] Marginal prices are aggregated to the county level. All other variables are aggregated to the block group level.

the Southwest, the coefficient on price is imprecisely estimated due to a large standard error; coefficients for poverty remain large and within the range of the national averages.

With greater standard errors, the coefficients on race/ethnicity covariates are not statistically significant, though point estimates may be even larger in certain regions. The only exception we observe is in the Southeast: a 10 percentage point increase in the proportion of Hispanic/Latino population in a block group is associated with a 1 percentage point decrease in water unaffordability. This result is in contrast with the national averages and illustrates how the mechanisms of demographic factors may play different roles across different regions.

C.4 Sensitivity analysis

Two key parameters in our analysis are the income and price elasticities of household water consumption. Our baseline scenario assumes an income elasticity of 0.1 and a price elasticity of -0.3, both values taken from point estimates in the literature (Dalhuisen et al., 2003; Havranek et al., 2018); these values are also in line with our estimates for per capita consumption aggregated at the county level (Table C.2). To illustrate the impact of these parameters in our framework, we consider multiple scenarios with alternative combinations of elasticities.

In this sensitivity analysis, we take the lowest absolute value of elasticities to be zero; i.e., household water consumption does not change with income or prices. We discuss alternative scenarios for the upper absolute value of each

elasticities in the subsections below. The first part of this section examines the effect of different income elasticities on our assessment of current unaffordability based on estimated consumption. The second part presents the results for program simulations under four combinations of elasticities.

C.4.1 Affordability assessment at the estimated consumption levels

To assess water affordability at the estimated level of consumption[3], our model predicts the level of water consumption for a household. This prediction adjusts the household per capita consumption based on the county average, the household income relative to the county-average income, and the income elasticity (see Section C.2 for details). In this setting, the income elasticity determines the spread of household per-capita water consumption around the average. A higher income elasticity results in a larger difference in per-capita consumption between households above and below the average income; a zero income elasticity results in all households consuming water at the county average.

Based on meta-analyses (Dalhuisen et al., 2003; Havranek et al., 2018), we consider four scenarios for income elasticity: 0, 0.1, 0.4, and 0.7. The two last cases are based on higher point estimates and upper bounds of 95% interval confidence reported in the meta-analyses. We consider the values of 0 and 0.7 to be extreme cases, as they lead to rather unrealistic consumption patterns in our framework. With income elasticities equal to zero, there is no heterogeneity in

[3]The income elasticity assumption does not affect unaffordability metrics evaluated at the essential level of consumption (50 gppd) as these metrics use a fixed per capita consumption for all households.

water consumption within a county. With income elasticity at 0.7, low-income households adjust more aggressively, which results in approximately 29% of households consuming water below our assumed essential level of 50 gppd; for comparison, with elasticities of 0.1 and 0.4, the approximate percentages of households below 50 gppd are 7% and 20%.

Table C.5 summarizes how different income elasticities affect the number of households facing unaffordable water when evaluated at different resolutions. Once again, we observe that metrics based on the median income greatly underestimate the number of households with unaffordable water, even at the block-group level. As shown in the third column, unaffordability metrics are higher when the income elasticity is zero, as poorer households are assumed to consume at the county average. The fifth column shows that a four-fold increase in the income elasticity from the baseline leads to a decrease in the unaffordability metrics of roughly 5.5 percentage points (or 33%) for the high-resolution cases; despite this change, the metrics in this column are in line with those using the essential level of consumption. Though metrics decrease even further when elasticity is 0.7 (column 6), we note that this comes as a result of poorer households consuming well below 50 gppd.

Table C.5: Percentage of households who pay more than 4.5 percent of annual household income on combined water and sewer bills by income and consumption data resolution under different income elasticity assumptions

		Income elasticity			
Income metric[a]	**Estimated consumption**	ϵ_y **= 0**	ϵ_y **= 0.1**	ϵ_y **= 0.4**	ϵ_y **= 0.7**
Median household income	Income-adjusted	2.55%	1.88%	0.76%	0.36%
Income bracket center	Income-adjusted	15.14%	13.57%	9.11%	5.82%
Income distribution	Income-adjusted	15.14%	13.58%	9.08%	5.44%

[a] Results are presented for three income metrics: (i) "Median household income" represents median incomes at the block-group level; (ii) "Income bracket center" measures incomes at the midpoint of income brackets evaluated at the block-group level; and (iii) "Income distribution" represents the 16-node income distribution with interpolation between nodes evaluated within Census block groups.

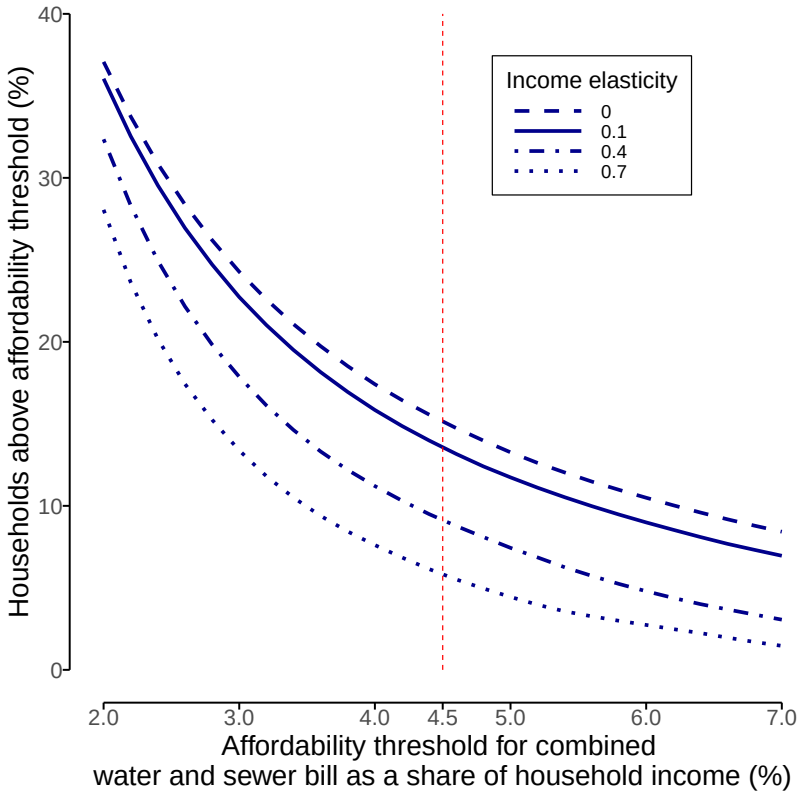

Figure C.3: Proportion of households above the affordability threshold based on water and sewer expenditure-shares of income for different assumptions of income elasticity.

Complementing Table C.5, Figure C.3 shows the distribution of expenditure shares implied by each scenario; the distributions are calculated using the center of income brackets within each block group. The slope of the curves around the 4.5% cutoff point is approximately preserved, meaning that the relative magnitude of the differences presented in Table C.5 is similar if we consider alternative affordability thresholds around 4.5% of income.

C.4.2 Program simulations

The outcomes of the affordability programs we simulate in our paper depend both the income and price elasticities. The income elasticity determines the estimated level of consumption and the behavior of households when their income changes due to programs based on rebates or income taxes. The price elasticity determines how households adjust their water consumption in price-based programs and the size of the programs.

