

Inequality and Environmental Policy: Essays on Climate Adaptation, Health, and
Renewable Energy

by

Emmett Bradley Reynier

A dissertation accepted and approved in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Economics

Dissertation Committee:

Mark Colas, Co-Chair

Edward Rubin, Co-Chair

Eric Zou, Core Member

Nicole Ngo, Institutional Representative

University of Oregon

Winter 2025

© 2025 Emmett Bradley Reynier
This work is openly licensed via CC BY-NC 4.0



DISSERTATION ABSTRACT

Emmett Bradley Reynier

Doctor of Philosophy in Economics

Title: Inequality and Environmental Policy: Essays on Climate Adaptation, Health, and Renewable Energy

This dissertation covers three topics related to inequality and environmental policy, covering, in particular, the distributional effects of climate change, the health effects of pesticides, and solar panel subsidies. Chapter I provides a summary of the work.

Chapter II, co-authored with John Morehouse, examines the distributional effects of climate change. We use a quantitative spatial equilibrium model to simulate the effects of climate change, allowing households to adapt via migration and energy use. We find that low-income and minority households have larger welfare losses from climate change to date and that those disparities in welfare effects will widen under future emissions scenarios. We then simulate the effects of a place-based, means-tested subsidy inspired by an Inflation Reduction Act program. Targeting subsidies to high-climate-exposure cities does a good job of helping the people who are worst off from a climate change perspective. However, it also incentivizes households to move into those high-exposure areas.

Chapter III, co-authored with Edward Rubin, quantifies the causal effect of a common herbicide, glyphosate, on perinatal health. To identify this causal effect, we leverage a natural experiment resulting from the release of genetically modified seeds designed to be resistant to glyphosate. We pair temporal variation in glyphosate use induced by the release of genetically modified seeds with spatial

variation in the suitability of land for growing crops with genetically modified varieties (corn, soy, and cotton). We find glyphosate causes significant decreases in birthweight and gestation length and increases in the probability of low birthweight. We find considerable heterogeneity in this effect, with the most at-risk births having the largest effects.

Chapter IV, co-authored with Mark Colas, determines the optimal spatial distribution of subsidies for residential solar panels. The benefits of solar panel subsidies vary considerably across space due to differences in sunlight, the dirtiness of the electricity grid, and household installation behavior. We pair a model of household solar panel installation behavior with a model of the electricity grid. This model allows us to quantify the optimal schedule of subsidies across states. We find that the current system of subsidies leads to a severe spatial misallocation of solar installations. Additionally, the overall level of subsidies is too high to be justified by the environmental benefits alone. This result suggests that there could be gains from reallocating funds to alternative programs.

This dissertation includes both previously published and unpublished and co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Emmett Bradley Reynier

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, OR, USA
University of Virginia, Charlottesville, VA, USA

DEGREES AWARDED:

Doctor of Philosophy, Economics, 2025, University of Oregon
Master of Science, Economics, 2021, University of Oregon
Bachelor of Science, Commerce, 2017, University of Virginia
Bachelor of Arts, Economics, 2017, University of Virginia

AREAS OF SPECIAL INTEREST:

Environmental Economics
Public Economics
Urban Economics

PROFESSIONAL EXPERIENCE:

ORISE Research Fellow, US EPA Office of Water, 2023-2025
Senior Business Analyst, Carmax, 2019
Business Analyst, Carmax, 2017-2019
Feeds and Search Engine Marketing Intern, Merkle, 2016

GRANTS, AWARDS AND HONORS:

Kleinsorge Summer Research Fellowship, 2022
Dale Underwood Outstanding Graduate Student Scholarship, 2020
Best First Year Econometrics Performance, 2020
First Year Fellowship, 2019-2020

PUBLICATIONS:

Colas, M. and E. Reynier (*Forthcoming*). “Optimal Subsidies for Residential Solar.” *Journal of Political Economy Microeconomics*.

Reynier, E. and E. Rubin (2025). “Glyphosate exposure and GM seed rollout unequally reduced perinatal health.” *Proceedings of the National Academy of Sciences*. 122, 3.

Colas, M. and E. Reynier (2023). “Vertical Migration Externalities.” *Regional Science and Urban Economics*. 101, 103900.

ACKNOWLEDGEMENTS

I owe many thanks to the co-chairs of my committee Mark Colas and Ed Rubin. Their years of guidance and support have been a major impetus for my growth as an economist and person. I would also like to thank Eric Zou, Nicole Ngo, Andrew Dickinson, Philip Economides, Tamara Ren, Giorgi Nikolaishvili, Keaton Miller, Jon Davis, Woan Foong Wong, Benjamin Hansen, Glen Waddell, Todd Doley, Julie Hewitt, Joel Corona, and Kristen Swedberg for their help in many aspects of my work.

The work in this dissertation was supported in part by an appointment to the Research Participation Program at the Water Economics Center, US Environmental Protection Agency, administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the US Department of Energy and EPA. Any views expressed are solely those of the authors and do not necessarily represent the opinions of the US EPA.

Finally, I'd like to thank my wife, Jocelyn, and our family. They have been incredibly supportive of me as I have crafted this dissertation.

To Jocelyn and Rainbow.

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	25
II. THE DISTRIBUTIONAL IMPACTS OF CLIMATE CHANGE ACROSS US LOCAL LABOR MARKETS	28
2.1. Introduction	28
2.2. Data and Descriptive Evidence	35
2.2.1. Data	35
2.2.2. Descriptive Evidence	38
2.3. Model	42
2.3.1. Households	43
2.3.1.1. The impact of climate change on household decisions	47
2.3.2. Firms	50
2.3.3. Rents	51
2.4. Estimation	52
2.4.1. Energy and Housing Demand	52
2.4.1.1. Elasticities of substitution	52
2.4.1.2. Effect of climate on comfort production	55
2.4.2. Labor Supply	56
2.5. Parameter Estimates	58
2.6. The welfare effects of climate change	68
2.6.1. Effect of climate change to-date	69
2.6.2. The black-white gap simulated under future emissions scenarios	74
2.7. IRA inspired place-based, means-tested subsidies	75

Chapter	Page
2.8. Conclusion	85
III. GLYPHOSATE EXPOSURE AND GM SEED ROLLOUT UNEQUALLY REDUCED PERINATAL HEALTH	87
3.1. Main	87
3.1.1. Background and motivation	89
3.1.1.1. Exposure	90
3.1.1.2. Health impacts	90
3.1.2. Empirical approach	92
3.1.2.1. Data	93
3.1.2.2. Estimation	96
3.1.3. Results	99
3.1.3.1. Reduced-form evidence of GM rollout's health effects	99
3.1.3.2. The effect of glyphosate on perinatal health	103
3.1.3.3. Heterogeneity in glyphosate's health impacts	105
3.2. Discussion	107
3.3. Methods	116
3.3.1. Infant health data	116
3.3.2. Pesticide use estimates	117
3.3.3. Attainable yield	117
3.3.4. Additional data	118
3.3.5. Empirical strategy	119
3.3.6. Identification	120
3.3.6.1. Reduced-form DID	120
3.3.6.2. Two-stage least squares	122
3.3.7. Birth-outcome index	124

Chapter	Page
3.3.8. Predicting birthweight	124
3.3.9. Exposure to glyphosate sprayed upstream through water	127
IV. OPTIMAL SUBSIDIES FOR RESIDENTIAL SOLAR	129
4.1. Introduction	129
4.2. Model	137
4.2.1. Households	138
4.2.2. Electricity Production	142
4.2.2.1. Background	142
4.2.2.2. Model: Electricity Production	143
4.2.2.3. Damages	146
4.2.3. Government's Problem and Nationally-Optimal Subsidies . .	146
4.3. Quantitative Model	150
4.3.1. Household preferences	150
4.3.2. Dispatchable Power Plant Production	152
4.3.3. Damages	154
4.4. Data and Estimation	156
4.4.1. Data Sources	156
4.4.2. Descriptive Patterns	159
4.4.3. Estimation	165
4.5. Estimation Results and Model Fit	168
4.5.1. Households	168
4.5.1.1. Parameter Estimates	168
4.5.1.2. Model Fit (Installations)	169
4.5.1.3. Comparison to Existing Estimates	171
4.5.2. Power Plants	172

Chapter	Page
4.6. Counterfactuals and Nationally-Optimal Subsidies	176
4.6.1. Welfare-Maximizing Subsidies	176
4.6.2. Damage-Minimizing Reforms	180
4.6.3. Tract-level Subsidies	182
4.6.4. Unconstrained Reforms	182
4.6.5. Marginal Subsidy Increases	185
4.7. Extensions, Robustness and Further Issues	186
4.7.1. Alternative Specifications of Household Utility	186
4.7.2. Alternative Discounting of Future Subsidy Payments	186
4.7.3. Line Losses	187
4.7.4. Transmission Constraints	188
4.7.5. Improved Storage of Nondispatchable Technology	189
4.7.6. Cleaner Electricity Production	190
4.7.7. Distributional Effects	191
4.7.8. Installation Elasticities with Historical Subsidy Measures	193
4.8. Conclusion	194

APPENDICES

A. APPENDIX FOR THE DISTRIBUTIONAL IMPACTS OF CLIMATE CHANGE ACROSS US LOCAL LABOR MARKETS	195
A.1. Data Appendix	195
A.1.1. Sample Selection	195
A.1.2. City-Level Indices	195
A.1.2.1. Energy Use	195
A.1.2.2. Wages	197

Chapter	Page
A.1.2.3. Housing Quantity and Rent	198
A.2. Additional Descriptive Results	199
A.2.1. Energy demand and climate	201
A.3. Model	207
A.3.1. Derivation of Comfort and Energy Prices	207
A.3.2. Marginal effect of climate on comfort demand	210
A.3.3. Firm FOC derivation	211
A.3.4. Labor Demand Parameters	211
A.3.5. Housing Supply Calibration	212
A.3.6. Equilibrium Definition	213
A.3.7. Equilibrium Outline	215
A.4. Estimation Appendix	216
A.4.1. Alternative specifications for effect of climate on comfort	216
A.4.2. Energy and comfort prices	218
A.4.3. Energy Efficiency: Evidence	222
A.4.4. Moving Cost Parameters for All Years	227
A.4.5. Climate amenities	227
A.4.6. Model fit	227
A.5. Additional Results	229
A.5.1. IRA Subsidies	233
B. APPENDIX FOR GLYPHOSATE EXPOSURE AND GM SEED ROLLOUT UNEQUALLY REDUCED PERINATAL HEALTH	240
B.1. Background	240
B.2. Data	243
B.3. OLS results	245

Chapter	Page
B.4. Comparing rescaled reduced-form DID results with 2SLS results . . .	245
B.5. Shift-share specification	245
B.6. Robustness of first-stage and reduced form results	246
B.7. Robustness of 2SLS results for other outcomes	246
B.8. Prediction performance	247
B.9. 2SLS results and the shape of the glyphosate damage function . . .	248
B.10. Demographic trends	250
B.11. Other forms of heterogeneity	250
B.12. Effect of GM on acreage and yield	251
B.13. Other socioeconomic outcomes	253
B.14. Effects of upstream glyphosate in water	254
B.14.1. Predicting glyphosate in water with machine learning	254
B.14.2. Results: Effect from upstream glyphosate in water	258
C. APPENDIX FOR OPTIMAL SUBSIDIES FOR RESIDENTIAL SOLAR	306
C.1. Data Appendix	306
C.1.1. Deep Solar	306
C.1.2. Google Project Sunroof	306
C.1.3. Tracking the Sun	307
C.1.4. System Advisor Model	308
C.1.5. State Electricity Prices	309
C.1.6. Subsidies	309
C.1.7. Power Plants	311
C.2. Theory and Quantitative Appendix: For Online Publication	316
C.2.1. States Without Net Metering	316

Chapter	Page
C.2.2. Maximum Likelihood Estimation of Power Plant Policy Functions	317
C.2.3. Details: Cost-Neutral Reforms	318
C.2.4. Numerical Algorithm for Calculating Optimal Subsidies	320
C.2.5. Details: Damage-Minimizing Subsidies	320
C.3. Results Appendix: For Online Publication	322
C.3.1. Installation Prices	322
C.3.2. Relationship Between Installations and Monetary Incentives	322
C.3.3. Installation Size Regressions	324
C.3.4. Border Discontinuities in Household Characteristics	326
C.3.5. Border Regressions with Controls for Non-Subsidy Incentives	326
C.3.6. Power Plant Model	327
C.3.7. Type of Subsidy in Nationally-Optimal System	327
C.3.8. Average Panel Size Across Counterfactuals	333
C.3.9. State-Level Results	333
C.3.10. Unconstrained Nationally-Optimal Subsidies with $\lambda = 1$	333
C.4. Extensions and Robustness Appendix: For Online Publication	337
C.4.1. Alternative Specifications of Household Utility	337
C.4.2. Alternative Discounting of Future Subsidy Payments	337
C.4.3. Line Losses	343
C.4.4. Transmission Constraints	344
C.4.5. Improved Storage of Nondispatchable Electricity	347
C.4.6. Cleaner Electricity Production	348
REFERENCES CITED	352

LIST OF FIGURES

Figure		Page
1.	Change in the number of hot and cold days between 1990 and 2019.	38
2.	Relationship between income and climate change.	40
3.	Relationship between share non-white and climate.	41
4.	Estimation results for the effect of climate on comfort production	61
5.	Moving cost estimates for Black and White households	62
6.	Simulations use actual climate data across all years.	68
7.	General equilibrium changes in welfare, energy prices, and population across states	71
8.	Compensating Variation by income decile and race in 2019	72
9.	Time series of the Black-white welfare gap under different emissions scenarios.	75
10.	CV from subsidies relative to CV from climate change to date for different spatial distributions of IRA subsidies.	81
11.	Change in climate with subsidies relative to without subsidies for households in the lowest income decile.	84
12.	GM-crop suitability predicts glyphosate increases after GM-seed introduction.	95
13.	GM-seed introduction increased glyphosate intensity in GM-crop suitable areas; birthweight reductions were also higher in GM-suitable counties and match GM-seed timing.	100
14.	Glyphosate’s direct and policy effects reduced birth outcomes for the average rural birth.	103
15.	Birthweight losses due to glyphosate and GM are largest for births with the lowest expected birthweights.	106

Figure	Page
16. Female infants, children of Black and non-White parents, and children of unmarried parents have lower predicted birthweights. . .	110
17. Expected subsidies and monetary benefit for a 15-panel system in each state.	160
18. Installed solar systems per 1000 individuals.	161
19. Border Discontinuities in Monetary Benefits and Installations.	162
20. Household solar installation model fit.	170
21. Model fit at the interconnection level.	173
22. Model fit at the interconnection level by hour and season.	174
23. Fuel mix of production by interconnection.	175
24. Estimated marginal damage of electricity production by region.	177
25. State-level nationally-optimal subsidies and misallocation for welfare-maximizing reforms.	180
26. State-level nationally-optimal subsidies and misallocation for damage-minimizing reforms.	181
27. Damages offset per additional dollar of government funds	185
A.1. Mean fire and flood risk scores from the First Street Foundation data by census tract.	200
A.2. Change in the no rain days and annual precipitation.	200
A.3. Relationship between share non-white and climate.	201
A.4. Relationship between share non-white and climate.	202
A.5. Black-white energy expenditure gap by income decile.	203
A.6. Heterogeneity in the effect of degree days on energy expenditures by income decile.	205
A.7. Heterogeneity in the effect of degree days on energy expenditures by income decile and race.	206
A.8. Estimation results for the effect of climate on comfort production	217

Figure	Page
A.9. The effect of an additional day at temperature on comfort price by race.	218
A.10.Rent and Energy Prices versus Comfort Price by demographic group	219
A.11.Energy-efficiency by education group across years and CBSAs	223
A.12.Energy-efficiency by education-group across years and CBSAs	224
A.13.Energy-efficiency by demographic-group across years and CBSAs	225
A.14.Estimate of MWTP as a percent of income for an additional day at the given temperature	228
A.15.Model fit when using observed climate to simulate the model.	229
A.16.Change in prices and demand across states.	230
A.17.Compensating Variation by race in 2019 if the climate was the same as in 1990.	231
A.18.College and non-college wages comparing present day climate to that of 1990.	233
A.19.Share of the population living in a disadvantaged community.	234
A.20.Distribution of baseline CV by income group.	234
A.21.CV for 2019 climate with and without subsidies relative to the 1990 climate without subsidies.	235
A.22.CV for 2019 climate with and without subsidies relative to the 1990 climate without subsidies for different magnitudes, G	236
A.23.Percent change in population under different subsidy program magnitudes and distributions.	237
A.24.Percent change in rents under different subsidy program magnitudes and distributions.	238
A.25.Percent change in wages under different subsidy program magnitudes and distributions.	239
B.1. GM crop suitability and increases in glyphosate for rural counties.	262

Figure	Page
B.2. GM seeds were rapidly adopted after their 1996 release.	263
B.3. Interpolation of fertilizer data.	264
B.4. Counties with high suitability for GM crops increased glyphosate intensity and reduced non-glyphosate pesticides with the introduction of glyphosate-resistant seeds.	265
B.5. Perinatal health declined in GM-crop suitable counties after the introduction of glyphosate-resistant seeds	266
B.6. The estimated effect of glyphosate on birthweight is robust to alternative specifications.	269
B.7. Glyphosate effects for infants born to non-white mothers are larger for birthweight and for the probabilities of preterm birth, LBW, and VLBW.	270
B.8. Predicted birthweights closely match actual birthweights across the predicted birthweight distribution.	272
B.9. First-stage event study coefficients are similar across predicted BW quintiles, reduced form shows larger effects in lower quintiles.	274
B.10. Heterogeneity in Glyphosate Effect is consistent across various predicted birthweight bin sizes, greater disparities among birthweight outcomes.	275
B.11. Limited evidence of heterogeneous marginal effects by sex within predicted BW quintile.	276
B.12. Robustness of birthweight effect to alternative economic controls.	279
B.13. Robustness of birthweight effect to alternative farm controls.	280
B.14. Robustness of birthweight effect to alternative instruments.	281
B.15. Robustness of birthweight effect to alternative fixed effects.	282
B.16. Robustness of birthweight effect to alternative fixed effects—within state.	283
B.17. Heterogeneity in Birthweight Effect by Geographic Subsets.	284

Figure	Page
B.18. Birthweight event studies by rural and non-rural counties.	285
B.19. The estimated effect of glyphosate on gestation is robust to alternative specifications.	286
B.20. The estimated effect of glyphosate on the health index is robust to alternative specifications.	287
B.21. Robustness to spatial subsets, all outcomes.	288
B.22. Demographic event studies.	289
B.23. Reduced form heterogeneity by mother's race.	290
B.24. Reduced form heterogeneity by mother's education.	291
B.25. Heterogeneity in effect by different month of birth, all outcomes.	292
B.26. Effect of local Max GM attainable yield on crop acreage.	293
B.27. Effect of local GM max attainable yield on crop yield.	294
B.28. Coefficients from an event study regression of various socioeconomic variables on GM suitability for rural counties.	295
B.29. Predicted glyphosate disaggregated into watersheds in 2004.	296
B.30. Spatial variation in water ML predictors.	297
B.31. Capturing upstream and downstream watersheds.	298
B.32. Cross Validation Results.	299
B.33. Out-of-sample prediction performance for LASSO and Random Forest models.	300
B.34. Density of out-of-sample predictions relative to the actual values.	300
B.35. Aggregating watersheds to counties.	301
B.36. Predicted county-level AMPA in July of 2004.	302
B.37. Effect of upstream glyphosate by distance bin.	304
B.38. Effect of upstream glyphosate by distance bin by high and low soil erodibility and precipitation.	305
C.1. Daily solar radiation ($kWh/m^2/day$) by census tract from Deepsolar. . .	306