We consider four scenarios: two-by-two combinations of "low" and "high" magnitudes of each elasticity. The "low" case for both elasticities is zero. Following the previous subsection, we define the "high" case for income elasticity as 0.4. For price elasticity, we define the "high" magnitude (in absolute terms) as -0.8. Even though price elasticities below -1 are possible in theory, the programs considered in this paper rely on water demand being inelastic; when the price elasticity is equal to or smaller than -1, water utilities cannot raise revenue to fund programs by increasing prices. Moreover, for elasticities above but close to -1, the implied price increase for non-assisted households becomes unreasonably large.

Figure C.4 shows the average fraction of income spent in water and sewer for each alternative scenario. Bar charts represent averages by income bracket and program; whiskers represent interquartile ranges, with dots representing medians. When the price elasticity is zero, all four programs deliver similar outcomes, as illustrated in panels (a) and (c). If households receiving a 50% rate discount do not increase their consumption, the discount is fully passed on to their expenditures. Moreover, price increases to fund programs are relatively smaller, as non-assisted households do not adjust to a lower consumption level.

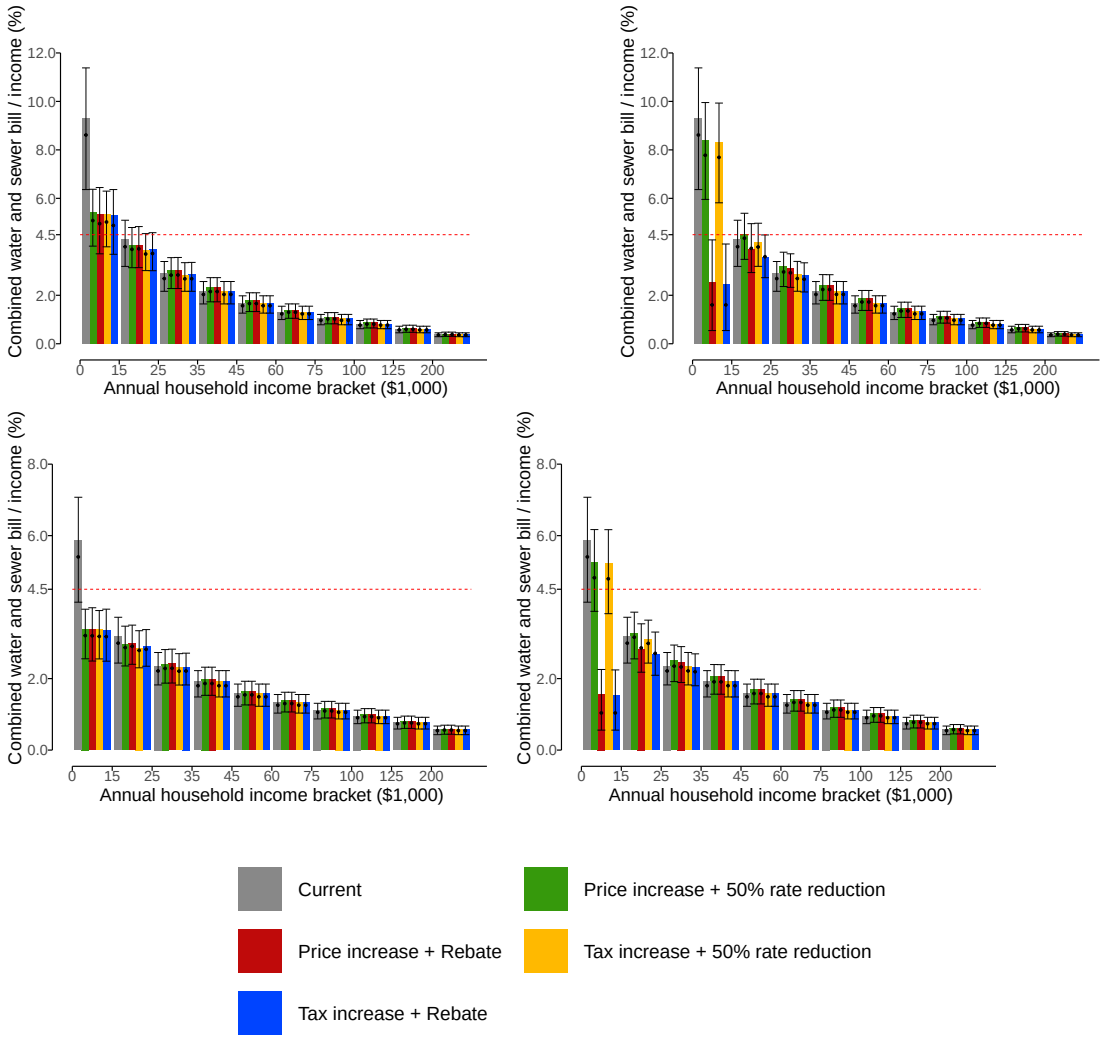

Figure C.4: Average expenditure shares on water and sewer service by income bracket for business-as-usual and each policy option assuming different values of elasticities.

Notes: **Top left panel:** Income elasticity = 0 and price elasticity = 0. **Top right panel:** Income elasticity = 0 and price elasticity = -0.8. **Bottom left panel:** Income elasticity = 0.4 and price elasticity = 0. **Bottom right panel:** Income elasticity = 0.4 and price elasticity = -0.8. In all bar charts, whiskers show the interquartile ranges and dots represent median values.

As shown in panel (c), these results hold even when income elasticity is high (0.4): rebates and taxes represent a tiny fraction of income and result in negligible consumption adjustments.

When price elasticity is high (in absolute value), price-based programs do

little to alleviate unaffordability concerns (panels (b) and (d)). A program based on price discounts funded by price increases may even make the problem worse, as non-assisted households near the eligibility threshold may experience substantially higher water bills (panel (b)); for instance, the average volumetric cost for non-assisted households jumps from around $12 per thousand gallons to above $20 per thousand gallons.

Overall, this analysis shows that price elasticity drives the difference in outcomes between programs, whereas income elasticity has little effect on their relative performances. Hence, except when the price elasticity is zero, we expect programs based on rebates/taxes to be more effective than programs based on price changes in reducing the burden of water bills for assisted households.

C.5 Robustness checks

In this section, we explore two alternatives to the first-stage estimation. The first option is to restrict the sample to counties in which we observe the average per capita water consumption—most counties included in the AWWA survey. The second option is to rely on the assumption that per capita water consumption is fixed across all households. Both approaches avoid estimating levels of consumption, so that conditional correlations between water affordability and socio-economic-demographic factors can be estimated in one stage.

C.5.1 Estimation with observed average consumption

Panel A of Table C.6 reports the results for the subsample in which we observe the average per capita water consumption. We note that most estimates are reasonably close to our main specification, with the possible exception of the effect for households identified as Hispanic. Nevertheless, the point estimates in Table C.6 are all within the 95% confidence intervals reported in Table 3.

C.5.2 Estimation with fixed per capita consumption

Instead of estimating the current per capita water use, another possibility is to assume a fixed, uniform quantity ω for all households. Then, with w_{ibc} equal to ω for all households, we compute U^{ω}_{ibc}, the affordability metric under fixed water use. Finally, we estimate $U^{\omega}_{bc} = \mathbf{X}'_{bc}\mathbf{\Gamma} + u_{bc}$. As the assumption on ω replaces the first stage of the estimation procedure described above, the robust asymptoptic standard errors will be consistent under the usual regularity conditions.

We present results for two different assumptions on ω. The first choice is to set ω at 50 gppd, in which case we assess the conditional correlations between affordability of essential water services and socioeconomic factors. The second choice of ω assumes consumption at 100 gppd. Estimates for both cases are presented in Table C.6, panels B and C. These results are qualitatively similar to our main specifications.