Figure	Page
C.2. Hourly electricity generation for a standard solar panel for six example counties.	309
C.3. State electricity prices (\$/ <i>kWh</i>)	310
C.4. Expected subsidy for a 15-panel system by subsidy type in 2014 dollars.	310
C.5. Map of region service areas	314
C.6. Estimation results for solar system price regression	324
C.7. Border Discontinuities in Household Characteristics.	326
C.8. Model fit at the region level in 2020.	330
C.9. Model fit at the region level by hour and season in 2020.	331
C.10. Fuel mix of production by region.	332

LIST OF TABLES

Table	Page
1. Elasticity of substitution estimates.	59
2. Estimation results for non-temperature climate amenities.	64
3. Estimated climate amenity rankings by city for 2019.	65
4. Change in climate amenities from 1990 to 2019.	66
5. Decomposition of welfare effects.	74
6. Regression of log installations per capita on the net present value of subsidies for a 15-panel installation within 10 miles of state borders.	165
7. Parameter estimates for household utility function.	168
8. Welfare maximizing results	178
9. Unconstrained Nationally-Optimal Subsidies.	183
10. Regression of log installations per capita on “historically- adjusted” measures of subsidy generosity.	192
A.1. Effect of CDD and HDD on Energy Expenditures.	204
A.2. Effect of CDD and HDD on Energy Demand.	208
A.3. Cities ranked by the percent difference between their average comfort expenditure and average housing expenditure in 2019.	220
A.4. Cities ranked by the percent difference between their average comfort price and average housing price in 2019.	221
A.5. Cities ranked by their average aggregate energy demand $\mathcal{E}$ and comfort demand $\mathcal{C}$ in 2019.	222
A.6. Household moving costs estimated from Equation (2.31) for 1990.	225
A.7. Household moving costs estimated from Equation (2.31) for 2000.	226

Table	Page
A.8. Household moving costs estimated from Equation (2.31) for 2010.	226
A.9. Household moving costs estimated from Equation (2.31) for 2019.	227
A.10. Decomposition of welfare effects. Values are CV as a percent of income.	231
A.11. Changes in population using 1990 climate as a counterfactual in 2019	232
B.1. Summary statistics for high- and low-yield rural counties and urban counties.	260
B.2. Correlation between GAEZ suitability measures and pre-period acreage for GM crops.	261
B.3. 2SLS estimates of the policy and direct GLY effects on perinatal health.	267
B.4. Comparing difference-in-differences and 2SLS results	268
B.5. Effect of glyphosate on birthweight, gestation, and health index: Robustness to age and race share controls	271
B.6. ML prediction performance	273
B.7. OLS estimates of the policy and direct GLY effects on perinatal health.	277
B.8. Effect of GLY on perinatal health estimated with 2SLS shift-share instrument.	278
B.9. Effect of predicted GLY or AMPA in water on birthweight.	303
C.1. Summary statistics on power plants by region	314
C.2. Summary statistics on power plants by fuel category	315
C.3. Solar system installation prices	323
C.4. Regression of log installations on the net present value of total monetary benefits associated with solar panel installations.	325
C.5. Regression of average panel size on the total monetary benefits associated with solar panel installations.	325

Table	Page
C.6. Regression of log installations per capita on the net present value of subsidies	328
C.7. Regression of log installations per capita on the net present value of subsidies	329
C.8. Percent of total subsidy value from each type of subsidy.	329
C.9. Average number of solar panels in a solar installation by Census region.	333
C.10Cost neutral state level subsidy results.	334
C.11Damage minimizing state level subsidy results.	335
C.12Unconstrained state level subsidy results.	338
C.13Counterfactual results under alternative model specifications.	339
C.14Nationally-optimal subsidies with alternative household discounting. . . .	340
C.15Percent of total subsidy value from each type of subsidy	341
C.16Damages offset per additional dollar of government funds	342
C.17Nationally-optimal subsidies when accounting for line losses.	345
C.18Nationally-optimal subsidies when accounting for transmission constraints.	346
C.19Nationally-optimal cost-neutral and unconstrained subsidies with improved storage technology.	349
C.20Nationally-optimal cost-neutral and unconstrained subsidies under alternative assumptions about central generation energy production.	351

CHAPTER I

INTRODUCTION

My dissertation provides insights into environmental policy-relevant questions. It contains previously published and unpublished co-authored material. Chapter II is joint with John Morehouse, Chapter III with Edward Rubin is published in the *Proceedings of the National Academy of Sciences*, and Chapter IV with Mark Colas is forthcoming in *The Journal of Political Economy Microeconomics*.

Chapter II quantifies the distributional effects of climate change. The spatial distribution of the population is such that low-income and minority households face greater exposure to the changing climate. Further, there is evidence that the same disadvantaged households who face greater exposure to climate change are also less able to adapt. We construct a quantitative spatial equilibrium model where households choose where to live and, conditional on that choice, how much electricity, natural gas, and housing to consume. Thus, households are able to adapt via changes in their energy and housing consumption or by migrating. The climate affects the amenity value of each location and household energy demand.

We find that welfare losses for the lowest income decile are twice as large as those of the highest income decile. Black households have 20% larger welfare losses relative to White households, and these occur throughout the income distribution. These gaps will continue to grow with further emissions and more extreme climate change. We perform a model-based decomposition and find the spatial distribution of the population and differential ability to migrate drive these gaps. Finally, we evaluate a policies that could be used to address these gaps, adding a place-based, means-tested subsidy to the model. Giving places with high climate exposure

the largest subsidies does a good job of targeting the households with the largest welfare losses from climate change but also incentivizes households previously living in other parts of the country to move into these high climate exposure areas.

Chapter III investigates the perinatal health effects of glyphosate. The widespread adoption of genetically modified crops engineered for herbicide tolerance led to a dramatic rise in glyphosate use. Despite glyphosate use increasing by over 750% since the early 90s, we know relatively little about the human health effects of glyphosate. This chapter estimates the causal effect of glyphosate on several infant health outcomes: birthweight, gestation length, and indicators of low and very low birthweight. To identify the causal effect, we leverage temporal variation driven by the release of genetically modified seeds in 1996, paired with spatial variation in suitability for crops with genetically modified varieties. We find that glyphosate reduces average birthweight and gestation length and increases the probability of low and very low birthweight. We find considerable heterogeneity in these effects along the expected birthweight dimension. The most vulnerable births, with the lowest expected birthweight, have effects that are an order of magnitude larger than the least vulnerable births.

Chapter IV studies the optimal design of spatially differentiated subsidies for residential solar panels. We build a structural model of solar panel demand and electricity production across the US and estimate the model by combining 1) remotely sensed data on residential solar panels, 2) power-plant-level data on hourly production and emissions, and 3) a state-of-the-art air pollution model. The current subsidies lead to severe spatial misallocation. National funding for subsidies under the current system exceeds the unconstrained optimum by over

70%. Our results suggest that there could be large welfare gains to redistributing funds towards other programs.

CHAPTER II

THE DISTRIBUTIONAL IMPACTS OF CLIMATE CHANGE ACROSS US LOCAL LABOR MARKETS

This chapter is co-authored with John Morehouse. We started working on this together while we were both graduate students at the University of Oregon. We have jointly worked on idea formulation, data collection, writing code, analyzing results, and writing the paper.

2.1 Introduction

Climate change has led to adverse economic outcomes across the globe as households face rising temperatures, sea levels, and natural disaster risks (U. S. Global Change Research Program, 2023). However, exposure to climate change differs considerably across space, both globally and within the United States (Cruz & Rossi-Hansberg, 2023; Hsiang et al., 2017; Wing et al., 2022). For example, climate change has given cold areas like New England milder winters and hot regions like the Southeast more extreme summers. These changes have significant distributional implications, as hotter locations within the US are generally poorer and less white than the cold locations.

Assessing the welfare effects of climate change requires accounting not only for exposure but also for households' capacity to adapt. Adaptation may exacerbate disparities in climate-related damages across demographic groups if economically advantaged households have access to more effective mitigation strategies. Meanwhile disadvantaged households with limited resources to adapt may experience larger impacts even when facing similar levels of exposure (e.g., (Heilmann, Kahn, & Tang, 2021)). These differences in exposure and adaptive capacity may then deepen the environmental and social inequalities documented in

the environmental justice literature (Cain, Hernandez-Cortes, Timmins, & Weber, 2023).

In this paper, we use a quantitative spatial equilibrium model of US local labor markets in the spirit of Rosen (1979) and Roback (1982) to quantify heterogeneity in the effects of climate change, allowing households to adapt through migration and home-energy use. In the model, heterogeneous households make static choices over (1) where they live and (2) how much energy and housing they consume. Potential wages earned, housing and energy costs, and amenities available in each city affect household’s location choices. Wages and rents are determined in equilibrium. Firms with varying productivity levels across cities demand labor, using college- and non-college-educated workers as imperfect substitutes. Each city has a housing supply curve, with rents responding endogenously to household location choices.

Households adapt to climate change in two ways in the model—through home energy demand and migration. Air conditioning and residential energy use are critical forms of adaptation to climate change (Barreca, Clay, Deschenes, Greenstone, & Shapiro, 2016). We model demand for “comfort,” which households produce using electricity, natural gas, and housing. This notion of comfort production captures the process by which people combine a house’s physical characteristics and the use of appliances such as air conditioners or heaters powered by electricity or natural gas to create a comfortable indoor space (Quigley & Rubinfeld, 1989). The local climate is an additional input to the comfort production function, generating differences in the marginal productivity of electricity, natural gas, and housing. For example, if outdoor temperatures increase, households must use their air conditioners more intensely to maintain the same

indoor temperature, thus increasing the cost of maintaining that temperature. Recent energy justice work suggests that the energy costs associated with climate change will be higher for minority and low-income households (e.g., Lyubich (2020); Reames (2016)). We capture these observed differences in energy demand, allowing households to endogenously adjust their housing, electricity, and natural gas demand in response to changes in outdoor climate.

The climate also affects the shared amenities of each location. These amenities capture how households value the climate when assessing the desirability of a location—for example, households may prefer milder summers. We characterize the amenity value of each city’s climate using heating and cooling degree days, precipitation patterns, and natural disaster risk. Spatial heterogeneity in exposure to climate change will alter some households’ relative rankings of locations—allowing them to mitigate the impact of climate change by moving to a new location (Cruz & Rossi-Hansberg, 2023). However, there is considerable heterogeneity in household mobility, with the literature consistently finding that lower-income households are less mobile than higher-income households (e.g., Bound and Holzer (2000); Colas and Hutchinson (2021); Depro, Timmins, and O’Neil (2015); Piyapromdee (2021)). Thus, low-income households may have a more difficult time migrating to mitigate the welfare effects of climate change.

We quantify the model with publicly available Census data from Ruggles, Flood, Goeken, Schouweiler, and Sobek (2022), historical climate data from PRISM (PRISM Climate Group, 2022), and natural disaster risk data from First Street Foundation (First Street Foundation, 2022a, 2022b). We specify the household “comfort” production function and estimate its parameters using derived demand functions, which are affected by prices of each good and the climate. The estimated

comfort production function allows us to predict energy and housing demand in counterfactual climates. We then estimate the model’s key household location choice parameters using the two-step estimator proposed in Berry, Levinsohn, and Pakes (2004), including the amenity value of climate. The first step uses household-level location choices to estimate moving costs and the component of city-level utility common among households of the same demographic group, which we call the “mean utilities.” The second step decomposes the mean utilities into contributions from wages, rents and energy bills, and climate amenities.

We use our estimated model to simulate the welfare effects from climate change across households in the US, both that have already occurred and that we predict to occur over the next century using state-of-the-art downscaled future climate models shared by the Climate Impact Lab (Gergel et al., 2023). We find significant distributional consequences of climate change both presently and in the future. Black and low-income households are particularly vulnerable, experiencing negative impacts not only due to their initial exposure to climate change but also due to disparities in their adaptive capacity. Black households have suffered 20% greater climate damages compared to white households from climate change to date, and we project this disparity to widen under future emissions scenarios. In the most extreme future emissions projections, we estimate that the Black-white damage gap will increase by more than eightfold by the end of the century. Additionally, the welfare losses from climate change accrue disproportionately to lower-income households. The lowest income decile has welfare losses twice as large as the highest income decile. Notably, the Black-white gap remains throughout the income distribution.

We then decompose the factors contributing to the unequal impacts. First, we simulate a counterfactual where we fix household locations, wages, and rents. We call this the “Mechanical” effect since it does not allow households to reoptimize across space, thus capturing differences in exposure and comfort costs. However, we can eliminate differences in exposure by comparing welfare within cities, therefore isolating differences in comfort costs. Next, we simulate counterfactuals where households can sort across locations, but we fix wages and rents at their baseline levels. We call this the “Sorting” effect as it reflects how differences in mobility impact inequality relative to the Mechanical effect. We can again condition on location in order to isolate differences in mobility from differences in exposure. Finally, we allow complete flexibility—households can sort across locations, and all prices are determined in equilibrium—to demonstrate the effect of endogenous changes in rents and wages on the welfare effects of climate change.

Our model decomposition reveals that the ex-ante distribution of the population is the most significant contributor to these gaps, followed by differences in mobility and, to less of an extent, differences in energy efficiency. When we fix locations, we find that Black households are worse off than white households by 0.8% of income. However, this gap disappears after conditioning on baseline location, suggesting that differences in comfort costs are not contributing to the average gap. When we allow households to sort across locations, the average gap remains at the same level as when fixing locations, while the gap conditional on baseline location increases to 0.5% of income. This result suggests that Black households are less able to mitigate welfare losses through migration relative to similarly exposed white households.

We then use the model to evaluate the effects of a means-tested, place-based policy inspired by the US Environmental Protection Agency’s Community Change Grant program. The Inflation Reduction Act (IRA) allocated \$3 billion to the Community Change Grant program, which funds projects in disadvantaged communities with the stated goal of improving climate resiliency (US EPA, 2023). Our results suggest that while the distribution of funds under the current program helps disadvantaged households generally, it does not target households with the largest climate damages. We then test alternative spatial distributions of the subsidies. Subsidizing places with high climate exposure benefits the households with the greatest losses from climate change. However, these subsidies attract marginal households from elsewhere in the country to high-climate-damage cities, thus increasing total exposure to climate change, increasing rents, and decreasing wages in equilibrium. This spatial reallocation results in some of the intended benefits of the program dissipating as deadweight loss (Kline & Moretti, 2014). Alternatively, we test distributing the subsidies to “climate havens,” places where climate amenities have improved. While this distribution of subsidies reduces exposure to climate change, most of the funds go to households that are already relatively well off from a climate perspective. Policymakers must weigh the benefits of spatially targeted, place-based subsidies with their associated effect on the distribution of households across space, thus increasing climate exposure and resulting in deadweight loss.

Our paper follows a rich literature using spatial equilibrium models to estimate the economic impact of the climate or other environmental goods (e.g., Albouy, Graf, Kellogg, and Wolff (2016); Bayer, Keohane, and Timmins (2009); Hamilton and Phaneuf (2015); Wrenn (2023)). In particular, Albouy et al.

(2016) estimate the amenity value of days across the temperature distribution in the US, a methodology rely on in our own estimation approach. Wrenn (2023) estimates household marginal willingness to pay for natural disasters using a Rosen-Roback model but does not run counterfactuals or explore environmental justice considerations. Another set of papers considers migration an adaptation mechanism to climate change in global spatial equilibrium models (e.g., Cruz and Rossi-Hansberg (2023); Desmet et al. (2021); Desmet and Rossi-Hansberg (2015)), but is more focused on quantifying differences in the welfare impacts across space rather than by demographic group. Rudik, Lyn, Tan, and Ortiz-Bobea (2021) estimate a dynamic spatial equilibrium model with daily temperature affecting amenities and firm productivity. Their estimation strategy relies on migration flows between states, limiting their ability to quantify differential effects between households. However, they incorporate additional mechanisms through which the climate affects the economy—namely it’s effect on firm productivity. Thus allowing for additional adaptation through trade and sectoral switching. We complement their work by formally analyzing the distributional effects of climate change, modeling heterogeneity by race, education, and income.

This paper is also related to work about the effects of climate on energy use (e.g., Auffhammer (2022); Davis and Gertler (2015); Doremus, Jacqz, and Johnston (2022); Rode et al. (2021)). This work inspires our strategy for estimating the impact of climate on energy demand. Auffhammer and Mansur (2014) emphasize a gap in our understanding of climate change’s long-run extensive margin effects on energy demand. Our estimation strategy is able to capture these extensive margin effects, as it reflects changes in energy use due to climate over a nearly 30 year period. Notably, we also address two concerns with the existing intensive margin

literature, (1) we explicitly deal with sorting across locations and (2) we do not assume a constant interior temperature to measure welfare effects.

We also contribute to the environmental justice literature, recently reviewed in Cain et al. (2023). Several papers in this literature study the interaction between household migration decisions and environmental justice (e.g., Bayer et al. (2009); Depro et al. (2015); Hausman and Stolper (2021)). These analyses focus on how sorting contributes to differences in measured exposure to environmental pollutants, as described in Banzhaf, Ma, and Timmins (2019a). We extend this sorting mechanism to climate change.

Methodologically, our model builds on the static quantitative spatial equilibrium literature with heterogeneous workers (e.g., Colas and Morehouse (2022); Diamond (2016); Piyapromdee (2021)). These papers analyze other settings with spatial consequences, allowing rents and wages respond endogenously in general equilibrium to household location choices. Using a similar framework, Morehouse (2022) demonstrates that a fear of distributional consequences is first-order concern for enacting serious climate-change regulation (e.g., a carbon tax). Here, we show that there are also significant distributional effects from a lack of climate policy.

2.2 Data and Descriptive Evidence

We begin by describing the main data sources used throughout the analysis. We break these into three categories: individual household data, climate data, and energy use data. We then show descriptive evidence of the spatial heterogeneity in climate change to-date and how the changes correlate with demographics.

2.2.1 Data.

Household Data. We use repeated cross-sections from the 1990 and 2000 censuses and the 2010 and 2019 5-year aggregated ACS surveys (Ruggles et al., 2022). These give us household demographic information, city, income, employment status, housing, electricity, and natural gas expenditures. We follow standard sample selection and data-cleaning techniques described in Appendix A.1. From these data, we estimate city-education-level wages, city-level rents, and city-demographic group-level energy expenditure for each year. As these indices are standard in the literature, we leave the details in Appendix A.1.2.

Climate Data We collect historical weather data from the PRISM climate group (PRISM Climate Group, 2022). Specifically, we use the 4-kilometer grid of daily average temperature and precipitation to create a panel of daily temperature and precipitation for each of the CBSAs in our model.¹ We aggregate the gridded data to the CBSA level using the weighted average of grid cells within each CBSA, where we weight by the fraction of the CBSA covered by the grid cell and the 1990 population in the grid cell. Population rasters come from SEDAC (CIESIN, 2017). We then construct various annual summary measures of climate in each city. Following the recent literature, we count the number of days each year in a set of discrete temperature bins (Auffhammer, 2022). We also calculate the percent of days with no rain as those with less than one mm of precipitation. Since, in some cases, we are interested in the effects of changes in climate, not short-run weather shocks, we take the 5-year moving averages of each of our annual weather variables.

¹Since the PRISM data only cover the continental US, we collect daily weather station observations from NOAA for stations in Honolulu (Menne et al., 2012), then take the population-weighted average of those observations.

We pair this factual climate data with future climate simulations from the Climate Impact Lab’s global downscaled projections for climate impacts research, CIP-GDPCIR (Gergel et al., 2023). These data result from downscaling and debiasing² 25 climate models for four emissions scenarios, each with daily precipitation and maximum and minimum surface temperature to 2100. We aggregate these to CBSA-level values by taking a population-weighted average as we do with the PRISM data. We then take the average of the 25 climate models to form an ensemble prediction of daily average temperature and total precipitation.

Finally, we use data on natural disaster risks—specifically, First Street Foundation’s fire and flood scores aggregated to the zip-code level (First Street Foundation, 2022a, 2022b). They use state-of-the-art fire and flood models to calculate a risk score between 1 and 10 associated with the respective disasters for every property in the US. We calculate various aggregations of the scores for each CBSA—for example, the median score or the percent of properties in specific score ranges. Appendix A.2 has maps of the average risk scores by census tract. Unfortunately, we only have access to a single set of scores—we use these scores in estimation but cannot vary them in our counterfactual simulations.

Energy Data. We supplement data on household energy expenditure from the census with state-level annual electricity prices from the EIA. We use these prices to back out energy usage from total expenditures. To address the concern that some households may not report energy expenditures separately from rent³, we follow Glaeser and Kahn (2010) in using the Residential Energy Consumption

²The original models produce output on a 1-degree grid, and the CIP-GDPCIR data downscale this to a 0.25-degree grid.

³If their utilities are included as part of rent, for example.




Figure 1. Change in the number of hot and cold days between 1990 and 2019.

Hot days are those with an average temperature above the 90th percentile of the temperature distribution, or about 80F. Cold days are below the 10th percentile, or about 32F. Difference taken between 5 year moving averages and are censored at the 5th and 95th percentiles of grid cells.

Survey (RECS) to correct this issue. The details of this estimation are in Appendix A.1.2.1. We first estimate city-demographic group-level energy use for single-family homeowners from the Census and ACS data, as those households are most likely to report their expenditures accurately. We then use the RECS to estimate the relationship between energy use for single-family homeowners and multi-family homeowners, single-family renters, and multi-family renters.

2.2.2 Descriptive Evidence.

Geography of climate change Climate change is often talked about as increases in average global temperatures, but this masks significant spatial heterogeneity in the degree and impact of change. Figure 1 shows the change in the number of hot and cold days, above the 90th and below the 10th percentiles of the 1990 temperature distribution, respectively, between 1990 and 2019. The hotter portions of the US have seen significant increases in hot days, while the cold places

in the US have seen significant decreases in cold days. Appendix A.2 shows that the eastern US has generally gotten wetter—higher annual precipitation and fewer days with no rain.

Heterogeneity in the impact of climate change is important because it means climate change alters the relative attractiveness of cities, not just the absolute level across all cities. Some cities, such as those in the Northeast, may have become more attractive to households as their winters are now milder than previously. Other cities, such as Oklahoma City or Dallas, might have become less attractive now that their already-hot summers have become hotter. These distortions in climate will lead to changes in the share of the population that chooses to live in each city in the long run, with households taking climate into account when considering the utility they would get from living in a particular city.

Demographics of climate change We are interested in determining whether climate change has differentially affected demographic groups. Here we look at how variation in demographic makeup between cities in 1990 compare to changes in the climate between 1990 and 2019. By looking at demographics in 1990, we avoid endogenous city-choice sorting that has resulted from changes in the climate.

Figure 2 shows a strong positive association between the share of a city’s population in the lowest income quintile in 1990 and how cold and hot days have changed since 1990. Cities with larger shares of low-income households have seen the largest increases in hot days and little change in the number of cold days. Meanwhile, cities with a smaller share of low-income households have seen bigger decreases in cold days and a smaller increase in hot days. Both of these suggest that climate change is regressive—cities with more poor people have had changes in their climate that require spending more on heating and cooling to maintain




Figure 2. Relationship between income and climate change.

Each point is a city, where the size of the point represents the city's 1990 population. Low income defined as households in the bottom quintile of the national income distribution in 1990. Regression lines are weighted by population.

the same indoor temperature relative to cities with fewer poor people. A similar relationship holds for the share of non-white households in cities, as seen in Figure 3. Appendix A.2 also demonstrates that a similar relationship holds for total precipitation.

We can also summarize the change in climate experienced by the average white and non-white household. A common measure of cooling demand is cooling degree days (CDD), which is calculated based on the difference between the daily average temperature and 65°F, summed over all days above 65°F. The average non-white household experienced 38 (5.1%) more CDDs than the average white household in 1990. Between 1990 and 2019, CDDs increased by 110 (14.7%) for white households and 124 (15.9%) for non-white households. Thus, the average




Figure 3. Relationship between share non-white and climate.

Each point is a city, where the size of the point represents the city's 1990 population. Low income defined as households in the bottom quintile of the national income distribution in 1990. Regression lines are weighted by population.

non-white household experienced an additional 14 (13.5%) CDD increase relative to white households, while already having a higher baseline CDD level.

These results suggest that disadvantaged demographic groups have faced disproportionate exposure to the adverse effects of climate change. Lower-income and less white cities have more hot days, while higher-income and more white cities have fewer cold days. However, the welfare effects of climate change may be more or less equitable once we account for adaptation. While the ability to adapt is greater among more economically advantaged households, households that are more exposed have a greater need for adaptation.

Energy use and demographics There is robust evidence from the energy justice literature that Black households spend more on energy than observably similar white households (Lyubich, 2020; Reames, 2016). This gap suggests

that there may also be differences across demographic groups in the response of household energy demand to climate. We explore this in depth in Appendix A.2.1. We first replicate the results in Lyubich (2020), finding a conditional energy expenditure gap between Black and white households throughout the income distribution. We then find that energy expenditures for Black households are more responsive to additional cold weather and less responsive to additional hot weather than white households.

2.3 Model

In this section, we develop a Rosen (1979) and Roback (1982) style spatial equilibrium model where households make a static, discrete choice of where to live. Conditional on that location, households pick consumption of electricity, natural gas, housing, and the numeraire. Cities vary in their wages, rents, and amenities. The climate impacts household choices through two mechanisms: the amenities available in each city and the marginal benefits of electricity, gas, and housing consumption.

Amenities in the model reflect non-market characteristics of cities that make them more or less desirable, including the climate. We use a rich characterization of the climate as amenities to capture a broad range of values that households have for the climate, for example, outdoor recreation or the risk of physical damage from natural disasters. Additionally, households have an imperfect ability to transform the outdoor environment into a comfortable indoor environment with housing and heating or cooling appliances. In the model, households use housing, electricity, and natural gas as inputs to produce indoor “comfort”.⁴ The climate affects the marginal benefits of inputs to the comfort production function. For example, if

⁴We use a similar notion of comfort to that found in Quigley and Rubinfeld (1989).

there are fewer cold days in a year, households require less natural gas to heat their homes, which we capture as a decrease in the marginal benefit of natural gas use.

We allow for rich heterogeneity in the parameters governing household mobility. We specify an individual-level moving cost function to capture the psychic costs of relocating.⁵ Additionally, households have a location-specific idiosyncratic preference shock that allows us to relax the spatial indifference condition of traditional Rosen-Roback models.⁶ The variance of this preference shock governs the elasticities of household location choices. A higher variance means that households have stronger individual preferences for locations relative to the shared components of utility.

Firms in each city produce the tradable numeraire with imperfectly substitutable skilled and unskilled labor. Each city also has an upward-sloping housing supply curve, the slope of which reflects differences in difficulty in developing new housing. These firms and housing supply curves complete the labor and housing markets that allow wages and rents to respond endogenously to household location choices.⁷

2.3.1 Households. Let $j \in J$ index cities and $d \in D$ index demographic groups. We omit t subscripts for exposition, but re-introduce them when describing the estimation. Household i living in city j receives utility from a composite numeraire good, X , “comfort in dwelling” $\mathcal{C}$, and location-specific amenities $\mathbb{A}$. We specify this utility in traditional Cobb-Douglas form,

⁵Bayer et al. (2009) emphasize the importance of including these moving costs.

⁶The spatial indifference condition requires households to receive the same utility from all cities in equilibrium. Busso, Gregory, and Kline (2013) and Kline and Moretti (2014) demonstrate the importance of relaxing such a constraint.

⁷We define an equilibrium for our model in Appendix A.3.6.

$$U_{ij} = X^{1-\alpha_d^c} \mathcal{C}_{dj}^{\alpha_d^c} \exp(\mathbb{A}_{ij}).$$

Households produce comfort using electricity, gas, and housing, which we specify as a nested CES production function. This production function captures the utility services provided by electricity, gas, and housing. For example, electricity does not give people utility directly, but an air conditioner can use electricity to provide cool indoor temperatures in a bedroom. Specifically, let comfort $\mathcal{C}$ be:

$$\mathcal{C}(H, \mathcal{E}|d, j) = (H^{\rho_c} + \mathcal{E}^{\rho_c})^{1/\rho_c}. \quad (2.1)$$

where H is housing and $\mathcal{E}$ is energy. The measure of housing, H , represents both the quantity and quality of housing—i.e. both an additional room in a house and a more recently built house could be interpreted as higher values of H . $\sigma_{\mathcal{C}} = \frac{1}{1-\rho_c}$ is the elasticity of substitution between housing and energy. This aggregation reflects the imperfect complementarity or substitutability between energy and housing. For example, larger houses may require more energy to heat or cool, or a newer home may be better insulated, requiring less energy to heat or cool, all else equal. $\mathcal{E}(\cdot)$ is the household's energy production function, which we parameterize as

$$\mathcal{E}(E, G|d, j) = (\theta_{dj}^E E^{\rho_{\mathcal{E}}} + \theta_{dj}^G G^{\rho_{\mathcal{E}}})^{1/\rho_{\mathcal{E}}}, \quad (2.2)$$

where $\sigma_{\mathcal{E}} = \frac{1}{1-\rho_{\mathcal{E}}}$ is the elasticity of substitution between electricity, E , and natural gas, G . Climate impacts the production of comfort through θ_{dj}^m , where $\theta_{dj}^m = f(\mathbf{Z}_j; \boldsymbol{\kappa}_d^m)$ for $m \in \{E, G\}$, where $\mathbf{Z}_j$ is a vector of local climate variables and $\boldsymbol{\kappa}_d^m$ is a parameter vector governing the potentially non-linear shape of $f(\cdot)$. Since we do not normalize these factor intensities to sum to one, we can think of them as

affecting the productivity of electricity and gas in both relative and absolute terms, as well as the importance of energy relative to housing. We provide more intuition on the effect of climate on housing and energy demand in Section 2.3.1.1.

Finally, we specify amenities as

$$\mathbb{A}_{ij} = \boldsymbol{\alpha}_d^Z \cdot \mathbf{Z}_j + \xi_{dj} + g(j, \mathbf{b}_i) + \sigma_d \varepsilon_{ij},$$

where $\boldsymbol{\alpha}_d^Z$ is a vector of parameters on a vector of climate variables $\mathbf{Z}_j$, ξ_{dj} is a shared, unobservable component of amenities, $g(j, \mathbf{b}_i)$ are moving costs as a function of the household's birth state, $\mathbf{b}_i$ and location j , and ε_{ij} is an idiosyncratic preference draw with variance σ_d . A higher variance in the idiosyncratic utility parameter reduces household mobility, as larger changes in the other components of utility would be required to overcome the idiosyncratic term. We parameterize moving costs, $g(\cdot)$, as:

$$g(j, \mathbf{b}_i) = \gamma_d^{st} \mathbb{I}(j \in b_i^{st}) + \gamma_d^{\text{dist}} \phi(j, b_i^{st}) + \gamma_d^{\text{dist}^2} \phi^2(j, b_i^{st}), \quad (2.3)$$

where $\mathbb{I}(j \in b_i^{st})$ is an indicator for j being in household i 's birth state, $\phi(j, b_i^{st})$ is the Euclidean distance between location j and the agent's birth state b_i^{st} , and $\phi^2(j, b_i^{st})$ is the squared Euclidean distance between j and b_i^{st} . Thus, households get a utility premium, γ_d^{st} , for living in their birth state, and pay an increasing utility cost the further they move from their birth state, captured by the quadratic in distance.

Households are subject to the following budget constraint, conditional on choosing location j ,

$$I_{dj} + \Upsilon_d = X + P_j^E E + P_j^G G + P_j^H H,$$

where I_{dj} is earned income for demographic group d in city j , Υ_d is unearned income for demographic group d , P_j^E and P_j^G are prices of electricity and gas, and P_j^H is the price of housing. We normalize the price of the composite to one.

We can find the household's utility maximizing choices of electricity, gas, housing, and the numeraire by first solving the nested cost minimizations associated with comfort production, the details of which are in Appendix A.3.1. Households choose the cost-minimizing combination of electricity and gas to produce a given amount of energy, which we can then use to solve for a unit cost function for energy. We call this unit cost function the "price of energy",

$$P_{dj}^{\mathcal{E}} = \left(\theta_{dj}^{E^{\sigma_{\mathcal{E}}}} P_j^{E^{1-\sigma_{\mathcal{E}}}} + \theta_{dj}^{G^{\sigma_{\mathcal{E}}}} P_j^{G^{1-\sigma_{\mathcal{E}}}} \right)^{\frac{1}{1-\sigma_{\mathcal{E}}}}. \quad (2.4)$$

Given this price of energy, households choose a cost-minimizing combination of housing and energy to produce a given amount of comfort. Similarly, we can use the conditional demand functions for housing and energy to solve for the unit cost function of comfort, which we'll call the "price of comfort,"

$$P_{dj}^{\mathcal{C}} = \left(P_{dj}^{H^{1-\sigma_{\mathcal{C}}}} + P_{dj}^{\mathcal{E}^{1-\sigma_{\mathcal{C}}}} \right)^{\frac{1}{1-\sigma_{\mathcal{C}}}}. \quad (2.5)$$

The price of comfort is a useful and rarely quantified object that reflects not just the total costs of housing and energy, but also captures how households ability to substitute between housing, electricity, and gas when they face changes in relative prices or productivities.

We can now transform the problem into a simpler form where, conditional on their location choice, households choose comfort and composite good consumption subject to a simplified budget constraint,

$$\max_{X, \mathcal{C}} X^{1-\alpha_d^c} \mathcal{C}^{\alpha_d^c} \exp(\mathbb{A}_{ij}) \quad \text{s.t.} \quad I_{dj} + \Upsilon_d = X + P_{dj}^c \mathcal{C}_{dj}.$$

This gives us the familiar demand functions for consumption and comfort, $X_{dj}^* = (1 - \alpha_{dj}^c)(I_{dj} + \Upsilon_d)$ and $\mathcal{C}_{dj}^* = \alpha_{dj}^c(I_{dj} + \Upsilon_d)/P_{dj}^c$. These demand functions yield the following conditional indirect utility function after taking logs,

$$v_{ij} = \log(I_{dj} + \Upsilon_d) - \alpha_{dj}^c \log(P_{dj}^c) + \boldsymbol{\alpha}_d^Z \cdot \mathbf{Z}_j + \xi_{dj} + g(j, \mathbf{b}_i) + \sigma_d \varepsilon_{ij}. \quad (2.6)$$

Assuming ε_d is distributed type 1 extreme-value, we can write the probability household i chooses location j as:

$$P_{ij} = \frac{\exp(\tilde{v}_{ij})}{\sum_j \exp(\tilde{v}_{ij})}, \quad (2.7)$$

where $\tilde{v}_{ij} = v_{ij} \sigma_d^{-1} - \varepsilon_{ij}$ is the non-idiosyncratic portion of utility.

2.3.1.1 The impact of climate change on household decisions.

In order to clarify the mechanisms through which climate affects household choices in the model, we provide some partial equilibrium derivations here. Consider the effect of climate variable $z_j^l \in \mathbf{Z}_j$ on household conditional indirect utility, holding income, rent, electricity, and gas prices fixed:

$$\frac{\partial v_{ij}}{\partial z_j^l} = \alpha_d^l - \frac{\alpha_d^c}{P_{dj}^c} \times \frac{\partial P_{dj}^c}{\partial z_j^l}. \quad (2.8)$$

There are two channels through which a change in climate affects utility, the amenity value of a location and the price of producing comfort. The change

in amenity value of city j is reflected by $\alpha_d^l \in \boldsymbol{\alpha}_d^Z$. If the change in climate is not desirable, such that α_d^l is negative (i.e., temperatures get hotter in the summer), then the utility households get from living in city j decreases. The second term $\frac{\alpha_d^c}{P_{dj}^c} \times \frac{\partial P_{dj}^c}{\partial z_j^l}$ reflects the effect of the change in the cost of producing comfort inside one's home. Suppose, that the climate variable is one that affects cooling demand, where $\partial \theta^E / \partial z^l = \kappa_E$ is the marginal effect on the electricity share parameter and $\partial \theta^G / \partial z^l = 0$. Taking the derivative of the price of comfort from equation (2.5), we have

$$\frac{\partial P_{dj}^c}{\partial z_j^l} = \frac{-\kappa_E}{\rho_\varepsilon} \frac{P_j^E E_{dj}}{C_{dj}}. \quad (2.9)$$

The sign of this derivative hinges on the elasticity of substitution between electricity and gas, $\sigma_\varepsilon = \frac{1}{1-\rho_\varepsilon}$. If electricity and gas are gross substitutes, such that $\rho_\varepsilon > 0$ and $\sigma_\varepsilon > 1$, then $\partial P_{dj}^c / \partial z_j^l < 0$. If electricity and gas are gross complements, such that $\rho_\varepsilon < 0$ and $\sigma_\varepsilon < 1$, then $\partial P_{dj}^c / \partial z_j^l > 0$. In the gross complements case, increases in the intensity parameters θ^E or θ^G caused by changes in the climate lead to a decrease in produced energy for given levels of electricity and gas, thus decreasing the productivity of the household's comfort function, and increasing the price of comfort.