Table C.6: Conditional correlations between water affordability and select socioeconomic characteristics[a]

	(A) Observed consumption only			(B) Fixed consumption at 50 gppd			(C) Fixed consumption at 100 gppd		
	Coef.	SE	95% CI	Coef.	SE	95% CI	Coef.	SE	95% CI
log(Population density) (Persons/Sq. mi)	-0.5057	0.0358	[-0.5758, -0.4355]	-0.5067	0.0274	[-0.5604, -0.4530]	-0.8163	0.0308	[-0.8767, -0.7560]
Average household size (Persons)	1.5908	0.0892	[1.4159, 1.7657]	0.9252	0.0708	[0.7865, 1.0639]	1.9230	0.0815	[1.7632, 2.0827]
log(Volumetric rate at 5–10 ccf) (USD/1,000 gallons)	7.2115	0.1441	[6.9291, 7.4940]	8.4189	0.1116	[8.2002, 8.6377]	14.5106	0.1336	[14.2487, 14.7726]
Base charge relative to CWSB at essential level (%)	0.0711	0.0032	[0.0647, 0.0774]	0.1150	0.0026	[0.1100, 0.1200]	0.0867	0.0027	[0.0814, 0.0920]
Households below poverty level (%)	0.6079	0.0049	[0.5982, 0.6176]	0.4922	0.0042	[0.4840, 0.5005]	0.6976	0.0040	[0.6897, 0.7055]
Households between 1 and 2× poverty level (%)	0.1622	0.0041	[0.1541, 0.1703]	0.1002	0.0033	[0.0938, 0.1066]	0.2608	0.0038	[0.2534, 0.2682]
Median gross rent relative to income (%)	0.1053	0.0042	[0.0972, 0.1135]	0.0580	0.0034	[0.0514, 0.0646]	0.1215	0.0038	[0.1141, 0.1288]
Occupied units that are rented (%)	0.0232	0.0022	[0.0190, 0.0275]	0.0116	0.0017	[0.0082, 0.0150]	0.0196	0.0021	[0.0155, 0.0237]
Median age of housing unit (Years)	0.0096	0.0023	[0.0052, 0.0141]	-0.0048	0.0019	[-0.0086, -0.0011]	0.0069	0.0022	[0.0027, 0.0112]
Population identified as Black/African American (%)	0.0317	0.0024	[0.0270, 0.0363]	0.0188	0.0019	[0.0150, 0.0226]	0.0416	0.0020	[0.0377, 0.0456]
Population identified as Hispanic/Latino (%)	0.0129	0.0028	[0.0075, 0.0183]	-0.0231	0.0023	[-0.0275, -0.0186]	-0.0025	0.0025	[-0.0074, 0.0024]
Constant	-32.3545	0.8503	[-34.0210, -30.6881]	-28.5992	0.6017	[-29.7785, -27.4199]	-35.2743	0.6727	[-36.5928, -33.9557]
State fixed efects	Yes			Yes			Yes		
Climate zone fixed efects	Yes			Yes			Yes		
Observations	58400			76240			76240		
R^2	0.6554			0.5715			0.7280		
F-statistic	1880.71			1693.31			3398.06		

[a]Dependent variable is the proportion of households above the 4.5% water affordability threshold within a block group. Marginal prices are aggregated to the county level. All other variables are aggregated to the block group level.

C.6 Comparison to other metrics in the literature

In this section, we present a detailed comparison of our methodology and findings with two recent studies on the topic of water affordability in the US. The first part replicates the method proposed by Mack and Wrase (2017) to our sample and discusses the results. The second part contrasts our metrics with those proposed by Teodoro (2018) and their results for large US cities.

C.6.1 National affordability assessment in Mack and Wrase

The study by Mack and Wrase (henceforth MW) published in 2017 is one of the first affordability assessments of water and sewer costs at the national level for the US (Mack & Wrase, 2017). Adopting the EPA criterion of 4.5% of income spent on water and sewer bills, MW identify at-risk households based on the median household income at the Census tract level. They find that 11.9% of all households in the continental US face unaffordable water and sewer bills; the estimation is based on 2015 data from the AWWA Survey and the US Census American Community Survey. Furthermore, the study indicates that the number of households facing unaffordable water and sewer bills could triple in five years if the trend of increasing water rates persists.

The results in MW have brought greater attention to the debate of water affordability in the US. However, the validity of their estimates is affected by the overly restrictive assumptions made in their approach, especially those related to water demand.

First, MW assume that all households in the US consume the same amount of

water: 12,000 gallons per month. Such an assumption is the result of assuming a fixed individual consumption of 100 gppd and a fixed household size of 4. In contrast, the 2016 AWWA survey shows the average consumption is 78 gppd, and the 2016 ACS shows the average household size is approximately 2.6. The resulting monthly consumption using 2016 averages is around 6,100 gallons per month—roughly half of the amount assumed in MW.

Second, the approach in MW assumes that all households in the US pay the same for water services. As shown in Figure C.1, however, combined water and sewer bills (CWSBs) can vary by a factor of five across counties with a population above 500,000. Using the AWWA survey, MW estimate the average unit cost of water and sewer services to be \$0.01/gallon. Combined with the volumetric assumption, this implies that all households pay \$120/month in water and sewer bills. Then, based on the EPA's 4.5% criterion, any household with an annual income at or below \$32,000 is deemed to face unaffordable water and sewer. Thus, in practice, the affordability metric proposed in MW is a fixed income threshold.

Third, MW's affordability threshold is evaluated at the median household income of Census tracts. As our results in the main chapter point out, estimates based on median income may largely underestimate affordability concerns, as they are not able to capture the lower end of the income distribution.

To illustrate the consequences of the limiting assumptions in MW, we replicate their approach using our sample. We consider both the unit cost in MW (\$0.01/gallon based on the 2015 AWWA survey) and a unit cost of \$0.0125/gallon, based on the 2016 AWWA survey used in our sample. We highlight that the set of utility companies participating in each survey is different; thus, we cannot

attribute this 25% increase to an average rate increase across utilities. These two unit costs result in fixed monthly CWSBs of $120 and $150, respectively.

Figure C.5 presents the comparison between MW's approach and our assessment using estimated consumption and income at the bracket center. At the 4.5% threshold, these approaches produce dramatically different outcomes: MW's method using the updated unit cost (black line) indicates that 19.4% of households in our sample face unaffordable water and sewer, whereas our preferred assessment (blue line) indicates 13.6%. Furthermore, the distribution of expenditure shares implied by MW's method is substantially steeper for values below 5%. Such a difference in slope is the result of assuming a fixed CWSB, which exaggerates the expenditure of low-income households. As a consequence, the assessment in MW is highly sensitive to changes in the threshold: at 3%, MW's method using the updated unit cost indicates that 49.4% of households in our sample face unaffordable water and sewer, whereas our assessment indicates 22.7%.

Considering MW's original unit cost (gray line in Figure C.5), the proportion of households facing unaffordable water and sewer is 8.6%—below the 11.9% found in MW. This difference is driven by our sample, which skews towards urban areas and, thus, with higher incomes. Figure C.6 demonstrates the difference in the income distributions in our sample and for the entire United States. Compared to the national distribution, households with annual incomes below $75,000 are underrepresented in our sample. Nevertheless, using the full national income distribution may also provide a biased assessment of affordability because rural areas are less likely to receive water and sewer services from utility companies.