Plugging equation (2.9) into (2.10) and simplifying, we get

$$\frac{\partial v_{ij}}{\partial z_j^l} = \alpha_d^l + \frac{\kappa_E}{\rho_\varepsilon} \frac{P_j^E E_{dj}}{I_{dj} + \Upsilon_d}. \quad (2.10)$$

We generally expect both of the terms in Equation (2.10) to have the same sign, since a change in climate that makes outdoor amenities less desirable typically also makes it more expensive to produce comfort inside.⁸

Importantly note that $\partial v_{ij}/\partial z_j^l$ is decreasing in income I_{dj} and increasing in electricity use E_{dj} . Thus, all else equal, households with lower incomes, or higher electricity usage conditional on income will be more affected by changes in the climate. Appendix A.3.2 has derivations for the marginal effect of climate on electricity, gas, and housing demand.

Now, we consider the effect of the climate variable on a household's city choice probability P_{ij} . Taking the derivative of equation (2.7), again holding income, rent, electricity, and gas prices constant, we have,

$$\frac{dP_{ij}}{dz_j^l} = \frac{1}{\sigma_d} \frac{\partial v_{ij}}{\partial z_j^l} P_{ij}(1 - P_{ij}). \quad (2.11)$$

This derivative makes it clear that larger values for variance of the idiosyncratic utility shock, σ_d decrease mobility. Higher values of σ_d mean that the unobserved, idiosyncratic portion of utility an individual has a higher relative weight as compared to wages, rents, and shared amenities. If an individual has a particularly strong affinity for a city, for example due to their social connections or cultural ties, this will make them less sensitive to changes in other aspects of the city, such as wages, rent, or the climate. Note that in general equilibrium, a change in the climate in one city can affect indirect utility in all other cities through endogenous changes in wages and rents. Colas and Saulnier (2023) provide a general derivation

⁸The sign of the second term is determined by the sign of κ_E and ρ_E , the rest of the variables in the second term will always be positive. Using our previous example, $\kappa_E > 0$ captures hotter weather requiring more electricity maintain comfort. Thus, it would be opposite sign as α_d^l . If electricity and gas are gross complements, then $\rho_E < 0$.

for this case—the intuition remains the same, higher values of the variance of idiosyncratic portion of utility decrease responsiveness to the shared components of utility.

2.3.2 Firms. Firms competing in perfectly competitive markets combine college-educated labor and non-college-educated labor to produce the composite numeraire good. We parameterize the firms' production function as:

$$Y_j = B_j K_j^\alpha \mathcal{L}_j^{1-\alpha}, \quad (2.12)$$

where B_j is city-specific total factor productivity, K_j is capital use, and $\mathcal{L}_j$ is a CES aggregator between the two labor types. More specifically,

$$\mathcal{L}_j = (\lambda_j S_j^{\rho_l} + (1 - \lambda_j) L_j^{\rho_l})^{\frac{1}{\rho_l}}, \quad (2.13)$$

where λ_j is the college labor input use intensity in city j , S_j is efficiency units of college labor, L_j is the efficiency units of non-college labor, and $\rho_l = \frac{\sigma_l - 1}{\sigma_l}$ is the elasticity of substitution between college and non-college labor.

Assuming capital supply is perfectly elastic and supplied on an international market (at rate $\bar{r}$), we can write the firm's optimal capital demand as

$$K_j^* = \frac{B_j \alpha Y_j}{\bar{r}}.$$

Plugging this into the firm's production function and then solving for the first order-conditions yields the firm's labor demand curves for college- and non-college educated labor:

$$\begin{aligned}
W_j^S &= \tilde{B}_j \mathcal{L}_j^{1-\rho_l} \lambda_j S_j^{\rho_l-1} \\
W_j^L &= \tilde{B}_j \mathcal{L}_j^{1-\rho_l} (1 - \lambda_j) L_j^{\rho_l-1}
\end{aligned} \tag{2.14}$$

where $\tilde{B}_j = (1 - \alpha) B_j^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{\bar{r}}\right)^{\frac{\alpha}{1-\alpha}}$. Appendix A.3.3 has additional details on the firm's profit maximization problem.

2.3.3 Rents. Each city is characterized by a long-run, upward sloping rental supply curve, which we parameterize as

$$\log(P_j^h) = a_j + \zeta_j \log(H_j), \tag{2.15}$$

where P_j^h is the price of housing in city j , a_j is a city-specific intercept capturing construction costs, ζ_j is the inverse elasticity of housing supply, and H_j is the total housing demand in the city. The intercept, a_j , varies across cities due to differences in baseline construction costs, such as the cost of building materials or labor. The elasticity, ζ_j , captures how responsive housing prices are to changes in the quantity of housing demanded. The availability of land for development and local regulatory constraints, such as zoning laws and land-use restrictions, affect the inverse elasticity—more restrictive housing development conditions lead to steeper increases in housing prices.

We assume that rental profit, Π , is distributed back to landowners lump-sum as unearned income, $\Upsilon_d = s_d \Pi$, where s_d is demographic group d 's share of investment income. Rental profits are

$$\Pi = \sum_j \left(P_j^H H_j - \frac{\exp a_j}{\zeta_j + 1} H_j^{\zeta_j+1} \right) \tag{2.16}$$

2.4 Estimation

In this section, we detail our estimation procedure. We focus our exposition on the estimation of the household demand and labor supply parameters. We use city-level energy, wage, and rent indices in estimating various parts of the model. Details for how we construct these indices can be found in Appendix A.1.2.

2.4.1 Energy and Housing Demand. We estimate the parameters of the household's nested CES comfort function in two steps. First, we use the relative demand functions to estimate the elasticity of substitution between electricity and gas, $\sigma_{\mathcal{E}}$, and then the elasticity of substitution between energy and housing, $\sigma_{\mathcal{C}}$ in first differences. Then, we use the implied values for the θ 's to estimate the effect that climate has on comfort production.

2.4.1.1 Elasticities of substitution. We first derive the first-order conditions from households choosing the cost minimizing combination of electricity and gas to produce a certain quantity of energy. The details of this cost-minimization problem are in Appendix A.4.1, which yields the following relative demand function for electricity and gas,

$$\log \left(\frac{E_{djt}}{G_{djt}} \right) = \sigma_{\mathcal{E}} \log \left(\frac{P_{jt}^G}{P_{jt}^E} \right) + \sigma_{\mathcal{E}} \log \left(\frac{\theta_{djt}^E}{\theta_{djt}^G} \right). \quad (2.17)$$

Since we observe quantities demanded and prices for both electricity and gas, we can estimate $\sigma_{\mathcal{E}}$ using linear regression. We take first differences of Equation (2.17) to arrive at our estimating equation,

$$\Delta \log \left(\frac{E_{djt}}{G_{djt}} \right) = \sigma_{\mathcal{E}} \Delta \log \left(\frac{P_{jt}^G}{P_{jt}^E} \right) + \Delta u_{djt}, \quad (2.18)$$

As we are concerned with endogeneity when regressing quantities of electricity and gas demanded on the prices of electricity of gas due to the simultaneity of supply and demand, we follow Burke and Abayasekara (2018) and instrument the log ratio of retail electricity and natural gas prices with the log ratio of lagged commercial electricity and natural gas prices, where $P_{jt}^{m,Comm}$ for $m \in \{E, G\}$ is the price of electricity or natural gas for commercial customers. Formally, the exclusion restriction is

$$E \left[\Delta \log \left(\frac{P_{j,t-1}^{G,Comm}}{P_{j,t-1}^{E,Comm}} \right) \Delta u_{djt} \right] = 0, \quad (2.19)$$

where u_{djt} is any other factor affecting the change in the ratio of electricity and gas demand. By instrumenting with lagged commercial prices, we use variation in energy prices that is driven by supply side factors and not correlated with other factors affecting residential energy demand, including the climate.

We can then back out $\theta_{djt}^E/\theta_{djt}^G$ using equation (2.17), the estimate of $\sigma_{\mathcal{E}}$, and our data on the prices and quantities of both electricity and gas: $\log(\theta_{djt}^E/\theta_{djt}^G) = \frac{1}{\sigma_{\mathcal{E}}} \log(E_{djt}/G_{djt}) - \log(P_{jt}^G/P_{jt}^E)$. This allows us to calculate normalized versions of energy, $\tilde{\mathcal{E}}_{djt}$, and the price of energy, $\tilde{P}_{dj}^{\mathcal{E}}$, where we have normalized energy by θ_{djt}^G .⁹

Next, we estimate $\sigma_{\mathcal{C}}$, the elasticity of substitution between housing and energy from the comfort production function in a similar manner. Agents produce comfort in their dwelling using housing and energy as inputs.¹⁰ Once again, we use

⁹These are $\tilde{\mathcal{E}}_{djt} = \left(\left(\theta_{djt}^E/\theta_{djt}^G \right) E_{djt}^{\rho_{\mathcal{E}}} + G_{djt}^{\rho_{\mathcal{E}}} \right)^{\frac{1}{\rho_{\mathcal{E}}}}$ and $\tilde{P}_{dj}^{\mathcal{E}} = \left(\left(\theta_{dj}^E/\theta_{dj}^G \right)^{\sigma_{\mathcal{E}}} P_j^{E^{1-\sigma_{\mathcal{E}}}} + P_j^{G^{1-\sigma_{\mathcal{E}}}} \right)^{\frac{1}{1-\sigma_{\mathcal{E}}}}$.

¹⁰We estimate the quantity of housing, H , by regressing gross rent reported in the data on a set of observable characteristics about the dwelling and a city fixed effect in logs. We interpret the city fixed effect as the cost per unit of housing, and then back out the quantity of housing by dividing gross rent by this per unit cost of housing. Appendix A.1.2.3 explains this process in detail.

the relative conditional demand functions, but now for energy and housing, that result from households choosing the cost-minimizing combination of housing and energy to produce a given amount of comfort,

$$\log \left(\frac{\tilde{\mathcal{E}}_{djt}}{H_{djt}} \right) = \sigma_c \log \left(\frac{P_{jt}^H}{\tilde{P}_{jt}^\mathcal{E}} \right) + \sigma_c \frac{\rho_c}{\rho_\mathcal{E}} \log \theta_{djt}^G. \quad (2.20)$$

Taking first differences yields our estimating equation,

$$\Delta \log \left(\frac{\tilde{\mathcal{E}}_{djt}}{H_{djt}} \right) = \sigma_c \Delta \log \left(\frac{P_{jt}^H}{\tilde{P}_{jt}^\mathcal{E}} \right) + \Delta \epsilon_{djt}. \quad (2.21)$$

Note that we do not observe amenities, and many amenities are correlated with both housing demand and rent. For example, consider a city with higher quality restaurants. These restaurants will cause increases in housing prices and housing demand, confounding our parameter estimates. Therefore, we use instrumental variables to isolate exogenous variation in the ratio of rent to energy price. Our instrument interacts labor demand shocks with land availability, where the labor demand shocks are from Katz and Murphy (1992), and the interaction with land availability follows Diamond (2016). We construct the labor demand shocks by multiplying the historical share of workers with the change in the number of workers in the rest of the country. Formally, the labor demand shock is

$$\Delta Z_{djt} = \sum_{\iota \in n} \omega_{djt}^{1990} \times (\Delta N_{d,-j,\iota,t}), \quad (2.22)$$

where ω_{djt}^{1990} is the share of workers in demographic group d from city j that worked in industry ι in 1990 as a fraction of the total number of workers across all industries in demographic group d in city j . $\Delta N_{d,-j,\iota,t}$ is the change in the number of workers in industry ι between years t and $t - 1$ outside of city j . We

interact this with the measure of housing supply elasticity from Saiz (2010), which reflects geological and bureaucratic constraints on space to build homes in a city.

The exclusion restriction is

$$E [\Delta Z_{djt} \Delta \epsilon_{djt}] = 0 \quad (2.23)$$

$$E [\text{Elasticity}_j \Delta Z_{djt} \Delta \epsilon_{djt}] = 0. \quad (2.24)$$

Any factors that effect relative energy and housing demand besides the relative prices are included in ϵ_{djt} , this includes climate and other amenities as described above. The instruments—labor demand shocks interacted with a static measure of land availability in the city—are unlikely to be correlated with changes in these other factors.

We can then use our estimate of σ_C , data on quantities and prices of housing and (normalized) energy to back out θ_{djt}^G using equation (2.20). Since we already have estimates of $\log(\theta_{djt}^E/\theta_{djt}^G)$ from the first step, we are able to calculate both θ_{djt}^G and θ_{djt}^E .

2.4.1.2 Effect of climate on comfort production. The next step of our estimation is to measure the effect of climate on the θ_{djt}^m 's, capturing how climate effects the productivity of electricity, gas, and housing. Note that while we normalize both parameters in the outer nest of the comfort production CES to be equal to one, we place no restrictions on θ_{djt}^G and θ_{djt}^E in the inner nest. Thus, θ_{djt}^G and θ_{djt}^E tell us not only the relative importance of electricity and gas, but also the importance of energy relative to housing.

We parameterize the effect of climate on these intensity parameters as:

$$\log \theta_{djt}^E = \sum_{\tau} \kappa^E(\tau) D_{\tau jt} + \psi_j^E + \psi_d^E + \nu_{djt}^E \quad (2.25)$$

$$\log \theta_{djt}^G = \sum_{\tau} \kappa^G(\tau) D_{\tau jt} + \psi_j^G + \psi_d^G + \nu_{djt}^G, \quad (2.26)$$

where $D_{\tau jt}$ is the number of days at temperature τ in city j and year t , ψ_j are city fixed effects, ψ_d are demographic group fixed effects, and the ν_{djt} 's are random disturbances. $\kappa^m(\tau)$ for $m \in \{E, G\}$ is the marginal effect of an additional day at temperature τ on electricity or gas. As we have limited data, we choose to allow for flexible non-linearity with relatively few parameters by estimating the $\kappa(\tau)$'s with cubic splines. Specifically, for $m \in \{E, G\}$, let $\kappa^m(\tau) = \sum_s \kappa_s^m S_s(\tau)$, where $S_s(\tau)$ for $s \in \{1, \dots, S\}$ are the standard basis functions of a cubic B-spline of degree S . Substituting this into the above parameterization of the θ 's gives us our two estimating equations,

$$\log \theta_{djt}^E = \sum_s \kappa_s^E \left(\sum_{\tau} S_s(\tau) D_{\tau jt} \right) + \psi_j^E + \psi_d^E + \nu_{djt}^E \quad (2.27)$$

$$\log \theta_{djt}^G = \sum_s \kappa_s^G \left(\sum_{\tau} S_s(\tau) D_{\tau jt} \right) + \psi_j^G + \psi_d^G + \nu_{djt}^G, \quad (2.28)$$

Thus, we identify the marginal impacts of climate on comfort demand using random deviations of weather outcomes from the long-run average climate within cities. For robustness, we use various degrees for the spline, as well as alternative parameterizations of climate.

2.4.2 Labor Supply. We quantify the rest of the parameters in the household utility function using maximum likelihood estimation and then linear regression. First, we define the mean utility associated with each choice as:

$$\delta_{djt} = \beta_e^I \log(I_{djt} + \Upsilon_{dt}) + \beta_e^c \log(P_{djt}^c) + \beta_e^Z \cdot Z_{jt} + \xi_{djt}, \quad (2.29)$$

where $\beta_e^I = \frac{1}{\sigma_e}$, $\beta_e^c = -\frac{\alpha_e^c}{\sigma_e}$ and $\beta_e^Z = \frac{\alpha_e^Z}{\sigma_e}$. We then write the household's choice probability as:

$$P_{ijt}(\delta_{djt}; \gamma_{dt}) = \frac{\exp(\delta_{djt} + \tilde{\gamma}_{dt}^{st} \mathbb{I}(j \in b_i^{st}) + \tilde{\gamma}_{dt}^{\text{dist}} \phi(j, b_i^{st}) + \tilde{\gamma}_{dt}^{\text{dist}2} \phi^2(j, b_i^{st}))}{\sum_{j'} \exp(\delta_{dj't} + \tilde{\gamma}_{dt}^{st} \mathbb{I}(j' \in b_i^{st}) + \tilde{\gamma}_{dt}^{\text{dist}} \phi(j', b_i^{st}) + \tilde{\gamma}_{dt}^{\text{dist}2} \phi^2(j', b_i^{st}))}, \quad (2.30)$$

where the tildes over the parameters indicate they are normalized by the variance of the idiosyncratic preference shock, σ_e . Given these choice probabilities, the likelihood function is:

$$LL(\delta_{djt}; \gamma_{dt}) = \sum_{i \in d} \sum_j \mathbb{I}_{ijdt} \log(P_{ijt}(\delta_{djt}; \gamma_{dt})), \quad (2.31)$$

where $i \in D$ is the set of all households in demographic group d and $\mathbb{I}_{ijdt}$ is an indicator equal to one if the household chose city j in the data. We recover the δ_{djt} 's conditional on γ_{dt} using the contraction mapping proposed in Berry et al. (2004).

We then decompose δ_{djt} from step one to estimate climate amenities. First, we use estimates from Diamond (2016) to calibrate β_e^w to 5.22 for college and 4.15 for non-college. Additionally, we use the observed comfort expenditure shares for each city, demographic group, and year to calibrate β_{djt}^c . With these parameters, we calculate the residual deltas as $\hat{\delta}_{djt} = \delta_{djt} - (\beta_e^I \log(I_{djt}) + \beta_{djt}^c \log(P_{djt}^c))$. We then estimate the effects of climate on the mean-utilities using a rich specification of the climate that includes heating and cooling degree days, total precipitation,

the number of days without precipitation, and measures of fire and flood risk. Our estimating equation is

$$\hat{\delta}_{djt} = \beta^Z \cdot \mathbf{Z}_{jt} + \beta^x \cdot \mathbf{X}_{jt} + \omega_d + \omega_t + \xi_{djt}, \quad (2.32)$$

where $\mathbf{Z}_{jt} = \{\text{Fire}_j, \text{Flood}_j, \text{CDD}_{jt}, \text{HDD}_{jt}, \text{Precip}_{jt}, \text{NoRain}_{jt}\}$ is a vector of climate variables. Fire_j and Flood_j are median fire and flood risk scores in city j that are static over time. CDD_{jt} and HDD_{jt} are heating and cooling degree days. Precip_{jt} is total precipitation and NoRain_{jt} is the proportion of days with no precipitation. $\mathbf{X}_{jt}$ is a vector of geographic and demographic controls including average elevation, average slope, percent of population married, average age, and percent of population Black in each city and year. ω_d and ω_t are demographic group and year fixed effects. This is a parsimonious, but rich characterization of the climate that allows us to capture the nuanced heterogeneity in the effects of climate change.

2.5 Parameter Estimates

Comfort Demand. Here, we report our parameter estimation for comfort demand. First, Table 1 shows the results from estimating Equations (2.18) and (2.21) with instrumental variables, where the columns of the table have results for alternative fixed effects included in the regressions. While the results are presented together for concision, they were run as two separate regressions. Our preferred specification is in column 1, with demographic group and city fixed effects.

We estimate that electricity and gas are net complements, as we estimate the elasticity of substitution, σ_E , to be less than one. Intuitively, a house that requires more electricity for cooling also requires more gas to heat. The instrument for the log ratio of electricity and gas prices, which is the log ratio of lagged

Model:	(1)	(2)	(3)	(4)	(5)
<i>Parameters</i>					
$\sigma_{\mathcal{E}}$	0.501*** (0.063)	0.525*** (0.066)	0.439*** (0.070)	0.432*** (0.086)	0.432*** (0.087)
$\sigma_{\mathcal{C}}$	0.594*** (0.079)	0.463*** (0.091)	0.316*** (0.092)	0.293 (0.348)	0.570 (0.436)
<i>Fixed-effects</i>					
Dem Group	Yes	Yes			Yes
City	Yes		Yes		
Year				Yes	Yes
<i>Fit statistics</i>					
Observations	1,260	1,260	1,260	1,260	1,260
F-test (1st stage), $\log(P_g/P_e)$	378.82	363.51	392.65	391.88	391.25
F-test (1st stage), $\log(P_h/P_{\mathcal{E}})$	147.47	65.63	143.79	1.78	0.90

Clustered (City) standard-errors in parentheses


*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The elasticity of substitution between electricity and gas is $\sigma_{\mathcal{E}}$ and the elasticity of substitution between energy and housing is $\sigma_{\mathcal{C}}$. Coefficients are shown in the same column, but come from separate regressions. Regressions are weighted by 1990 population. We exclude the division aggregation CBSA's in estimation.

Table 1. Elasticity of substitution estimates.

commercial electricity and gas prices, is strong, with an F-stat of well over 300 in our preferred specification. Housing and energy are net complements as well. Suppose the price of housing goes down relative to energy prices. In that case, the housing quantity demanded will increase, but so will the energy demanded since more energy is required to keep that increased amount of housing comfortable. Here, we instrument the log ratio of rent to energy price with the same instrument from energy estimation, and the Katz-Murphy index interacted with housing supply elasticities from Saiz (2010). Again, our first stage F-stat is very strong.

With the estimated elasticities of substitution in hand, we can calculate values for $\log \theta_{djt}^E$ and $\log \theta_{djt}^G$, which we then regress on climate variables to determine the impact of the climate on the comfort production function. We use cubic B-splines with 7 degrees of freedom to flexibly estimate the marginal effect of an additional day at a given temperature. For electricity, an additional day at either a hotter or colder temperature than 65 degrees leads to additional energy demand, reflecting that households use electricity for both heating and cooling. Additionally, the demand response is stronger further into temperature extremes. Meanwhile, natural gas has coefficients significantly different from zero only when the temperature is colder than 65 degrees. We investigate heterogeneity in the climate’s effect on comfort production in appendix A.4.1, finding differences that are small in magnitude. We then use the above estimates to calculate the marginal effect of climate on comfort prices, finding that the effect is larger for Black households than white—though not outside of bootstrapped confidence intervals. Appendix A.4.1 provides other figures describing patterns in comfort production across cities.



(a) Estimate of $\kappa_E(\tau)$, the marginal effect of temperature on $\log \theta^E$.

(b) Estimate of $\kappa_G(\tau)$, the marginal effect of temperature on $\log \theta^G$.

Figure 4. Estimation results for the effect of climate on comfort production

Estimating parameters $\kappa^E(\tau)$ and $\kappa^G(\tau)$ using cubic splines with 7 degrees of freedom. The effect is normalized to a day with an average temperature of 65 degrees F. Standard errors are calculated with clustered bootstrapping using 10,000 draws, clustered by city-year. We have limited the range of temperature to between the 1st and 99th percentile of average daily temperature, weighted by 1990 population.

Labor Supply. Figure 5 plots the estimated moving cost function (2.3) using parameters for estimated by equation (2.31). We plot the moving cost estimates over the range of possible distances from a household's birthstate.¹¹ We find that Black households are generally less mobile than other households. The birth-state premium parameter is positive, indicating a utility premium, and larger for Black households. Furthermore, the distance parameter is negative, indicating a utility penalty for moving further away from one's birth state, and more negative for Black households. We also find that non-college-educated workers are generally less mobile than college-educated workers. Our estimates are quantitatively very close to those in Colas and Morehouse (2022) and Diamond (2016). We include moving cost estimates for all years in Appendix A.4.4.

¹¹In both Figure 5 and in our simulations, we set the moving cost function equal to its minimum value for distances beyond its estimated minimum.




Figure 5. Moving cost estimates for Black and White households

We use estimates from Table A.9, but hold moving costs constant at the minimum for distances further than the estimated minimum. The intercept reflects the parameter estimates of $\tilde{\gamma}^{st}$. The histogram shows the distribution of *potential* moves in 2019, excluding the census divisions—so each household in the model shows up 70 times, once for each of the cities we include in the model. We set moving costs equal to the minimum value of the moving costs function at distances further beyond its estimated minimum.

Table 2 shows the parameter estimates for climate amenities. We show the robustness of our controls to the inclusion of different geographic and demographic controls. Households dislike both heating and cooling degree days and like days without rain. These estimates are larger in magnitude than Albouy et al. (2016), but smaller than Rudik et al. (2021). Households dislike both fire and flood risk—they are willing to pay 24% of income for a one standard deviation decrease in fire risk and 6% of income for a one standard deviation decrease in flood risk. We emphasize that we do not presently have natural disaster risk indices that vary over time or climate scenarios. Appendix A.4.5 contains estimates for alternative specifications for temperature portion of climate amenities, showing qualitatively similar results to our degree day specification.

Table 3 shows the top and bottom ten cities according to their climate amenities in 2019 ($\hat{\beta}^Z \cdot \mathbf{Z}_{j,2019}$), with category-specific rankings also broken out. These rankings demonstrate our model’s ability to richly characterize a city’s climate. For example, the top ten has cities from California with mild summers, winters, and risks, as well as San Antonio—which ranks poorly due to its hot summer weather but makes up for it by having very few cold days. Most striking is the inclusion of Baton Rouge in the top ten and nearby New Orleans in the bottom ten. This is due to the difference in flood risk between the two cities—nearly every property in New Orleans has at least moderate flood risk, whereas less than half do in Baton Rouge. Table 4 shows how climate amenities have changed between 1990 and 2019, using 5-year moving averages. Increased hot temperatures have decreased amenities in most cities, while fewer cold days have improved amenities. In already-hot places, the shape of the temperature-amenity curve means that the decrease in

Table 2. Estimation results for non-temperature climate amenities.

Dep. Var. Model:	Mean Utility Residual			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
HDD (1000's)	-1.73*** (0.14)	-1.62*** (0.15)	-1.60*** (0.15)	-1.99*** (0.23)
CDD (1000's)	-1.54*** (0.28)	-1.44*** (0.30)	-1.03*** (0.28)	-1.63*** (0.43)
Median Fire Risk	-0.75*** (0.13)	-0.60*** (0.14)	-0.55*** (0.10)	-0.62*** (0.11)
Median Flood Risk	-0.83*** (0.15)	-0.94*** (0.14)	-0.68*** (0.11)	-0.75*** (0.11)
Annual Precip (m)	-2.81* (1.69)	-2.30 (1.61)	-2.03 (1.59)	-1.17 (1.75)
Annual Precip Squared (sq m)	1.55* (0.81)	1.42* (0.79)	0.75 (0.73)	0.53 (0.77)
Pr(No Rain)	7.00** (2.99)	7.06** (3.03)	7.10** (2.86)	7.18** (2.86)
Geographic Controls		Yes		Yes
Demographic Controls			Yes	Yes
<i>Fixed-effects</i>				
Dem Group	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	1,680	1,680	1,680	1,680
R ²	0.94	0.94	0.95	0.95
Within R ²	0.27	0.29	0.33	0.34

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Regressions are weighted by population. We exclude the division aggregation CBSA's in estimation. Geographic controls include city's average slope and elevation. Demographic controls include percent married, average age, and percent black.

amenities from more hot days outweighs the improvement in amenities from fewer cold days. The opposite generally holds for places that are relatively cold.

Table 3. Estimated climate amenity rankings by city for 2019.

The Hot category is from temperatures over 65F, Cold from temperatures below 65F, Rain from percent no rain days and annual precipitation, and Risks from median fire and flood risk. Full sample degree day specification with geographic and demographic controls.

Rank	City	Category Rankings			
		Hot	Cold	Rain	Risks
Best Cities					
1	San Jose-Sunnyvale-Santa Clara, CA	5	18	7	32
2	Los Angeles-Long Beach-Santa Ana, CA	40	7	2	54
3	San Diego-Carlsbad-San Marcos, CA	30	10	5	65
4	San Francisco-Oakland-Fremont, CA	2	21	10	59
5	Oxnard-Thousand Oaks-Ventura, CA	18	15	1	67
6	Fresno, CA	53	17	6	53
7	Sacramento-Arden-Arcade-Roseville, CA	38	22	9	57
8	Baton Rouge, LA	59	13	61	4
9	San Antonio, TX	61	12	13	58
10	Phoenix-Mesa-Scottsdale, AZ	69	6	3	64
Worst Cities					
61	Grand Rapids-Wyoming, MI	7	67	42	14
62	Buffalo-Cheektowaga-Tonawanda, NY	10	64	58	7
63	Scranton-Wilkes-Barre, PA	12	59	51	43
64	Springfield, MA	16	62	25	44
65	Worcester, MA	8	63	32	48
66	Albany-Schenectady-Troy, NY	14	66	41	37
67	Rochester, NY	9	65	56	42
68	Minneapolis-St. Paul-Bloomington, MN-WI	11	70	20	22
69	New Orleans-Metairie-Kenner, LA	65	8	60	70
70	Syracuse, NY	4	68	67	45


Table 4. Change in climate amenities from 1990 to 2019.

Full sample degree day specification with geographic and demographic controls.

City	Change (%)	Category Change (%)			
		Rain	Hot	Cold	
Improving Cities					
Hartford-West Hartford-East Hartford, CT	26.2	7.1	-22.4	41.5	
Grand Rapids-Wyoming, MI	25.7	0.3	-2.7	28.1	
Honolulu, HI	23.8	17.8	6.0	-0.1	
New Haven-Milford, CT	23.1	5.5	-17.7	35.3	
Denver-Aurora, CO	21.6	-12.2	-16.0	49.8	
San Antonio, TX	19.8	15.3	-14.2	18.7	
Worcester, MA	19.2	-11.9	-15.5	46.7	
Portland-Vancouver-Beaverton, OR-WA	18.9	-3.0	-11.7	33.6	
Buffalo-Cheektowaga-Tonawanda, NY	18.2	10.1	-10.1	18.3	
Bridgeport-Stamford-Norwalk, CT	17.1	1.3	-17.3	33.2	
Worsening Cities					
St. Louis, MO-IL	-27.6	-34.2	-9.4	16.0	
Dayton, OH	-28.3	-40.6	-11.2	23.4	
Houston-Baytown-Sugar Land, TX	-31.1	-27.2	-35.0	31.1	
Louisville, KY-IN	-32.8	-47.8	-21.4	36.4	
Baton Rouge, LA	-35.9	-30.5	-35.5	30.1	
Tampa-St. Petersburg-Clearwater, FL	-62.8	-44.6	-28.6	10.4	
Jacksonville, FL	-68.6	-55.3	-39.9	26.5	
Orlando, FL	-70.4	-41.9	-42.7	14.2	
Miami-Fort Lauderdale-Miami Beach, FL	-73.0	-40.3	-39.8	7.1	
New Orleans-Metairie-Kenner, LA	-73.4	-54.4	-54.8	35.8	

Model Fit We use our estimated parameters to simulate baseline results for each year. Appendix A.3.7 provides detailed steps to solve the model’s equilibrium given climate $\tilde{\mathbf{Z}}$. In summary, we begin with population share and housing price guesses, use those guesses to calculate household energy and housing demand, wages, and comfort prices, and then recalculate population shares given those wages and comfort prices. If the population shares match the original guess,


then we have found the equilibrium. If they do not match, we update our initial guess and repeat until they do match. In the baseline simulations, we plug in the actual climate values as $\tilde{\mathbf{Z}}$ to validate the model's predictions against actual data from each of our four years of data. Figure 6 shows the simulated shares of each demographic group that choose to live in a particular city against the actual population shares and simulated housing, electricity, and gas demand relative to actual electricity, housing, and gas demand. Each observation represents a demographic group in a city for the given year.




(a) Simulated vs actual population shares.



(b) Simulated vs actual housing demand.



(c) Simulated vs actual electricity demand.



(d) Simulated vs actual natural gas demand.

Figure 6. Simulations use actual climate data across all years.

The size of the observations represents the population of that demographic group.

2.6 The welfare effects of climate change

We calculate the welfare effects from climate change using compensating variation (CV), defined as the percent of baseline income that households would need to receive in order to maintain the same level of utility under the

counterfactual climate, $\tilde{\mathbf{Z}}$, as they do under the current climate, $\mathbf{Z}$.¹² Household i 's compensating variation is given by:

$$CV_i = \left(\mathbb{E}[V_i(\tilde{\mathbf{Z}})] - \mathbb{E}[V_i(\mathbf{Z})] \right) \times \frac{1}{\beta_e^I}, \quad (2.33)$$

where $V_i(\mathbf{Z}) = v_{ij^*}(\mathbf{Z})$ is household i 's indirect utility with climate $\mathbf{Z}$, evaluated at equilibrium choices j^* and prices. We calculate these equilibrium choices using the same algorithm referenced in the model fit section, which is detailed in Appendix A.3.7. Taking the expectation of indirect utility over the idiosyncratic shocks gives us

$$\mathbb{E}[V_i(\mathbf{Z})] = \bar{\gamma} + \log \left(\sum_{j' \in J} \exp \left(\beta_e^I \log I_{dj'} + \beta_{dj'}^C \log P_d^C(Z_{j'}) + \beta_d^Z \cdot Z_{j'} + \xi_{dj'} + g(j', \mathbf{b}_i) \right) \right),$$

where $\bar{\gamma}$ is Euler's constant. Note that the price of comfort, $P_d^C(Z_j)$, is a function of local climate variables.

2.6.1 Effect of climate change to-date. We begin with a simulation where counterfactual climate comes from the average climate between 1986 and 1990 before there were any significant human-caused changes to the climate. We first analyze how the effects of climate change vary across space. Figure 7a shows compensating variation by state—showing that already hot places are hurt by more than cool places. Generally, the South and Midwest are worse off than the Northeast and West, owing to increases in temperature being unpleasant in an already hot location and pleasant in a cool location. Additionally, fewer rain-free days in the Midwest have decreased amenities. The effect of climate on the energy price drives some of this heterogeneity, shown in Figure 7b. Producing

¹²In all counterfactuals, we hold natural disaster risk constant, as we only have a single set of natural disaster risk scores.

comfort has become more expensive in hot places and less expensive in cold places. Figure 7c shows that households migrate in away from parts of the South and Midwest and into the Northeast and West.¹³ These spatial changes suggest that the baseline distribution of the population is essential in determining differences in welfare effects across demographic groups.

There are significant differences in the welfare effect of climate change across the income distribution. We calculate the average compensating variation by income decile and race in Figure 8. The welfare effects are worse for lower-income households among all races. CV as a percent of income is twice as high for the lowest income decile relative to the highest—the lowest income decile has welfare losses of 4%pts of income, whereas the highest income decile is only worse off by 2%pts of income. The welfare effects for Black households are worse for Black households relative to white households within every income decile.¹⁴ The mean welfare loss for Black households is twenty percent larger than that of white households—2.5%pts of income for white households relative to 3.0%pts for Black households.

We decompose the proposed mechanisms driving the observed gaps in CV between white and Black households. The results of this decomposition are in Table 5.¹⁵ We run two simulations in addition to the fully flexible model to do this. First, we fix locations and wages—reflecting the “Mechanical” effect of changing the climate without allowing households to migrate. In this simulation, households still optimally choose their comfort demand. Therefore, we allow

¹³Table A.11 has the top and bottom 10 cities by change in population. Maps of rent, wages, and comfort demand are available in Figure A.16. Additionally, Figure A.18 shows how the skill ratio affects wages in the counterfactual.

¹⁴Appendix Figure A.17 shows the full distribution of CV_i across all households.

¹⁵Table A.10 has these results in levels rather than differences between races.




Figure 7. General equilibrium changes in welfare, energy prices, and population across states

We compare the 2019 climate to the 1980s climate. Averages are weighted by population. For states without any CBSAs in them, we use the value for the census division in which the state is located.




Figure 8. Compensating Variation by income decile and race in 2019

Simulations run using 2019 economy, comparing 2019 and 1990 climate. Income deciles are calculated relative to the entire population.

local housing markets to clear since housing demand responds to climate as part of comfort production. Column (1) of the first row of Table 5 shows that the average gap in CV between Black and white households is -0.8%pts, reflecting both disparities in the sensitivity in comfort price and in exposure. We can compare *within* baseline cities to control for differences in exposure. Column (2) shows that after conditioning on baseline city, there is no welfare gap between Black and white households. Thus, differences in sensitivity to comfort prices are not driving the gap.

In a second simulation, we allow households to sort across locations, fix wages to their original level, and keep rent at the same level as in the Mechanical

simulation. This “Sorting” simulation tells us the effect of allowing households to reoptimize their location choice across space. As seen in the second row of Columnn (1) Table 5, allowing for sorting marginally increases the unconditional gap. However, two-thirds of the gap remains after controlling for baseline location. Thus, white households can better take advantage of migration as an adaptation mechanism conditional on exposure. These differences are not explained by differences in the ratio of college- to not-college-educated households within the race groups. Column (3) shows the CV gap within baseline city choice and education level, only reducing the gap from the Sorting simulation by 0.1%pts.

Finally, the third row of Table 5 shows the effects of adjustments to rents and wages in equilibrium. The unconditional welfare gap shrinks—households migrating from the South to the West or Northeast depress wages and inflate rents in their new cities, while the opposite happens in the South. These changes reduce the unconditional gap due to the differences in the baseline populations in those regions. The within-city gap remains, but education level explains more of that gap than in the Sorting simulation. Since the ratio of college- to non-college educated workers determines wages in each city, greater out-migration of college educated workers from high-climate-damage cities actually decrease wages for non-college educated households in those cities. This is true despite the overall population decreasing.

In summary, we find considerable gaps in the welfare effects of climate change to date across incomes and races. These gaps are explained by both the ex-ante population distribution and differences in mobility. Differences in energy efficiency do not seem to be a major contributor.

Simulation	Black-white CV Gap (%pt of income)		
	Unconditional	Within City	Within City-Education
	(1)	(2)	(3)
Mechanical	-0.8	0.0	0.0
Sorting	-0.9	-0.6	-0.5
Full Effects	-0.5	-0.3	-0.1

Table 5. Decomposition of welfare effects.

Column (1) calculates the average difference in CV between Black and white households from the climate change to date simulations. Column (2) calculates the average welfare gaps after taking out the average effect for each city (equivalent to adding a city fixed effect). Column (3) calculates the average welfare gaps within city and education level (college and non-college). The “Mechanical” simulation fixes locations and wages. Local housing markets are allowed to clear in response to changes in housing demand from change in climate. The “Sorting” simulation keeps wages fixed and rents the same as the Mechanical simulation, but allows households to reoptimize across space. “Full Effects” is the fully flexible general equilibrium model.