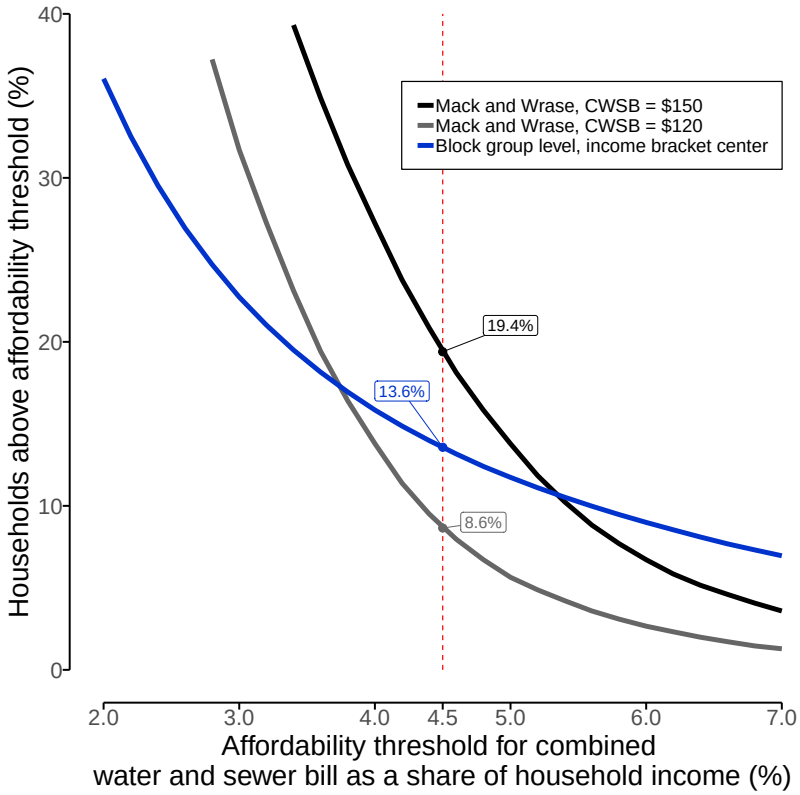

Figure C.5: A comparison of the proportion of households above affordability threshold for Mack and Wrase's approach (MW) and our proposed method.
Notes: The gray line follows MW, considers a unit cost of \$0.01/gallon and a monthly combined water and sewer bill (CWSB) is \$120 for all households in the sample. The black line uses an updated unit cost of \$0.0125/gallon and a monthly CWSB of \$150 for all households in the sample. The blue line considers estimated, income-adjusted consumption based on income bracket centers at the block-group level.

Based on the comparisons above, we contend that our approach delivers a more reliable assessment of water and sewer affordability than MW. Moreover, our method allows policymakers to go one step further and identify which households to target in policies to ameliorate unaffordability concerns.

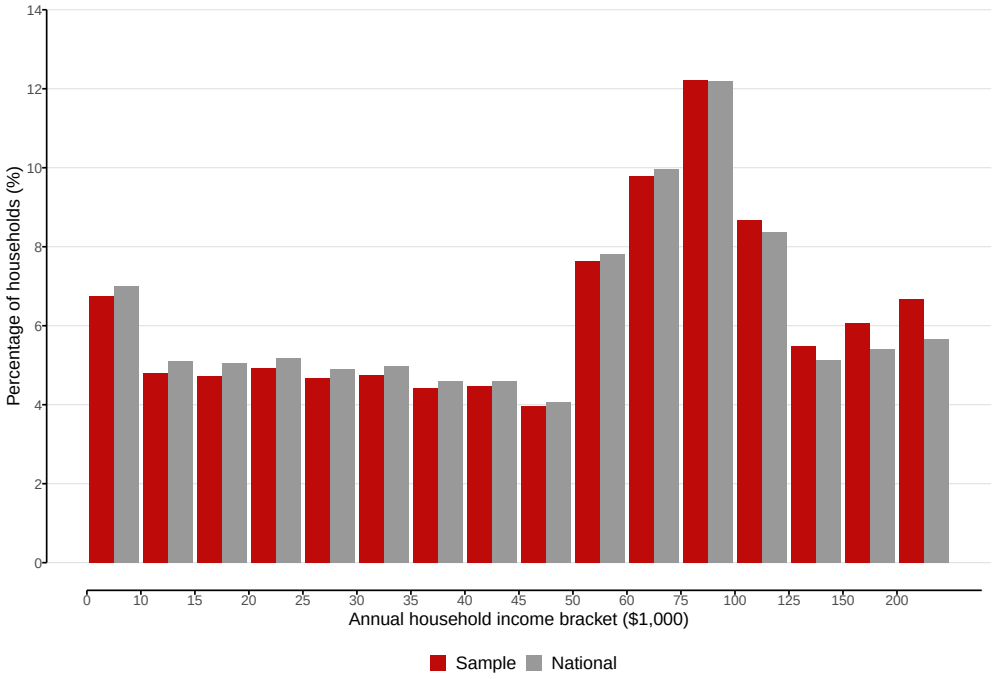

Figure C.6: Household annual income distributions in our sample and for the entire United States.
Source: US Census, American Community Survey.

C.6.2 Local affordability assessments using alternative metrics

A recent study offers a critique of EPA's affordability criterion (i.e., 4.5% of median household income) and proposes two metrics that bring additional considerations to the discussion of water affordability (Teodoro, 2018). In this paper, Teodoro argues that EPA's criterion possesses at least four flaws in its ability to provide an accurate picture of water affordability: (i) it focuses on average use rather than basic needs; (ii) it focuses on median income, thus overlooking poor households; (iii) it disregards other essential living costs; and (iv) it poses a binary standard set at an arbitrary threshold.

To overcome the limitations of EPA's criterion, Teodoro proposes two com-

plementary metrics. The first metric is the affordability ratio, or AR_{20}. This metric is the ratio between the basic monthly water and sewer bill and disposable income evaluated at the 20th income percentile. The disposable income is defined as the household income minus essential nonwater/sewer expenses, such as taxes, health care, food, housing, and home energy.[4] The second metric is hours of labor at minimum wage, or HM. This metric calculates the number of labor hours at the local minimum wage rate necessary to pay for basic water and sewer services.

Teodoro calculates both metrics for the 25 largest cities in the US to illustrate how water and sewer affordability varies across urban areas. To compare our approach to Teodoro's metrics, we compute our metrics for 20 of those cities which are present in our sample.[5] In line with Teodoro's focus on basic consumption, our comparisons assume the essential level of consumption (50 gppd) and use interpolated local income distributions.

Table C.7 presents a comparison between the results following our approach and Teodoro's analysis. Based on the metrics calculated in Teodoro's study, we rank the 20 cities overlapping our sample. Columns 3 and 4 are directly obtained from Teodoro (2018). In column 5, we present our calculated metric, denoted as U_{50}. Columns 6 and 7 display the rank-order of each city according to metrics AR_{20} and HM. Column 8 shows the average rank-order for both of Teodoro's metrics combined. Finally, column 9 shows the rank order based on our U_{50} metric.

To compare the outcomes of each metric, the bottom rows of Table C.7 show

[4]Typically, disposable income is defined as after-tax income.

[5]Our sample does not include New York, NY, Houston, TX, Indianapolis, IN, Washington, DC, and Detroit, MI.