2.6.2 The black-white gap simulated under future emissions

scenarios. Next, we simulate the model under different emissions scenarios to 2100 using the CIL-GDPCIR data. We then calculate the Black-white gap as the change in welfare for Black households relative to their baseline welfare minus the change in welfare for white households relative to their baseline welfare. Figure 9 shows the evolution of the Black-white welfare gap, which generally grows with more emissions.

Under the lowest emissions scenario (SSP1-2.6), which assumes zero emissions by 2050 and 1.8 degrees of warming by the end of the century, climate change is expected to hurt Black households by an additional 0.5%pts of income relative to white households. Under the more likely, higher emissions scenarios, this gap grows to between 2%pts to 4%pts depending on the emissions scenario. This

gap highlights the need for equity to be an important part of climate policy in the future.




Figure 9. Time series of the Black-white welfare gap under different emissions scenarios.

We simulate general equilibrium in the model using future climate scenarios with model parameters set to their 2019 values. We then calculate CV for future years relative to the 2020 climate. The CV gap is the difference in CV between Black and white households.

2.7 IRA inspired place-based, means-tested subsidies

Given the unequal impacts of climate change, we now study how government policy might be used to address these inequalities. We incorporate a means-tested, place-based subsidy into our main model, the structure and magnitude of which are inspired by a new Inflation Reduction Act (IRA) program administered by the US EPA. We compare the effects of climate change to date with and without these subsidies and test alternative distributions of these funds across space.

The IRA contains many programs designed to address climate change. One such program is the US EPA’s Community Change Grant Program, for which the IRA allocated \$3 billion to fund projects in disadvantaged communities related

to climate resiliency and adaptation, mitigating climate and health risks, and several other similar categories (US EPA, 2023). In describing the program, Brenda Mallory, the Chair for the White House Council on Environmental Quality, said, “As part of the President’s Justice40 Initiative, these grants will help disadvantaged communities tackle environmental and climate justice challenges they face by reducing pollution, increasing resilience to impacts from climate change, and building community capacity to see these projects through” (US EPA, 2024).¹⁶ Applications for the primary round of funding are due in November of 2024.¹⁷ There are few restrictions on how the EPA distributes Community Change Grant funds across space. In principle, any disadvantaged community is eligible for a project.¹⁸ We do not model the effects of these projects directly, both because the EPA has not yet awarded most of the funds and because the projects themselves will vary considerably in their structure and goals. However, we are able to capture

¹⁶President Biden’s Justice40 initiative seeks to provide 40 percent of the benefits of federal rulemaking on climate, clean energy, affordable housing, and other related categories to disadvantaged communities (Executive Office of the President, 2021). The administration defines a disadvantaged community as places “historically marginalized and overburdened by pollution and underinvestment in housing, transportation, water and wastewater infrastructure, and health care” Executive Office of the President (2021). In practice, the Climate and Economic Justice Screening Tool classifies disadvantaged tracts based on eight different categories: climate, energy, health, housing, legacy pollution, transportation, or workforce development. For all but workforce development, communities are disadvantaged if they are above the 65th percentile for the share of low-income households and above a threshold for exposure to various environmental stressors (e.g., above the 90th percentile expected agricultural, building, or population loss, projected flood, or fire risk for the climate change category).

¹⁷An initial round of \$325 million in funding was awarded in July of 2024. While the specific projects take on a wide array of activities, the most common types are providing weatherization assistance, creating “community resilience hubs”, improving public parks, and installing solar panels and energy storage.

¹⁸Appendix Figure A.19 shows the proportion of the population living in a disadvantaged community by state. While there is some variation, all states have a significant portion of their population living in disadvantaged communities. Several target areas place some spatial restrictions on funds: \$150 million for tribes in Alaska, \$50 million for territories, and \$100 million for southern border communities.

two important aspects of the program, which is that these projects are *place based* and only disadvantaged communities are eligible.

Inspired by the Community Change grants, we incorporate means-tested, place-based subsidies into the model. The government decides the total size of the program, G , the portion of the total given to each city, wt_j , where $\sum_j wt_j = 1$, and the income threshold for eligibility, I_{thresh} . Households who choose to live in city j and are below the income threshold, I_{thresh} , divide the subsidy equally. $N_j^{disad} = \sum_i P_{ij} \mathbb{I}(I_{dj} < I_{thresh}) di$ gives the total population in city j that below the income threshold, where P_{ij} is household i 's probability of choosing city j , as defined by Equation (2.7). We set the income threshold equal to the 41st percentile of income, matching the share of the population living in a disadvantaged community according to the 2020 Decennial Census. Subsidies for each city and demographic group are $IRA_{dj} = \frac{Gwt_j}{N_j^{disad}} \mathbb{I}(I_{dj} < I_{thresh})$. These subsidies are given to eligible households lump-sum. A flat tax paid by all households funds the subsidy program, $G = \sum_i \tau I_{dj} di$. Income in city j for demographic group d is now $(1 - \tau)I_{dj} + \Upsilon_d + IRA_{dj}$.

Our modeled policy deviates from the actual Community Change grant program in several important ways. First, the actual program funds projects in disadvantaged communities, while the modeled policy directly transfers cash to low-income households. It is unlikely that households, if given the cash transfer, would use that cash on the same project funded by Community Change grant program. Thus, we present results that are normalized by the size of the program, G , and run separate simulations of the model under various G 's. This is consistent with treating the marginal value of public funds as a constant different from one. Second, we do not know which projects will actually be funded, nor where the

funding will go. Around 10% of the programs funding has already been awarded, while applications for the remaining funds are due in November of 2024. Thus, our results do not evaluate the projects or the program itself, but rather a hypothetical policy capturing most of the important features of the Community Change grant program.

We consider three policy alternatives that vary the spatial distribution of the subsidies across cities, wt_j . These distributions give subsidies to cities based on (1) the percent of the national disadvantaged population living in each city, (2) the share of losses in climate amenities for cities that experience losses in amenities, or (3) the share of gains in climate amenities for cities that experience gains in amenities. Distributing funds equally based on the disadvantaged population most closely mirrors the current program.

We design the second and third distributions to capture the tradeoff between helping people with high climate damages and the general equilibrium consequences of place-based subsidies (see e.g., Busso et al. (2013), Kline and Moretti (2014)). If households are immobile, then giving subsidies to places with high climate damages will also benefit people who have the largest welfare losses from climate change. However, since households are mobile, these subsidies induce some marginal households to migrate into these high climate damage cities—increasing rents and decreasing wages, thus dampening the intended welfare effects while also increasing exposure to future climate change. Therefore, we also consider giving subsidies to places that have benefited from climate change, which we’ll refer to as “climate havens.” Distributing funds to climate havens will decrease exposure to climate change but also gives subsidies to many “inframarginal” households who would

have lived in a climate haven regardless of the additional subsidies and thus are already better off from climate change relative to households from other locations.

We simulate the full model, allowing households to sort across space and for wages and rents to respond to these new choices. We consider the equilibrium with the 2019 climate and no subsidies as the baseline and then quantify changes from that equilibrium when adding the three different distributions of subsidies described above. Since the magnitude of the effects will depend on the total size of the program, we present results per dollar of total subsidies. Appendix A.5.1 compares results across a wide range of values for G , from \$30 million to \$10 billion.

Additionally, we calculate what we call “mechanical CV,” or the subsidies net of taxes that each household would expect to receive if we hold locations, wages, and rents fixed to equilibrium choices under the 2019 climate without subsidies. This allows us to determine how important migration and general equilibrium effects are relative to the mechanical effect of introducing subsidies. Each individual’s mechanical CV is the weighted average of subsidies net of taxes across cities expressed as a percent of income, weighted by the probability they choose each city in the 2019 climate equilibrium, denoted P_{ij}^* ,

$$CV_{idj}^{Mech} = \sum_j P_{ij}^* \frac{\text{IRA}_{dj}(P_{ij}^*) - \tau I_{dj}}{I_{dj}}$$

Figure 10 shows compensating variation from the introduction of subsidies relative to the baseline 2019 climate simulation. We first order all households according to the welfare effects from climate change to date, assigning each household to a baseline CV decile on the X-axis. Households worst off from climate change to date are in the first decile on the left, while those relatively well off are in the tenth decile on the right. We then break households up into three groups

according to their baseline income percentile—those below the 10th percentile, those between the 10th and 50th percentile, and those above the 50th percentile, corresponding to the three panels of the figure.¹⁹ We plot the average CV in each group per dollar of total government spending from both the fully flexible model (solid lines) and mechanical CV (dashed lines), where we fix locations, rents, and wages to baseline levels.

Disadvantaged Population The green lines in Figure 10 show results when we distribute funds according to cities based on their share of the national disadvantaged population. This distribution is most similar to the current program, given we do not yet know how the EPA will distribute funds across space.²⁰ This distribution helps low-income households generally but not proportionately to welfare losses from climate change. Households in the bottom income decile are better off by 0.22%pts of income per billion dollars spent, while households above median income are marginally worse off from paying the tax. Additionally, these subsidies do not alter the spatial distribution of households.²¹ While there is some variation in the share of the population in each city that lives in a disadvantaged community (See Figure A.19), this variation is not large enough to generate differences in subsidy levels that alter many household’s decision on where to live.

¹⁹Figure A.20 shows the distribution of households across baseline CV for each of these baseline income groups. We calculate these household-level income percentiles by taking the average income the household would earn in each city, weighted by the probability that the household chooses to live there.

²⁰The final distribution of funds may not be proportional to the disadvantaged population if there are differences across space in the likelihood of a group representing a disadvantaged community applying for or receiving funding.

²¹Appendix Figure A.23 shows changes in population caused by the subsidies.




Figure 10. CV from subsidies relative to CV from climate change to date for different spatial distributions of IRA subsidies.

The X-axis is a household’s decile in CV from the baseline simulation comparing 2019 and 1990 climates. The Y-axis is the average CV per dollar of total government spending with subsidies relative to without subsidies under the 2019 climate. We split the figure into three income bins, below the 10th percentile, 10th-50th percentile, and above the 50th percentile, based on household expected income in the baseline simulation. The “Mechanical” simulation takes choice probabilities from the baseline simulation and calculates expected subsidies given those choices—which is what would happen if households were not allowed to reoptimize their location choices after the government introduces subsidies. The “Fully flexible” simulations allow for full flexibility in choosing locations. See description in the text on the subsidy distributions

Thus, there is little difference in CV between the mechanical and fully flexible simulations.

Negative Climate Change The orange lines in Figure 10 show results when we distribute funds proportional to each city’s losses in climate amenities in 2019 relative to 1990. This distribution does the best job of helping households with the most significant welfare losses from climate change—it is the only distribution where CV from subsidies is decreasing in the baseline CV percentile.²² Households in the lowest income decile and lowest baseline CV decile are better off by 0.5%pts per billion spent. In contrast, those in the highest baseline CV decile are unaffected by the subsidies.

However, the large spatial differentials in subsidies lead some households to migrate into the high climate damage cities, leading to increases in rents and decreases in non-college-educated wages that dissipate a portion of the benefits to those most hurt by climate change in the baseline. The overall benefits to households in the lowest income decile are 19% lower in the fully flexible simulation than under the mechanical simulation.²³ This reallocation has important environmental justice implications. The orange lines in Figure 11 show the change in average CDD, HDD, fire risk, and flood risk for households in the lowest income decile. Generally, low-income households move to hotter and more flood-prone areas under the Negative Climate Change distribution. Since there is effectively

²²Appendix Figure A.21 shows the CV relative to the 1990 climate, along with the baseline CV. The effect of subsidies with a \$3 billion program is small relative to the overall effect of climate change to date.

²³Visually, this is the difference in area under the dashed and solid lines in Figure 10, weighted by the distribution of households in Figure A.20. This number is sensitive to the total magnitude of the program: with a \$300 million program, 16% is lost, and with a \$10 billion program, 26% is lost.

no reallocation of high-income households across space, these changes mean that this distribution of subsidies would increase the ex-ante exposure gap we found to drive the baseline welfare effects from climate change.

Positive Climate Change The purple lines in Figure 10 show results when we distribute the funds to places that have benefitted from climate change to date. This distribution is increasing in the baseline CV percentile, primarily benefitting households who are already relatively well off from climate change. For example, the average household in the lowest income decile and highest baseline CV decile has effectively zero welfare effects from climate change to date, relative to low-income households in the first baseline CV decile whose welfare losses are equivalent to 7%pts of income (See Figure A.21 for CV relative to the 1990 climate). These same households are the largest benefactors of the Positive Climate Change distribution, with CV of 1%pt of income. Many of these households are “inframarginal”—they would have chosen to live in these climate havens regardless of the subsidies. Well over 90% of subsidies in the Positive Climate Change distribution are given to households who would have lived in a city with increases in climate amenities absent the subsidies.

There is a 16% welfare loss among households in the lowest income decile due to migration and its resulting general equilibrium effects in the Positive Climate Change distribution.²⁴ Since this simulation distributes subsidies to “climate havens”, the observed in-migration decreases exposure to climate change. The purple lines in Figure 11 show that low-income households generally move to cooler and less flood-prone cities. These results emphasize that a program designed

²⁴This is the difference in area between the dashed and dotted lines in Figure 10. With a \$300 million program, 1% is lost, and \$10 billion program, 30% is lost.




Figure 11. Change in climate with subsidies relative to without subsidies for households in the lowest income decile.

The X-axis is a household's decile in CV from the baseline simulation comparing 2019 and 1990 climates relative to the entire population. We have filtered to just households below the 10th income percentile in the baseline simulation. We calculate the values on the Y-axis as a percent change in the respective climate variable relative to baseline values for each group, normalized by the total amount of subsidies dispersed. See description in the text on the subsidy distributions

to reduce climate exposure among low-income households must target marginal households in order to be successful. Place-based policies are not well suited to achieve this goal.

In summary, we find that place-based, means-tested subsidies distributed to places with high climate damages may be an effective tool to help households with the largest impacts from climate change. However, because households are mobile, some marginal households move into high-climate-exposure cities, both exacerbating the existing disparities in climate exposure and dissipating a portion of the program's benefits, such as deadweight loss. Appendix A.5.1 demonstrates that these results qualitatively hold across a wide range of total program magnitudes. Since the policies we analyze redistribute funds from high- to low-income households, they all reduce the aggregate disparities in welfare effects of climate change across race and income that we quantified in our main results. The Negative Climate Change distribution is the most effective of the three at reducing the welfare gaps.

2.8 Conclusion

We have found that climate change has already had and will continue to have significant distributional consequences. Black and non-college-educated households are worse off under climate change, owing not only to their ex-ante exposure to changes in climate but also due to a differential ability to adapt. Black households are worse off by 0.5% of income relative to white households from climate change to date, and that gap continues to grow under various future emissions scenarios. In the worst-case emissions scenario, the gap grows by over 4%pts of income by the end of the century. Low-income households are much worse off than high-income households, with the welfare effect of climate change to date

being two times larger for the lowest-income decile relative to the highest-income decile.

We evaluated the effectiveness of a place-based, means-tested subsidy program to address the observed inequalities in the effects of climate change. We found that there is a fundamental tradeoff when using such place-based policies between targeting the places that have the largest climate damages and increasing exposure to future climate damages for low-income households. Policymakers should weigh these tradeoffs when designing similar place-based policies.

Future work should address other forms of inequality that may arise due to climate change. For example, the urban heat island effect may lead to differential experienced climates within cities that our model is not well suited to analyze. One could allow for households to choose both their city and neighborhood in that city to directly model any within-city sorting.

CHAPTER III

GLYPHOSATE EXPOSURE AND GM SEED ROLLOUT UNEQUALLY REDUCED PERINATAL HEALTH

This chapter is joint work with Assistant Professor Edward Rubin published in the *Proceedings of the National Academy of Sciences*. The idea for this paper originated as my final project for Ed’s first-year PhD Econometrics course. We then continued working on the project together. I contributed significantly to the idea formulation, data and coding work, and writing.

3.1 Main

While the introduction of genetically modified (GM) crops profoundly transformed the agricultural landscape of the United States, our understanding of the implications for human health remains limited. At the heart of GM technology lies its resistance to the herbicide glyphosate, which enabled farmers to directly spray glyphosate onto GM crops, eliminating weeds while sparing the crops themselves. Recent research advertises the potential of this technology to improve farm productivity (Hansen & Wingender, 2023; Lusk, Tack, & Hendricks, 2017; Z. Xu, Hennessy, Sardana, & Moschini, 2013), but the pairing of GM seeds with glyphosate introduced a complex array of health externalities—with substantial uncertainty regarding the total effect. On the one hand, if glyphosate replaced more toxic herbicides common to non-GM cultivation, its introduction could yield positive health effects. However, the ensuing liberal application of glyphosate—enabled by the central innovation of herbicide resistance—has led to substantially higher volumes of chemicals sprayed, potentially worsening health (Kniss, 2017; S. Lee, Moschini, & Perry, 2023; Schulz, Bub, Petschick, Stehle, & Wolfram, 2021).

In this study, we use spatiotemporal variation in GM-crop adoption to quantify the sign and magnitude of the human health externality that resulted from the widespread adoption of GM crops and the rapid increase of glyphosate. We then document heterogeneity in the effects of the GM weed management regime on perinatal health. Our results suggest that glyphosate has adverse effects on perinatal health, net of any benefits associated with reductions in the use of other herbicides. These adverse effects concentrate among the most at-risk births.

The US first approved glyphosate for agricultural use in 1974, and the chemical subsequently gained prominence as a broad-spectrum herbicide used extensively in agriculture. Its effectiveness in weed control and relatively low toxicity contributed to its widespread adoption, becoming a critical component of weed-management practices. Pairing glyphosate with glyphosate-resistant GM crops removed a natural limiter in glyphosate use—glyphosate kills non-genetically modified crops and targeted weeds. Relaxing this constraint resulted in dramatic increases in glyphosate application intensity. Before the release of GM seeds in 1996, US farmers applied 0.1 kg of glyphosate per hectare of cropland; since GM seeds, the intensity has risen to over 1.3 kg/hectare (Baker & Stone, 2015). Meanwhile, in the EU, which never approved GM seeds, the glyphosate application rate remains near the US’s pre-GM levels: approximately 0.2 kg/hectare Antier et al. (2020). While potentially affected by other factors, this US-EU gap in glyphosate intensity illustrates how GM technology enabled substantially higher levels glyphosate intensity than otherwise available.

Since its approval, US regulators have maintained “there are no risks to human health from the current registered uses of glyphosate” (U.S. Environmental Protection Agency, 2020)—despite a dearth of population-wide, causally founded

studies. However, two recent studies document negative health impacts of glyphosate exposure in Brazil. Both Dias, Rocha, and Soares Dias, Rocha, and Soares (2023) and Skidmore, Sims, and Gibbs Skidmore, Sims, and Gibbs (2023) find that glyphosate exposure—driven by the expansion of GM seeds and transported through rivers—increased infant mortality and pediatric cancer deaths in Brazil. These two studies offer the first large-scale, population-level, plausibly causal estimates of the health costs of glyphosate exposure. Our study complements these analyses by considering glyphosate impacts in a substantively different socioeconomic setting—the US’s GDP per capita is approximately nine times greater than Brazil’s (World Bank Open Data, 2024)—with a potentially different exposure mechanism stemming from differences in intensity of glyphosate use, hydrology, geology, and meteorology. Brazil applies nearly twice the amount of glyphosate per hectare of cropland as the US (Dias et al., 2023). Thus, while the US and Brazil overlap as leaders in glyphosate application, key differences between the contexts warrant new study in the US and abroad. Finally, given the Brazilian context, the papers focus on effects driven by exposure through upstream glyphosate application; we find that local exposure drives negative health impacts in the US context.

3.1.1 Background and motivation. glyphosate and other herbicides pervade the environment, giving rise to multiple pathways of human exposure (Gillezeau et al., 2019) including water (Medalie et al., 2020), dust blown by the wind (Ramirez Haberkon, Aparicio, & Mendez, 2021), aerial drift (Hu, Wu, Lyu, Ning, & She, 2023), direct contact (Foucault et al., 2021), and food residue (Pappas & Foos, 2023). Following application, glyphosate exhibits a relatively short breakdown period, with a half-life ranging from 2 to 215 days (Battaglin, Meyer,

Kuivila, & Dietze, 2014). Although a significant portion of the herbicide binds to the soil, reducing runoff, its high water solubility allows the unbound remnant to enter both surface and groundwater (Duke & Powles, 2008; Hartmann et al., 2021). A comprehensive study across US waterways from 2015–2017 revealed the presence of glyphosate or its degradate, aminomethylphosphonic acid (AMPA), in 90% of samples (Medalie et al., 2020). Additionally, wind-dispersed dust particles containing soil-bound herbicide residues can contribute to air pollution (Ramirez Haberkon et al., 2021). While food residue is suspected to be another source of population-wide exposure to glyphosate (Landrigan, 2018), we cannot capture its effect in this study.

3.1.1.1 Exposure. The multiple exposure mechanisms—coupled with the breadth and volume of glyphosate application—have resulted in widespread detection of glyphosate in the urine and blood of US residents. The US Center for Disease Control (CDC) detected glyphosate in 81% of urine samples from a nationally representative cohort (Ospina et al., 2022). Multiple studies with pregnant women found glyphosate present in the urine of nearly every tested mother-to-be (Gerona et al., 2022; Parvez et al., 2018). This ubiquity of glyphosate exposure in the US population highlights the importance of understanding the impacts of glyphosate exposure at a national scale—particularly within populations likely exposed to higher levels. We focus on rural populations’ exposures to local glyphosate sources—glyphosate exposure through dust, drift, direct contact, or water originating within the county of residence. We explore the potential effects of upstream spraying in SI Appendix N.

3.1.1.2 Health impacts. A growing body of literature suggests that glyphosate has the potential to impact human health through a variety of biological

mechanisms. Existing evidence typically comes from either laboratory studies on non-human animals or human-focused observational studies. While laboratory-based studies offer well-identified causal effects, they often suffer challenges of external validity. Previous observational studies (Arbuckle, Lin, & Mery, 2001; Gerona et al., 2022; Parvez et al., 2018) primarily focus on associations between self-reported exposure and health outcomes and typically avoid causal claims. Dias, Rocha, and Soares Dias et al. (2023), Skidmore, Sims, and Gibbs Skidmore et al. (2023), and Camacho and Mejía Camacho and Mejia (2017) are exceptions—employing quasi-experimental methods to make causal statements.

Lab studies have established links between glyphosate and congenital anomalies in rats, developmental issues in frogs and chickens, and endocrine disruption for male reproduction in mice (Dallegrave et al., 2003; Paganelli, Gnazzo, Acosta, López, & Carrasco, 2010; Pham et al., 2019). Multiple additional studies link glyphosate exposure to endocrine disruption, which can affect developmental and reproductive health (Muñoz, Bleak, & Calaf, 2021) and often have nonmonotonic effects (Vogel, 2008). Research also documents glyphosate toxicity for placental cells—raising concerns for adverse effects in fetal development (Benachour et al., 2007).

Consistent with lab-based concerns for the effect of glyphosate on development and reproductive health, several observational studies report associations between glyphosate exposure and miscarriage (Arbuckle et al., 2001), gestational length (Parvez et al., 2018), and birthweight (Gerona et al., 2022)—related to a growing literature on health risks from occupational exposure (Tago, Andersson, & Treich, 2014) to glyphosate Zhang, Rana, Shaffer, Taioli, and Sheppard (2019) and other pesticides (Leon et al., 2019). Considering these

established mechanisms and documented associations, we evaluate the evidence for a causal effect of glyphosate on perinatal health outcomes—particularly birthweight and gestational length.

Beyond glyphosate-specific studies, a growing body of literature documents adverse health effects associated with pesticide exposure, even at low doses (Calzada, Gisbert, & Moscoso, 2023; Fletcher & Noghanibehambari, 2024; E. Frank & Taylor, 2022; E. G. Frank, 2024; Jones, 2020; Larsen, Gaines, & Deschênes, 2017; Taylor, 2022). One of the closest studies to the current paper—by Larsen, Gaines, and Deschenes Larsen et al. (2017)—finds pesticide exposure increases adverse birth outcomes (weight, gestation, and abnormalities) for California mothers. The authors highlight that this effect is driven by the sub-population exposed to the highest levels of pesticides. These results and the growing body of literature motivate three points for our paper. First, birthweight and gestation are plausible outcomes to test for the health effects of glyphosate. Second, if glyphosate affects health, rural populations are likely the most impacted. Third, the effect of glyphosate may be heterogeneous—varying within the exposed population. These three observations lay the foundation of our analysis.

3.1.2 Empirical approach. We employ two complementary empirical approaches—*reduced-form DID* and *two-stage*—to estimate the perinatal health impacts from the rollout of GM crops and the ensuing intensification of glyphosate. Both approaches employ similar strategies to isolate plausibly exogenous variation in county-level GM adoption and glyphosate exposure by combining temporal variation in the commercial release of GM seeds in the US with spatial variation in the suitability for growing the main crops for which GM seeds are available—corn, soy, and cotton. The temporal dimension of this strategy utilizes the arbitrary

timing of the release of GM seeds. GM seeds became commercially available in the US in 1996, and farmers rapidly adopted GM seeds and intensified glyphosate applications in the following years. The strategy’s spatial dimension uses the fact that these changes, on average, had larger impacts in places that were relatively more suitable for the crops—typically where the crops were already grown. Consequently, each approach leverages two comparisons to identify glyphosate’s effect on perinatal health outcomes: (1) *before* versus *after* GM-induced glyphosate expansion, and (2) *areas more suitable* for GM crops versus *less-suitable areas*.

The two empirical approaches provide complementary strengths. The *reduced-form DID approach* requires weaker assumptions for identification (discussed in *Methods* 3.3.5) but only provides comparisons between more and less GM-suitable regions—rather than direct quantification of glyphosate’s health impacts. The *two-stage approach* provides this quantification but comes with stronger assumptions.

3.1.2.1 Data. We measure perinatal health outcomes using the universe of individual-level birth data from the National Vital Statistics System (NVSS) during 1990–2013—the restricted-access natality files that allow us to match each birth to the counties of occurrence and the mother’s residence (National Center for Health Statistics, 1990–2013). These natality data include our primary perinatal health outcomes—birthweight and gestational length—and demographic and residence information from both parents and birth-location information. Our primary analyses focus on the over 10 million births that occurred between 1990 and 2013 in rural US counties or involved mothers residing in rural counties—as defined by the US Department of Agriculture (USDA). We focus on

this subset as it represents the births most likely to be impacted by the increase in glyphosate intensity and exposure induced by the release of GM seeds.

Using each infant’s mother’s county of residence, we match the birth to a measure of glyphosate exposure—glyphosate intensity, defined as the volume of annual, county-level estimates of agricultural glyphosate applications per square kilometer of total county area—from the USGS National Pesticide Synthesis Project spanning 1992–2017 (Baker & Stone, 2015). Our measures of crop suitability—attainable yield estimates for corn, soy, and cotton—come from the Food and Agriculture Organization of the United Nations Global Agro-Ecological Zones modelling framework (FAO-GAEZ). These time-invariant predicted yields result from modeling crop responses to environmental conditions such as soil type and climate—holding management practices constant (FAO & IIASA, 2021). We calculate a continuous GM suitability percentile from the GAEZ data using a process described in *Methods*.

Fig. 12 illustrates both the spatial and temporal variation that we use to identify the effect of glyphosate on infant health. Fig. 12a maps the spatial variation in GM-crop suitability. This suitability strongly correlates with the increase in glyphosate intensity (kg/km^2) after the introduction of GM seed—shown in Fig. 12b. Fig. 12c highlights how the commercial release of GM seeds (dashed vertical line in 1996) drove marked increases in glyphosate use—particularly for higher GM-suitability counties. While the figures show *all* counties, we only use rural counties in our main analysis; Fig. S1 maps these rural counties. Additionally, Fig. S2 shows state-level adoption rates of GM seeds over time—adoption was rapid among all states.




Figure 12. GM-crop suitability predicts glyphosate increases after GM-seed introduction.

(a) Percentile of attainable yield for GM crops equals the difference in attainable yield between high- and low-input scenarios from FAO GAEZ FAO and IIASA (2021) for corn, soy, and cotton. We rescale each crop to be a national percentile, take the maximum over the three crops, and finally scaling again to be a national percentile. **(b)** Change in glyphosate censored at the 1st and 99th percentiles. **(c)** Total glyphosate application by GM crop suitability quartile. We restrict our analysis to rural counties. See Fig. S1 for a map of the rural designation and Fig. S2 for state-level GM adoption over time.

3.1.2.2 Estimation. We now discuss our two empirical approaches.

Methods section *Empirical strategy* elaborates on these strategies’ assumptions and implementations.

Reduced-form DID approach Our *reduced-form DID approach* estimates how infant-health outcomes in higher- and lower-GM-suitable counties diverged after the introduction of GM seeds and the subsequent glyphosate intensification (conditional on controls). Accordingly, this approach employs a difference-in-differences (DID) estimator with continuous ‘treatment’ (GM-crop suitability)—mainly relying upon a parallel-trends assumption, i.e. if GM seeds had not been introduced, more and less GM-suitable counties would have continued on similar trajectories (Angrist & Pischke, 2009).

Two-stage approach Our second approach uses a similar set of comparisons to the reduced-form approach—high *vs.* low GM suitability, before *vs.* after the GM-seed rollout. However, our *two-stage approach*—implemented as two-stage least squares (2SLS)—scales reduced-form estimates by a first stage: the increase in glyphosate intensity (kg/km^2), estimated by the same double difference. We capture the pattern of adoption by interacting the instrument (GM-crop suitability) with indicators for year. 2SLS’s central assumption for a causal interpretation is the exclusion restriction (Angrist & Pischke, 2009), i.e. conditional on our controls, any divergence in perinatal health between high- and low-GM suitability counties after the GM-seed rollout only occurred through the channel of increased glyphosate exposure (discussed further in *Methods*).

Event studies Because the instrument in our 2SLS approach interacts GM-crop suitability with year indicators, its reduced-form and first-stage results provide convenient graphical summaries, i.e. event studies. We define *suitability* as a county’s maximum percentile of attainable yield across the GM crops corn, soy, and cotton—scaled between 0 and 1 (further discussed in *Attainable yield* in *Methods*). One can interpret the event study’s yearly coefficients as summarizing the difference between the average health outcome in high-GM-crop suitability counties relative to low-GM-crop suitability counties each year (relative to the 1995 difference). The event study coefficients depict how infant-health differences between high- and low-suitability counties evolved through time.

These event studies also provide evidence on the plausibility of the parallel-trends assumption upon which the reduced-form DID results rely. Because the event studies are also reduced-form results, causal interpretation only relies upon the parallel-trends assumption, rather than the stronger exclusion restriction required by 2SLS.

Separating glyphosate’s direct and policy effects Interpreting these causal effects requires some nuance. glyphosate-tolerant GM seeds allowed farmers to change their weed management practices—reducing their usage of non-glyphosate pesticides and mechanical tilling (S. Lee et al., 2023). These—and other—changes could potentially affect perinatal health, violating the exclusion restriction. Therefore, we present two effects in our results—a *policy effect* and a *glyphosate effect*.

The *policy effect* does not control for these other changes and thus captures the total (net) effect of the introduction of GM seeds. For our reduced-form DID results, the policy effect is potentially the more policy-relevant parameter

as it aggregates across all of the health effects brought about by GM seeds and glyphosate intensification relative to alternative practices in place before the introduction of GM seeds. For this reason, the exclusion restriction for the 2SLS-estimated policy effect represents a fairly strong assumption—violated if the GM-seed rollout induced employment or non-glyphosate agricultural changes that affected infant health.

Alternatively, by directly controlling for changes induced by the GM rollout—non-glyphosate pesticides, fertilizer, population composition, and various local economic conditions—the *glyphosate effect* provides an estimate of the direct causal effect of glyphosate on perinatal health. This direct glyphosate effect still requires an exclusion restriction: except for the pesticides, fertilizers, population composition, and economic outcomes for which we control, no other mechanisms affected perinatal that also correlated with both suitability for GM crops *and* the timing of the GM-seed rollout. Accordingly, the glyphosate effect is likely more appropriate for the two-stage approach—permitting a *per unit of glyphosate* interpretation with a plausible exclusion restriction.

Finally, glyphosate-based herbicides like Roundup typically contain other chemicals, e.g. surfactants that reduce surface surface tension. Our measured effects include impacts from other ingredients mixed with glyphosate in commercial herbicide formulations.

Heterogeneity In addition to estimating the *average* effect of glyphosate exposure (methods described above), we also estimate heterogeneous effects. Understanding heterogeneity in this setting is critical: health and policy implications can differ depending on *whom* glyphosate impacts and *how diffuse* the impacts are. Reductions in birthweight among infants already at risk for low

birthweight are likely more costly than decreases in birthweight among higher-weight infants.

We estimate glyphosate’s impact as a function of the infant’s *expected* birthweight, which we estimate using a random forest trained to predict infants’ expected birthweights. Specifically, this learning model predicts birthweight as a function of the infant’s and parents’ information—using data from all pre-1996 births, before the glyphosate-resistant seeds were widely available. (See SI Appendix 3.3.8 for detailed methodology.) The resulting predictions enable us to test whether glyphosate’s adverse health impacts affected infants equally or whether they concentrated in low- or high-weight births. Using *predicted* birthweight allows us to avoid bias from splitting the sample on the outcome (Elwert & Winship, 2014). Subsequently, we estimate our reduced-form and 2SLS results by predicted-birthweight percentiles.

3.1.3 Results.

3.1.3.1 *Reduced-form evidence of GM rollout’s health effects.*

The event study in Fig. 13a demonstrates that GM crop suitability strongly predicts post-1996 increases in glyphosate application levels. Because GM crop suitability is a continuous percentile, the plotted coefficients represent the difference in glyphosate intensity between the highest and lowest GM suitability counties relative to their difference in 1995. Before 1996, lower- and higher-GM-crop-suitable counties followed similar glyphosate trajectories: the event study remains relatively flat and near zero. However, after the 1996 introduction of GM seeds, glyphosate intensity in counties with higher attainable yields for corn, soy, and cotton quickly outpaced glyphosate intensity in less-suitable counties. As adoption accelerated, this glyphosate intensity gap between high- and low-attainable yield counties




Figure 13. GM-seed introduction increased glyphosate intensity in GM-crop suitable areas; birthweight reductions were also higher in GM-suitable counties and match GM-seed timing.

(a) Estimated event-study coefficients for the effect of local GM max attainable yield percentile on glyphosate by year relative to 1995. Pesticide data only go back to 1992 (no coefficients in 1990–1991). **(b)** Similar event study for birthweight. Both regressions represent the *policy-effect* specification: county and year-by-month fixed effects and controls for parent demographics (mother’s age, race, education, marital status, birth facility, resident status, previous births, fathers age and race, and sex of infant). Standard errors cluster by state and year. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.

widened. The event study confirms the strength of our instrument (percentile of GM-crop suitability) and illustrates the first stage of our 2SLS approach.

The event-study results mirror the reduced-form DID estimates for the *policy effect* on glyphosate in first row of *Panel A* of Fig. 14 (left columns). The introduction of GM seeds significantly increased the glyphosate intensity in high-GM-crop suitability counties by .023 kg/km² relative to low-suitability counties. This effect is large—approximately the same level as the 2012 mean glyphosate intensity.

These first-stage findings align with a plausible mechanism: higher attainable yields fostered greater GM crop adoption, increasing glyphosate application. *SI Appendix* Fig. S4 shows applications of several prevalent herbicides—alachlor, cyanazine, fluazifop, and metolachlor—decreased following the release of GM seeds, suggesting farmers substituted away from these herbicides and toward glyphosate. Consequently, we control for these pesticides—along with fertilizer, several employment measures, and population composition—when estimating the *glyphosate effect* and omit them when estimating the *policy effect*.

Fig. 13b depicts a similar event study for birthweight—plotting each year’s average birthweight difference between births in high and low GM-crop suitability counties (relative to the 1995 difference) after controlling for maternal demographics. Prior to the introduction of GM crops, the birthweight gap between births in higher- and lower-attainable-yield counties remained stable. However, beginning in 1996—coinciding with the release of glyphosate-tolerant seeds and the intensification of glyphosate application—birthweights in higher-GM-suitability counties began declining relative to lower-yield counties. In 2005, a decade after the

introduction of glyphosate-tolerant seeds, the average birthweight in the highest-yield county had fallen approximately 30 grams relative to the lowest-yield county.

Again, the reduced-form DID estimates for the policy effect on birthweight in Panel A of Fig. 14 match the event study. Relative to low-suitability counties, the average infant in high-suitability counties lost 29.4 grams from the introduction of GM seeds and the ensuing intensification of glyphosate.

The event study in Fig. 13b also visualizes the reduced-form estimates of our 2SLS estimator: the effect of high attainable yield for GM crops on birthweight over time. The flat trend prior to 1996 in the event study coefficients of Fig. 13b also supports the parallel-trends assumption that underpins our empirical approach: before the introduction of GM-tolerant seeds, counties followed similar trajectories independent of GM-crop suitability.

The remainder of Panel A in Fig. 14 reproduces similar reduced-form DID policy-effect estimates for five additional health outcomes. The results depict similar, statistically significant adverse health effects for perinatal health. In higher-suitability counties, the average birth's gestation length declines 1.03 days, probabilities of LBW and VLBW increased .51 pp and .11 pp, probability of preterm increased 1.64 pp, and health index declined by .011. Reassuringly, the probability of C-section was unaffected. *SI Appendix* Fig. S5 provides reduced-form event studies for each of these results.

The reduced-form DID estimates for the *glyphosate effect* on the right side of Panel A (Fig. 14) tell a very similar, slightly attenuated, story: the introduction of GM seed significantly reduced perinatal health in high-suitability counties (LBW and VLBW are significant at the 10% level). The fact that directly controlling for potential confounders of glyphosate does not substantially alter the results

lends plausibility to the exclusion restriction required by our 2SLS. Together, the reduced-form event-study and DID results bear considerable evidence that the release of GM seeds in 1996 led to significant deterioration of perinatal health in high GM-attainable-yield counties relative to low-yield counties.