Table C.7: A comparison between Teodoro's AR_{20} and HM metrics for water and sewer affordability and our method for most populous US cities[a]

Population Rank	City, State	AR_{20}[b] (%)	HM (hours)[c]	U_{50} (%)[d]	AR_{20} Rank[e]	HM Rank	AR_{20} & HM Avg. Rank	U_{50} Rank
1	New York, NY	14.1	6.8					
2	Los Angeles, CA	8.2	7	8.6	10.5	16.5	13.5	12.5
3	Chicago, IL	8.2	4.5	6.6	10.5	19	14.75	17
4	Houston, TX	11.7	10.3					
5	Phoenix, AZ	4.8	4	6.3	20	20	20	18
6	Philadelphia, PA	11.2	8.1	17.8	6	13	9.5	1
7	San Antonio, TX	5.9	7.6	9	19	14	16.5	10
8	San Diego, CA	17.1	9.5	15.4	3	6.5	4.75	4
9	Dallas, TX	8.7	8.3	4	8	12	10	19
10	San Jose, CA	8.8	9.9	10.7	7	5	6	9
11	Austin, TX	8.3	12.6	12.1	9	3	6	7
12	Jacksonville, FL	7.8	8.5	12.2	13	11	12	6
13	San Francisco, CA	26.9	13.6	16.8	1	1	1	2
14	Columbus, OH	12.7	13.1	8.9	5	2	3.5	11
15	Indianapolis, IN	13.5	13.5					
16	Fort Worth, TX	8	9.2	8	12	8	10	14
17	Charlotte, NC	6.6	9.5	8.6	17	6.5	11.75	12.5
18	Seattle, WA	18.8	12	14.4	2	4	3	5
19	Denver, CO	7.3	7	7.9	14	16.5	15.25	15
20	El Paso, TX	6.9	7.5	6.9	16	15	15.5	16
21	Washington, DC	14.3	9.8					
22	Boston, MA	16.5	9	16.5	4	10	7	3
23	Detroit, MI	24.4	10.4					
24	Nashville, TN	7.1	9.1	11.3	15	9	12	8
25	Memphis, TN	6.4	5.5	3.7	18	18	18	20
Spearman rank-order correlations								
AR_{20} and HM		0.64						
AR_{20} and U_{50}		0.65						
HM and U_{50}		0.62						
AR_{20} & HM Avg. and U_{50}		0.70						

[a] Our sample does not contain data for 5 of 25 most populous cities: New York, NY, Houston, TX, Indianapolis, IN, Washington, DC, and Detroit, MI. Ranks are calculated for the remaining twenty cities.

[b] As defined in Teodoro (2018), AR_{20} is the affordability ratio at the 20th income percentile: the share of disposable income spent on basic water and sewer for the income level at the 20th percentile.

[c] HM is the hours of labor at minimum wage required to pay for basic water and sewer.

[d] U_{50} is the percentage of households above with combined water and sewer bills greater than 4.5% of annual household income evaluated at an essential consumption level of 50 gppd.

[e] AR_{20} & HM Avg. rank denotes the average ranks of these two metrics combined. *Source:* AR_{20} and HM metrics are from Teodoro (2018). U_{50}, ranks, and correlations are authors' calculations.

Spearman rank-order correlations. As the results indicate, our approach generates similar outcomes and is strongly correlated with the average rank of AR_{20} and HM metrics, with a correlation coefficient of 0.7. It is also markedly correlated with metrics AR_{20} (0.65) and HM (0.62) separately. For a reference of magnitudes, the rank-correlation between both of Teodoro's metrics is 0.64.

The comparisons above illustrate the complementarities between approaches and are not designed to determine a single best metric. Nevertheless, we reiterate some relative advantages our method possesses. First, we

follow the recommendations of the National Academy of Public Administration (National Academy of Public Administration, 2017), with our metrics relying on data available from public sources. Second, we provide metrics that can be directly applied to simulate and implement policies, as our analysis is performed at the representative household level and accounts for the full distribution of income rather than specific percentiles; this feature identifies which household groups face a higher burden of unaffordable water and sewer services. Third, our method is scalable, and our metrics can be calculated at various levels of aggregation. Finally, our metrics are flexible, in the sense that they are not tied to a specific threshold, and present a complete picture of the potential distributional impacts of affordability programs. As illustrated in Figure 6, for example, our methods can inform policymakers about the consequences and potential size of affordability programs at any given income threshold, which is the norm in existing low-income rate assistance for water utilities, as well as for defining household eligibility in other means-tested programs.

REFERENCES

Aguirregabiria, V., & Ho, C.-Y. (2012). A dynamic oligopoly game of the US airline industry: Estimation and policy experiments. *Journal of Econometrics*, *168*(1), 156–173.

Alley, R. B., Marotzke, J., Nordhaus, W. D., Overpeck, J. T., Peteet, D. M., Pielke, R. A., ... others (2003). Abrupt climate change. *Science*, *299*(5615), 2005–2010.

American Water Works Association. (2012). Buried no longer: Confronting America's water infrastructure challenge. *American Water Works Association, Denver, CO*.

Anthoff, D., & Tol, R. S. (2010). On international equity weights and national decision making on climate change. *Journal of Environmental Economics and Management*, *60*(1), 14–20.

Aryal, G., Ciliberto, F., & Leyden, B. T. (2019). Coordinated capacity reductions and public communication in the airline industry. *Becker Friedman Institute for Research in Economics Working Paper*(2018-11).

Azar, C., & Johansson, D. J. (2012). Valuing the non-CO2 climate impacts of aviation. *Climatic Change*, *111*(3-4), 559–579.

Barnett, A. H. (1980). The Pigouvian tax rule under monopoly. *American Economic Review*, *70*(5), 1037–1041.

Bauschke, H., Deutsch, F., Hundal, H., & Park, S.-H. (2003). Accelerating the convergence of the method of alternating projections. *Transactions of the American Mathematical Society*, *355*(9), 3433–3461.

BEA. (2019). *Regional Economic Accounts.* Bureau of Economic Analysis. Retrieved July, 2020, from `https://www.bea.gov/data/economic-accounts/regional`

Belaia, M., Funke, M., & Glanemann, N. (2017). Global warming and a potential tipping point in the atlantic thermohaline circulation: the role of risk aversion. *Environmental and Resource Economics*, *67*(1), 93–125.

Belobaba, P., Odoni, A., & Barnhart, C. (2015). *The global airline industry*. John Wiley & Sons.

Bennett, A., Kallus, N., & Schnabel, T. (2019). Deep generalized method of moments for instrumental variable analysis. In *Advances in neural information processing systems* (pp. 3564–3574).

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Berry, S., Carnall, M., & Spiller, P. T. (2006). Airline hubs: costs, markups and the implications of customer heterogeneity. *Competition policy and antitrust*.

Berry, S., & Jia, P. (2010). Tracing the woes: An empirical analysis of the airline industry. *American Economic Journal: Microeconomics*, *2*(3), 1–43.

Borenstein, S. (1989). Hubs and high fares: dominance and market power in the US airline industry. *The RAND Journal of Economics*, 344–365.

Borenstein, S. (1991). The dominant-firm advantage in multiproduct industries: Evidence from the us airlines. *The Quarterly Journal of Economics*, *106*(4), 1237–1266.

Borenstein, S., & Rose, N. L. (2014). How airline markets work… or do they? regulatory reform in the airline industry. In *Economic regulation and its reform: What have we learned?* (pp. 63–135). University of Chicago Press.

Bovenberg, A. L., & Goulder, L. H. (1996). Optimal environmental taxation in the presence of other taxes: General-equilibrium analyses. *American Economic Review*, *86*(4), 985-1000.

Bovenberg, A. L., & Mooij, R. A. D. (1994). Environmental levies and distortionary taxation. *The American Economic Review*, *84*(4), 1085–1089.

Brueckner, J. K., & Abreu, C. (2017). Airline fuel usage and carbon emissions: Determining factors. *Journal of Air Transport Management*, *62*, 10–17.

Brueckner, J. K., Lee, D., & Singer, E. S. (2013). Airline competition and domestic us airfares: A comprehensive reappraisal. *Economics of Transportation*, *2*(1), 1–17.

Brueckner, J. K., Lee, D. N., Picard, P. M., & Singer, E. (2015). Product unbundling in the travel industry: The economics of airline bag fees. *Journal of Economics & Management Strategy*, *24*(3), 457–484.

BTS. (2019a). *Air Carrier Financial Reports (Form 41 Financial Data).* Bureau of Transportation Statistics. Retrieved July, 2020, from https://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=135&Link=0

BTS. (2019b). *Air Carrier Statistics (Form 41 Traffic) - U.S. Carriers.* Bureau of Transportation Statistics. Retrieved July, 2020, from https://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=110&DB_Name=Air%20Carrier%20Statistics%20(Form%2041%20Traffic)-%20%20U.S.%20Carriers

BTS. (2019c). *Airline On-Time Performance Data.* Bureau of Transportation Statistics. Retrieved July, 2020, from https://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=120&DB_Short_Name=On-Time&DB_Name=Airline%20On-Time%20Performance%20Data&Link=0&DB_URL=Mode_ID=1&Mode_Desc=Aviation&Subject_ID2=0

BTS. (2019d). *Airline Origin and Destination Survey (DB1B).* Bureau of Transportation Statistics. Retrieved July, 2020, from https://www.transtats.bts.gov/DatabaseInfo.asp?DB_ID=125&DB_Short_Name=Origin%20and%20Destination%20Survey&DB_Name=Airline%20Origin%20and%20Destination%20Survey%20%28DB1B%29&Link=0&DB_URL=Mode_ID=1&Mode_Desc=Aviation&Subject_ID2=0

Buchanan, J. M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review*, *59*(1), 174–177.

Cai, Y., Judd, K. L., Lenton, T. M., Lontzek, T. S., & Narita, D. (2015). Environmental tipping points significantly affect the cost- benefit assessment of climate policies. *Proceedings of the National Academy of Sciences*, *112*(15),

4606–4611.

Cai, Y., Lenton, T. M., & Lontzek, T. S. (2016). Risk of multiple interacting tipping points should encourage rapid co2 emission reduction. *Nature Climate Change*, *6*(5), 520–525.

Cai, Y., & Lontzek, T. S. (2019). The social cost of carbon with economic and climate risks. *Journal of Political Economy*, *127*(6), 2684–2734.

Calel, R., Stainforth, D. A., & Dietz, S. (2015). Tall tales and fat tails: the science and economics of extreme warming. *Climatic Change*, *132*(1), 127–141.

California State Assembly. (2015, October). *AB-401 Low-income Water Rate Assistance Program.* Retrieved April 16, 2018, from `https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201520160AB401`

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, *1*(1), 451–488.

Ciliberto, F., & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, *77*(6), 1791–1828.

Ciliberto, F., & Williams, J. W. (2010). Limited access to airport facilities and market power in the airline industry. *The Journal of Law and Economics*, *53*(3), 467–495.

Ciliberto, F., & Williams, J. W. (2014). Does multimarket contact facilitate tacit collusion? inference on conduct parameters in the airline industry. *The RAND Journal of Economics*, *45*(4), 764–791.

Ciliberto, F., & Zhang, Z. (2017). Multiple equilibria and deterrence in airline markets. *Economic Inquiry*, *55*(1), 319–338.

City of Philadelphia. (2015, November). *Amending Title 19 of The Philadelphia Code (Finance, Taxes, and Collections), Chapter 1600 (Water and Sewer Rents).* Retrieved April 16, 2018, from `https://phila.legistar.com/LegislationDetail.aspx?ID=1821444&GUID=EE8B7A07-A75F-4EBD-A0BF-4D71FEB5919B`

Conlon, C., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *RAND Journal of Economics*, *Forthcoming*.

Cremer, H., Gahvari, F., & Ladoux, N. (1998). Externalities and optimal taxation. *Journal of Public Economics*, *70*(3), 343–364.

Crost, B., & Traeger, C. P. (2013). Optimal climate policy: uncertainty versus Monte Carlo. *Economics Letters*, *120*(3), 552–558.

Crost, B., & Traeger, C. P. (2014). Optimal CO_2 mitigation under damage risk valuation. *Nature Climate Change*, *4*(7), 631–636.

Dalhuisen, J. M., Florax, R. J., De Groot, H. L., & Nijkamp, P. (2003). Price and income elasticities of residential water demand: a meta-analysis. *Land Economics*, *79*(2), 292–308.

Daniel, K. D., Litterman, R. B., & Wagner, G. (2019). Declining co2 price paths. *Proceedings of the National Academy of Sciences*, *116*(42), 20886–20891.

EIA. (2018). *Refiner Petroleum Product Prices by Sales Type.* U.S. Energy Information Administration. Retrieved July, 2020, from https://www.eia.gov/dnav/pet/pet_pri_refoth_dcu_nus_m.htm

EIA. (2019). *Carbon dioxide emissions coecficients.* Energy Information Agency. Retrieved from https://www.eia.gov/environment/emissions/co2_vol_mass.php

Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, 937–969.

European Parliament. (2015). *Emission reduction targets for international aviation and shipping.* Directorate General for Internal Policies, European Parliament. Retrieved August, 2020, from https://www.europarl.europa.eu/thinktank/en/document.html?reference=IPOL_STU(2015)569964

Evans, W. N., & Kessides, I. N. (1994). Living by the "golden rule": Multimarket contact in the US airline industry. *The Quarterly Journal of Economics*, *109*(2), 341–366.

FAA. (2020). *Airport & Airway Trust Fund.* Federal Aviation Administration. Retrieved August, 2020, from https://www.faa.gov/about/budget/aatf/

Fabra, N., & Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, *104*(9), 2872–99.

FCCC. (2018). *National greenhouse gas inventory data for the period 1990-2016.* Framework Convention on Climate Change, United Nations. Retrieved August, 2020, from https://unfccc.int/sites/default/files/resource/17e.pdf

Fitzpatrick, L. G., & Kelly, D. L. (2017). Probabilistic stabilization targets. *Journal of the Association of Environmental and Resource Economists*, *4*(2), 611–657.

Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, *124*(1), 249–302.

Fukui, H., & Miyoshi, C. (2017). The impact of aviation fuel tax on fuel consumption and carbon emissions: The case of the us airline industry. *Transportation Research Part D: Transport and Environment*, *50*, 234–253.

Fullerton, D., & Metcalf, G. E. (2002). Cap and trade policies in the presence of monopoly and distortionary taxation. *Resource and energy economics*, *24*(4), 327–347.

Ganapati, S., Shapiro, J. S., & Walker, R. (2016). The Incidence of Carbon Taxes in US Manufacturing: Lessons from Energy Cost Pass-Through. *NBER Working Paper*, *22281*.

Goolsbee, A., & Syverson, C. (2008). How do incumbents respond to the threat of entry? Evidence from the major airlines. *The Quarterly Journal of Economics*, *123*(4), 1611–1633.

Goulder, L. H. (1995). Environmental taxation and the double dividend: a reader's guide. *International tax and public finance*, *2*(2), 157–183.

Goulder, L. H. (1998). Environmental policy making in a second-best setting. *Journal of Applied Economics*, *1*(2), 279–328.

Goulder, L. H., Parry, I. W., Williams Iii, R. C., & Burtraw, D. (1999). The cost-effectiveness of alternative instruments for environmental protection in a second-best setting. *Journal of public Economics*, *72*(3), 329–360.

Greenstone, M., Kopits, E., & Wolverton, A. (2013). Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation. *Review of Environmental Economics and Policy*, *7*(1), 23–46.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 1029–1054.

Hassett, K. A., Mathur, A., & Metcalf, G. E. (2009). The Incidence of a US Carbon Tax: A Lifetime and Regional Analysis. *Energy Journal*, *30*(2).

Havranek, T., Irsova, Z., & Vlach, T. (2018). Measuring the income elasticity of water demand: the importance of publication and endogeneity biases. *Land Economics*, *94*(2), 259–283.

Hayashi, F. (2000). Econometrics. 2000. *Princeton University Press. Section*, *1*, 60–69.

Heutel, G., Moreno-Cruz, J., & Shayegh, S. (2016, December). Climate tipping points and solar geoengineering. *Journal of Economic Behavior & Organization*, *132*, 19–45.

Heutel, G., Moreno-Cruz, J., & Shayegh, S. (2018). Solar geoengineering, uncertainty, and the price of carbon. *Journal of Environmental Economics and Management*, *87*, 24–41.

Hope, C. (2011). *The social cost of CO2 from the PAGE09 model.* Working paper, Judge Business School, Cambridge University.

Howard, P. H., & Sterner, T. (2017). Few and not so far between: A meta-analysis of climate damage estimates. *Environmental and Resource Economics*, *68*(1), 197–225.

IATA. (2019). *World Air Transport Statistics.* International Air Transport Association. Retrieved August, 2020, from `https://www.iata.org/contentassets/a686ff624550453e8bf0c9b3f7f0ab26/wats-2019-mediakit.pdf`

ICAO. (2018). *Carbon Emissions Calculator Methodology, Version 11.* International Civil Aviation Organization. Retrieved July, 2020, from `https://www.icao.int/environmental-protection/CarbonOffset/Documents/Methodology%20ICAO%20Carbon%20Calculator_v11-2018.pdf`

IPCC. (2013). Summary for policymakers [Book Section]. In T. Stocker et al. (Eds.), *Climate change 2013: The physical science basis. contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change* (p. 1–30). Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Jensen, S., & Traeger, C. P. (2014). Optimal climate change mitigation under

long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, *69*, 104–125.

Jerch, R. (2019). *The Local Consequences of Federal Mandates: Evidence from the Clean Water Act* (Tech. Rep.). Cornell University.

Kelly, D. L., & Kolstad, C. D. (1999). Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, *23*(4), 491–518.

Kelly, D. L., & Tan, Z. (2015). Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management*, *72*, 98–122.

Kim, E. H., & Singal, V. (1993). Mergers and market power: Evidence from the airline industry. *American Economic Review*, 549–569.

Klaiber, H. A., Smith, V. K., Kaminsky, M., & Strong, A. (2014). Measuring Price Elasticities for Residential Water Demand with Limited Information. *Land Economics*, *90*(1), 100–113.

Kleven, H. (2020). Sufficient statistics revisited. *NBER Working Paper*, *27242*.

Knutti, R., Rugenstein, M. A. A., & Hegerl, G. C. (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, *6*, 727–736.

Kopp, R. E., Shwom, R. L., Wagner, G., & Yuan, J. (2016). Tipping elements and climate–economic shocks: Pathways toward integrated assessment. *Earth's Future*, *4*(8), 346–372.

Kriegler, E., Hall, J. W., Held, H., Dawson, R., & Schellnhuber, H. J. (2009). Imprecise probability assessment of tipping points in the climate system. *Proceedings of the National Academy of Sciences*, *106*(13), 5041–5046.

Lade, G. E., & Bushnell, J. (2019). Fuel subsidy pass-through and market structure: Evidence from the renewable fuel standard. *Journal of the Association of Environmental and Resource Economists*, *6*(3), 563–592.

Leach, A. J. (2007). The climate change learning curve. *Journal of Economic Dynamics and Control*, *31*(5), 1728–1752.

Lederman, M. (2007). Do enhancements to loyalty programs affect demand? The impact of international frequent flyer partnerships on domestic airline demand. *The RAND Journal of Economics*, *38*(4), 1134–1158.

Lederman, M. (2008). Are frequent-flyer programs a cause of the "hub premium"? *Journal of Economics & Management Strategy*, *17*(1), 35–66.

Lee, D., Fahey, D. W., Forster, P. M., Newton, P. J., Wit, R. C., Lim, L. L., … Sausen, R. (2009). Aviation and global climate change in the 21st century. *Atmospheric Environment*, *43*(22-23), 3520–3537.

Lee, D., & Luengo-Prado, M. J. (2005). The impact of passenger mix on reported" hub premiums" in the US airline industry. *Southern Economic Journal*, 372–394.

Lemoine, D., & Traeger, C. (2014). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, *6*(1), 137–166.

Lemoine, D., & Traeger, C. P. (2016a). Ambiguous tipping points. *Journal of Economic Behavior & Organization*, *132*, 5–18.

Lemoine, D., & Traeger, C. P. (2016b). Economics of tipping the climate dominoes. *Nature Climate Change*, *6*(5), 514–519.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., & Schellnhuber, H. J. (2008). Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences*, *105*(6), 1786–1793.

Leslie, G. (2018). Tax induced emissions? estimating short-run emission impacts from carbon taxation under different market structures. *Journal of Public Economics*, *167*, 220–239.

Lontzek, T. S., Cai, Y., Judd, K. L., & Lenton, T. M. (2015). Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nature Climate Change*, *5*(5), 441–444.

Lontzek, T. S., Narita, D., & Wilms, O. (2016). Stochastic integrated assessment of ecosystem tipping risk. *Environmental and Resource Economics*, *65*(3), 573–598.

Luby, I. H., Polasky, S., & Swackhamer, D. L. (2018). US Urban Water Prices: Cheaper When Drier. *Water Resources Research*, *54*(9), 6126–6132.

Mack, E. A., & Wrase, S. (2017). A Burgeoning Crisis? A Nationwide Assessment of the Geography of Water Affordability in the United States. *PLoS ONE*, *12*(1), e0169488.

MacKinnon, J. G. (2006). Bootstrap methods in econometrics. *Economic Record*, *82*, S2–S18.

Mansur, E. T. (2007). Upstream competition and vertical integration in electricity markets. *The Journal of Law and Economics*, *50*(1), 125–156.

Martins, R., Quintal, C., Cruz, L., & Barata, E. (2016). Water affordability issues in developed countries–The relevance of micro approaches. *Utilities Policy*, *43*, 117–123.

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*(673).

Metcalf, G. E., & Stock, J. H. (2017). Integrated assessment models and the social cost of carbon: A review and assessment of U.S. experience. *Review of Environmental Economics and Policy*, *11*(1), 80–99.

Miranda, M. J., & Fackler, P. L. (2004). *Applied computational economics and finance*. MIT press.

Morrison, S. A. (2001). Actual, adjacent, and potential competition estimating the full effect of Southwest Airlines. *Journal of Transport Economics and Policy*, *35*(2), 239–256.

Muehlegger, E., & Sweeney, R. L. (2017). Pass-Through of Input Cost Shocks Under Imperfect Competition: Evidence from the US Fracking Boom. *NBER Working Paper*, *24025*.

Mumm, J., & Ciaccia, J. (2017). Improving the Narrative on Affordability and the Measurements We Need to Take Us There. *Journal-American Water Works Association*, *109*(5), 42–48.

National Academy of Public Administration. (2017, October). Developing a New Framework for Community Developing a New Framework for Com-

munity Affordability of Clean Water Services. *National Academy of Public Administration*.

National Research Council. (2002). *Abrupt climate change: inevitable surprises*. Washington, DC: National Academy Press.

Nelson, J. P. (2004). Meta-analysis of airport noise and hedonic property values. *Journal of Transport Economics and Policy (JTEP)*, *38*(1), 1–27.

Newbold, S. C., & Daigneault, A. (2009). Climate response uncertainty and the benefits of greenhouse gas emissions reductions. *Environmental and Resource Economics*, *44*(3), 351–377.

Nocedal, J., & Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media.

Nordhaus, W. (1993). Rolling the 'dice': an optimal transition path for controlling greenhouse gases. *Resource and Energy Economics*, *15*(1), 27–50.

Nordhaus, W. (2008). *A question of balance: economic modeling of global warming*. Yale University Press New Haven.

Nordhaus, W. (2014). Estimates of the social cost of carbon: concepts and results from the DICE-2013R model and alternative approaches. *Journal of the Association of Environmental and Resource Economists*, *1*(1/2), 273–312.

NREL. (2001). *Investigation of ByproductApplication to Jet Fuel*. National Renewable Energy Laboratory. Retrieved July, 2020, from `https://www.nrel.gov/docs/fy02osti/30611.pdf`

Olmstead, S. M. (2010). The Economics of Managing Scarce Water Resources. *Review of Environmental Economics and Policy*, *4*(2), 179–198.

Olmstead, S. M., Hanemann, W. M., & Stavins, R. N. (2007). Water demand under alternative price structures. *Journal of Environmental Economics and Management*, *54*(2), 181–198.

Oppenheimer, M., Little, C. M., & Cooke, R. M. (2016). Expert judgement and uncertainty quantification for climate change. *Nature Climate Change*, *6*(5), 445–451.

Overpeck, J. T., & Cole, J. E. (2006). Abrupt change in Earth's climate system. *Annual Review of Environment and Resources*, *31*, 1–31.

Owen, B., Lee, D. S., & Lim, L. (2010). Flying into the Future: Aviation Emissions Scenarios to 2050. *Environmental Science & Technology*, *44*(7), 2255-2260.

Pagoni, I., & Psaraki-Kalouptsidi, V. (2016). The impact of carbon emission fees on passenger demand and air fares: A game theoretic approach. *Journal of Air Transport Management*, *55*, 41–51.

Parry, I. W. (1995). Pollution taxes and revenue recycling. *Journal of Environmental Economics and management*, *29*(3), S64–S77.

Preonas, L. (2017). Market Power in Coal Shipping and Implications for U.S. Climate Policy. *Energy Institute at Haas Working Paper*, *WP285*.

Puterman, M. L., & Shin, M. C. (1978). Modified policy iteration algorithms for discounted markov decision problems. *Management Science*, *24*(11), 1127–1137.

Reiss, P. C., & Spiller, P. T. (1989). Competition and entry in small airline markets. *The Journal of Law and Economics*, *32*(2, Part 2), S179–S202.

Ren, X., Fullerton, D., & Braden, J. B. (2011). Optimal taxation of externalities interacting through markets: A theoretical general equilibrium analysis. *Resource and Energy Economics*, *33*(3), 496–514.

Requate, T. (2006). Environmental policy under imperfect competition. *The International Yearbook of Environmental and Resource Economics*, *2007*, 120–207.

Roe, G. H., & Baker, M. B. (2007). Why is climate sensitivity so unpredictable? *Science*, *318*(5850), 629–632.

Rudik, I. (2020). Optimal climate policy when damages are unknown. *American Economic Journal: Economic Policy*, *12*(2), 340–73.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, *80*(3), 1019–1061.

Sanderson, E., & Windmeijer, F. (2016). A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*, *190*(2), 212–221.

Sandmo, A. (1975). Optimal taxation in the presence of externalities. *The Swedish Journal of Economics*, 86–98.

Schlenker, W., & Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, *83*(2), 768–809.

Stainforth, D., & Calel, R. (2017). Linking physical climate research and economic assessments of mitigation policies. In *Egu general assembly conference abstracts* (Vol. 19, p. 14168).

Stock, J. H., & Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andres & J. H. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas J. Rothenberg.* Cambridge University Press, Cambridge, UK.

Teodoro, M. P. (2018). Measuring Household Affordability for Water and Sewer Utilities. *Journal-American Water Works Association*, *110*(1), 13–24.

Traeger, C. P. (2014). A 4-stated dice: quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics*, *59*(1), 1–37.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

UN General Assembly. (2010). The human right to water and sanitation: resolution / adopted by the General Assembly, 3 August 2010, A/RES/64/292.

U.S. EPA. (1995). *Interim Economic Guidance for Water Quality Standards* (Tech. Rep. No. EPA 832-B-95-002). U.S. Environmental Protection Agency, Office of Water.

U.S. EPA. (1997). *Combined Sewer Overflows—Guidance for Combined Sewer Overflows – Guidance Financial Capability Assessment and Schedule Development* (Tech. Rep. No. EPA 832-B-97-004). U.S. Environmental Protection Agency, Office of Water.

Weil, P. (1990). Nonexpected utility in macroeconomics. *The Quarterly Journal of Economics*, *105*(1), 29–42.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, *91*(1), 1–19.

West, S. E., & Williams, R. C. (2004, March). *Empirical estimates for environmental policy making in a second-best setting* (Working Paper No. 10330). National Bureau of Economic Research. Retrieved from `http://www.nber.org/papers/w10330` doi: 10.3386/w10330

West, S. E., & Williams III, R. C. (2004). Estimates from a consumer demand system: implications for the incidence of environmental taxes. *Journal of Environmental Economics and Management*, *47*(3), 535–558.

Weyant, J. (2017). Some contributions of Integrated Assessment Models of Global Climate Change. *Review of Environmental Economics and Policy*, *11*(1), 115–137.

Whittington, D., Nauges, C., Fuente, D., & Wu, X. (2015). A diagnostic tool for estimating the incidence of subsidies delivered by water utilities in low-and medium-income countries, with illustrative simulations. *Utilities Policy*, *34*, 70–81.

Wichman, C. J. (2014). Perceived price in residential water demand: Evidence from a natural experiment. *Journal of Economic Behavior & Organization*, *107*, 308–323.

Wichman, C. J. (2017). Information provision and consumer behavior: A natural experiment in billing frequency. *Journal of Public Economics*, *152*, 13–33.

Wichman, C. J., Taylor, L. O., & von Haefen, R. H. (2016). Conservation policies: Who responds to price and who responds to prescription? *Journal of Environmental Economics and Management*, *79*, 114–134.

Winchester, N., Wollersheim, C., Clewlow, R., Jost, N. C., Paltsev, S., Reilly, J. M., & Waitz, I. A. (2013). The impact of climate policy on us aviation. *Journal of Transport Economics and Policy*, *47*(1), 1–15.