Outcome (<i>unit</i>)	Policy effect		Glyphosate effect		2012 Mean
	Estimate	Conf. Interval	Estimate	Conf. Interval	
<i>Panel A: Reduced-form DID results</i>					
Glyphosate (<i>kg/km²</i>)	.023	[.014, .033]	.011	[.006, .017]	.023
Birthweight (<i>g</i>)	−29.4	[−44.5, −14.4]	−23.8	[−38.9, −8.68]	3,271.1
Gestation (<i>days</i>)	−1.03	[−1.55, −.52]	−.83	[−1.32, −.34]	270.5
LBW (<i>pp</i>)	.51	[.10, .91]	.47	[−.08, 1.02]	8.0
VLBW (<i>pp</i>)	.11	[.03, .19]	.09	[−.06, .23]	1.4
Preterm (<i>pp</i>)	1.64	[.58, 2.70]	1.07	[.16, 1.97]	20.7
Health index	−.011	[−.016, −.006]	−.010	[−.015, −.005]	−.015
C-section (<i>pp</i>)	.11	[−1.33, 1.55]	−.03	[−1.40, 1.33]	27.8
<i>Panel B: 2SLS results</i>					
Birthweight (<i>g</i>)	−19.6	[−34.7, −4.6]	−29.8	[−52.5, −7.2]	3,271.1
Gestation (<i>days</i>)	−1.00	[−1.55, −.45]	−1.49	[−2.28, −.70]	270.5
LBW (<i>pp</i>)	.39	[.03, .75]	.65	[.18, 1.12]	8.0
VLBW (<i>pp</i>)	.11	[.03, .19]	.20	[.07, .34]	1.4
Preterm (<i>pp</i>)	1.41	[.39, 2.43]	2.14	[.73, 3.55]	20.7
Health index	−.0093	[−.015, −.0038]	−.013	[−.022, −.0039]	−.015
C-section (<i>pp</i>)	.76	[−.19, 1.70]	.87	[−.83, 2.58]	27.8

Figure 14. Glyphosate’s direct and policy effects reduced birth outcomes for the average rural birth.

All reported estimates are the effect at the weighted mean of glyphosate in 2012, where we weight by total births. *Panel A* reports difference-in-difference estimates comparing high- and low-GM suitability counties before and after the release of GM seeds in 1996. We use 1990-1995 as our pre-period and 2005-2010 as the post-period to drop the transition phase. *Panel B* reports 2SLS estimates from 2SLS regression of perinatal health on glyphosate, instrumenting for glyphosate with GM suitability interacted with year dummies. The *Policy effect* controls for family demographics. The *glyphosate effect* controls for other pesticides and unemployment. See text for details. *LBW* and *VLBW* give the probabilities of low birthweight (<2500g) and very low birthweight (<1500g) in percentage points (0–100). 95% confidence intervals calculated using standard errors clustered by year and state. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.

3.1.3.2 The effect of glyphosate on perinatal health. Our two-

stage results integrate the variation depicted in Fig. 13 into a 2SLS framework—

effectively dividing the health-outcome event study in Fig. 13b by the glyphosate event study in Fig. 13a. In fact, scaling the reduced-form DID by the corresponding glyphosate effects generates very similar estimates to 2SLS (*SI Appendix* Table S4). *Panel B* in Fig. 14 reports the estimated effect (and 95% confidence interval) of glyphosate exposure at the 2012 mean level of glyphosate intensity (0.023 kg/km²). *SI Appendix* Table S3 contains the corresponding regression coefficients.

As discussed above, the *glyphosate effect* (i.e. controlling for potential violations of the exclusion restriction) is likely the more appropriate parameter to consider for the two-stage results. Accordingly, we focus our discussion of the 2SLS results on the glyphosate effect. However, Table 14 provides 2SLS policy-effect estimates for comparison.

The 2SLS glyphosate-effect results in *Panel B* of Fig. 14 (right side) indicate that for the average rural birth, exposure to glyphosate at its 2012 mean intensity significantly reduced birthweight by 29.8 grams; reduced gestation 1.49 days; increased the probabilities of LBW, VLBW, and preterm by .65, .20, and 2.14 pp; and reduced the composite health index by .013. Comparing these effects to each outcome’s mean value suggests that the effects are indeed meaningful: relative to the outcomes’ means, the effects represent 0.9% of mean birthweight, 0.6% of mean gestation, 8.1% of LBW, 14.3% of VLBW, 10.3% of preterm probability, and 87.7% of the health index. Again, the effect on the probability of a C-section is positive but not statistically significant and small in magnitude. Together, these results indicate that even at the mean level of intensity in the US, glyphosate exposure significantly deteriorates infant health.

These glyphosate effects result from local (own-county) exposure. Unlike Dias, Rocha and Soares Dias et al. (2023) and Skidmore, Sims, and Gibbs

Skidmore et al. (2023), we do not find any significant effects of glyphosate sprayed upstream—SI Appendix N.2 contains the results of our water-based analysis.

These findings are robust to a number of modeling alternatives. *SI Appendix* Fig. S6 provides 2SLS estimates for the marginal effect of glyphosate on birthweight, where we vary the inclusion of controls, fixed effects, and the definition of attainable yield. Event studies for birthweight under these alternative models are also robust (SI Appendix F), and specification charts for other outcomes are in SI Appendix G. We also estimate OLS results without instruments in SI Appendix C and a “shift-share” model with slightly different identifying assumptions in SI Appendix E. Finally, *SI* Fig. S21 shows results from estimating the model on different geographic subsets of the United States. In each case, the results are qualitatively unchanged.

We also investigate various threats to identification. A primary concern for our approach is whether the release of GM seeds coincided with non-glyphosate socioeconomic effects which also affected birthweight—e.g. employment or income. Notably, there is no evidence that the introduction of GM seeds significantly affected average farm or non-farm income in the study counties. The event study for the unemployment rate does show a trend—suggesting the potential for bias if excluded (*SI Appendix* Fig. S28). However, controlling for unemployment does not meaningfully change our results, suggesting employment changes are not driving our results.

3.1.3.3 Heterogeneity in glyphosate’s health impacts. The results in the previous sections estimate the average effect of glyphosate within a heterogeneous population—where the effect of glyphosate may differ across individuals. Fig. 15 shows the two-stage estimates for policy and glyphosate

effects on birthweight vary considerably by decile of expected birthweight. *SI Appendix H* explains our prediction of birthweight. Both effects exhibit considerable heterogeneity that follows the same pattern: the magnitude of glyphosate’s estimated effect on birthweight is largest in the lowest decile (i.e. for births with the lowest expected birthweight), and the effect declines as the percentile of predicted birthweight increases. *SI Appendix Fig. S9* shows first-stage and reduced form results for birthweight by predicted birthweight quintile. The glyphosate effect on birthweight in the first decile is 12 times larger than in the tenth decile. Accordingly, glyphosate’s largest impact concentrates among the most vulnerable births: the estimated glyphosate effect in the first decile at mean 2012 glyphosate exposure is a loss of 75 grams, relative to just 6 grams in the tenth decile of predicted birthweight.




Figure 15. Birthweight losses due to glyphosate and GM are largest for births with the lowest expected birthweights.

Estimates of the policy and glyphosate effects on birthweight by predicted birthweight deciles. Each estimate is from a separate regression. All regressions include controls for family demographics, county fixed effects, and year by month fixed effects. Sample restricted to births occurring in a rural county or to mothers residing in a rural county. Standard errors are clustered by year and state. glyphosate instrumented with GM max attainable yield percentile.

We also find significant evidence that glyphosate’s effects vary with expected birthweights for several other health outcomes—gestation length and the probabilities of low and very-low birthweight. The increase in the probabilities of low birthweight and very-low birthweight are more than 60 times larger in the lowest decile than in the top decile of expected birthweight. *SI Appendix* Fig. S10 illustrates heterogeneity in the policy effect for all outcomes by predicted birthweight quintiles, deciles, and ventiles. As with birthweight, the most acute effects for gestation occur in lower percentiles, declining in magnitude with predicted birthweight percentile. Unlike our birthweight results, where we find no significant effect among higher percentiles, we find statistically significant evidence that glyphosate reduces gestation length in every decile. At the 2012 mean level of glyphosate exposure, the glyphosate’s effect on gestation ranges from -2 to -0.8 days for the average birth. Again, these results suggest the most vulnerable infants bear the largest perinatal health impacts.

3.2 Discussion

This study finds significant evidence that glyphosate adversely affected births across several measures of perinatal health throughout the rural United States in the last twenty years. To our knowledge, these findings are the first quasi-experimental evidence of glyphosate’s adverse health effects at a population scale in the US. However, they are consistent with a growing body of literature documenting glyphosate’s negative impact on development and reproduction. Using machine learning to predict expected birthweights, we find significant heterogeneity underlying the effect of glyphosate on perinatal health: glyphosate’s adverse effects are largest for infants with the lowest expected birthweights.

Other effects While this estimated effect of GM seeds and glyphosate on perinatal health highlights a critical consequence of herbicide application, it represents only part of the total potential externality associated with increased chemical usage. The previously discussed toxicology literature suggests several additional mechanisms and effects on human health that we do not measure here. Lab experiments (Dallegrave et al., 2003; Paganelli et al., 2010; Pham et al., 2019) and recent observational work suggest additional ecological costs, such as biodiversity loss (Engist, Guzman, Larsen, Church, & Noack, 2024).

Although it does not address all of glyphosate’s effects, using birthweight as our primary outcome offers several advantages. As discussed above, prior research establishes adverse effects of glyphosate on reproduction/development and finds glyphosate in the urine of nearly every expectant mother. Numerous studies then link birthweight to other later-life outcomes (Bharadwaj, Lundborg, & Rooth, 2017; Heijmans et al., 2008). Additionally, data on birthweight is widely available and reliable—the US has accurately and objectively measured the birthweight of nearly every child born for decades. Infants are also less prone to endogenous responses to health and environmental shocks. Similarly, infants have short histories over which they can accumulate exposure—simplifying exposure measurement. Birthweight thus provides both a key health outcome and a canary-in-a-coalmine-like indicator of glyphosate toxicity. Ample opportunities exist for future work to test additional costs and benefits of glyphosate and GM technology to more broadly assess welfare effects.

Heterogeneity, mechanisms, and inequity Several mechanisms potentially explain the observed heterogeneity. While we cannot fully disentangle these mechanisms, we use the dimension along which we observe heterogeneity—

predicted birthweight—to understand which attributes correlate with the most-affected births. As predicted birthweight results from a predictive model trained on the demographics of infants’ parents, we observe which parental attributes correlate with lower predicted birthweight and, consequently, more adverse glyphosate effects. Fig. 16 reveals that lower predicted-birthweight infants are more likely to be female, Black, and/or children of unmarried parents. The predictive model places nearly all births to Black parents into the first quintile, where we detect the most adverse impacts of glyphosate. If we estimate our main model separately for births to white and non-white mothers, we find that the *Glyphosate Effect* is 1.8 times larger for births to non-white mothers relative to white mothers (*SI Appendix* Fig. S7). This finding links to an extensive environmental justice literature that documents Black households’ unequal pollutant exposure (Banzhaf, Ma, & Timmins, 2019b).

We observe similar effects on male and female births *within* birthweight quintiles (*SI Appendix* Fig. S11). Therefore, the increased effect among low-predicted birthweights is not driven by a larger effect among females. However, female births still bear a larger health burden since they make up a larger proportion of low-predicted births.

The fact that the adversity of glyphosate’s effects correlates with race and sex highlights a potentially serious issue for equity. Further, we find the most adverse effects among the lowest expected-weight births—potentially magnifying short-term healthcare costs (Beam et al., 2020), later-life outcomes/welfare (Bharadwaj et al., 2017), and epigenetic/intergenerational consequences (Heijmans et al., 2008). Consequently, our results have potentially important implications for equity and justice in the US.




Figure 16. Female infants, children of Black and non-White parents, and children of unmarried parents have lower predicted birthweights.

Each line represents the percent of births (y -axis) within the demographic group (line color) at each predicted birthweight percentile (x -axis) for births to mothers with rural residences. Predicted birthweight percentile is calculated relative to the distribution of rural-residence births in years prior to 1996. Averages cover two-percentile bins.

Unfortunately, data limitations restrict us from further testing potential mechanisms. While our empirical approach recovers the causal effect of glyphosate on perinatal health, it does not provide causal estimates for the drivers of heterogeneity. One potential mechanism for the heterogeneity is differential exposure. Because the natality and glyphosate data are only resolved at the county, we cannot test whether infants with lower predicted birthweights face higher levels of glyphosate relative to other infants in their county. Differential baseline health (or healthcare access) could also contribute to the observed heterogeneity. Further, the shape of glyphosate’s damage function is unknown. If glyphosate has nonlinear effects on health—nonlinear in glyphosate exposure or due to interactions with other health risks/complications—differential exposure or heterogeneous non-glyphosate health risks could also produce these heterogeneous impacts. Finally,

our examination of heterogeneity focuses on only one of many possible dimensions. Future work could contribute to many of these issues with more resolved data.

Magnitude of results The results in this paper indicate glyphosate exposure adversely affects perinatal health. Whether the estimated effects are large in magnitude depends upon several dimensions of heterogeneity and potentially on additional nuance.

It is important to consider that the effects presented in Table 14 are average effects, in two senses, namely, averaged across all rural infants and calculated for the exposure to the average glyphosate intensity in 2012 (approximately .023 kg/km²). For birthweight, this average effect of glyphosate is a reduction of 20–40 grams (1 percent). The corresponding implied increases in probabilities of low birthweight and preterm birth (.4–.6 pp and 1.4–2.1 pp, respectively) represent percent increases of approximately 5–9% and 7–10%. If all infants faced a 20–40 gram reduction in birthweight (or 1%), there may be minor concern for public health, with the exception of the most vulnerable infants—e.g. the lowest-weight infants for whom marginal benefits of weight may be high. Indeed, Currie, Greenstone, and Meckel (2017) consider a similar birthweight effect small (Currie, Greenstone, & Meckel, 2017). However, our estimated average effect on birthweight is larger than impacts estimated by several recent studies on regulated pollutants—e.g. 13.2-gram reduction from a one-standard-deviation increase in late-term CO exposure (Currie, Neidell, & Schmieder, 2009) and 6–7 gram reduction for each ‘cheating’ diesel Volkswagen (per 1000 cars) (Alexander & Schwandt, 2022). The magnitude of our average effect is quite close (but opposite sign) to estimates for the effect the Food Stamp Program (Almond, Hoynes, & Schanzenbach, 2011).

Our estimates for the average effects on the probabilities of low birthweight and preterm births are very similar (both in level and percent) to recent work on exposure to leaking underground petroleum storage tanks (8.7% LBW; 7.4% preterm) (Marcus, 2021), automobile-based air pollution (11.8% LBW; 10.8% preterm) (Currie & Walker, 2011), late-term exposure to CO (8.6% LBW) (Currie et al., 2009), and in utero exposure to chemicals via drinking water (6.5% LBW) (Currie, Graff Zivin, Meckel, Neidell, & Schlenker, 2013). Our estimated average effect of glyphosate on the probability of low birthweight is approximately one third of estimated effects for infants born to families within 1 km of a fracking site (Currie et al., 2017).

However, Fig. 15 highlights significant heterogeneity underlying the *average* effects discussed above. Considering both the heterogeneity depicted in Fig. 15 and additional heterogeneity in glyphosate intensity is critical in assessing the magnitude of the estimated effects.

As Fig. 15 illustrates, the estimated impact of glyphosate on birthweight for infants in the lowest expected-birthweight decile are 5–12 times larger than the highest decile—and approximately 2.5 times larger than the average effect. At the mean level of 2012 glyphosate exposure, the estimated effects for infants in the first decile are 46–77 grams—suggesting infants in this already-low-weight group lost 2–4% of their expected birthweight due to glyphosate. We find even larger heterogeneity on the probability of low-weight birth: the effect (.7–1.1 pp increase) for the lowest decile is 4 times larger than the average effect and 50–60 times larger than the effect for the top decile.

The paper presents results as effects from exposure to the *mean level of 2012 glyphosate intensity*—e.g. Table 14 (approximately .023 kg/km²). However,

glyphosate intensity varies substantially in the US: exposure is 50% higher at the 75th percentile, 200% higher at the 90th percentile, and 300% higher at the 95th percentile. Accordingly, at the 90th percentile (i.e. for a quarter of rural US births), we estimate that the *average* (pooling across expected birthweight) loss in birthweight is 62–95 grams—i.e. 2–3% of the mean birthweight. Exposure to glyphosate at the 90th percentile of glyphosate intensity increases the probabilities of low birthweight and preterm birth by 1–2 pp (LBW) and 4–7 pp, respectively. At the 90th percentile, our estimated effects for the probability of low birthweight are on the same scale as the effect of living within 1 km of a fracking site (Currie et al., 2017). The relatively large size of effects implied at higher percentiles of pesticide intensity is consistent with the results of Larsen, Gaines, and Deschênes, who find adverse infant health effects from pesticides concentrate in the top 5 percent of exposure (Larsen et al., 2017). It is important to note that this discussion assumes glyphosate’s damage function is linear (*SI* Section I discusses further). Regardless of the damage function’s shape, only considering *average exposure* may obscure key insights into glyphosate’s health impacts within the US population.

The discussion above separately considers the two dimensions of heterogeneity. However, low-expected birthweight infants are not all exposed to the average glyphosate intensity. In fact, the correlation between expected birthweight and glyphosate intensity is quite low (.02). For a first-decile expected birthweight infant, exposed at the 90th percentile of glyphosate (affecting approximately one percent of rural births), the estimated weight lost due to glyphosate is 146–243 grams—a loss of 5–9 percent of the infant’s expected weight. The corresponding calculation for the probability of a low-weight birth implies an increase of 5–

8 percentage points (on a base). These estimates suggest a small subset of the population may bear particularly large health burdens from glyphosate exposure.

Finally, the paper’s results likely present a lower bound for the true impact of glyphosate due to aggregation bias (ecological fallacy) inherent to the data—understating glyphosate’s actual health burden for exposed individuals. Because the match between birth outcomes and glyphosate exposure is only possible at the county level, it is likely that some infants with little true exposure to glyphosate received high-glyphosate exposure in the data (the reverse is also likely). The resulting aggregation bias attenuates the differences between true high- and low-exposure individuals—likely attenuating our estimates (Banzhaf et al., 2019b). Future work with better measurements of mothers’ exposures could resolve this issue as recent work has done in other settings (Currie et al., 2017, 2009).

Broader considerations GM crops and the resulting glyphosate intensification profoundly changed agriculture. In this study, we quantify one health externality caused by these technological innovations, which reduced average birthweight by 23–32 grams at the average level of glyphosate exposure.

To put the estimated glyphosate health damages in perspective, we convert them to dollars. Waitzman, Jalali, and Grosse estimate that a preterm birth costs an additional 82 thousand USD (2023 dollars) relative to a full-term birth (Waitzman, Jalali, & Grosse, 2021). This estimate includes additional medical care following the birth, special education expenses, and lost labor market earnings later in the child’s life. These estimates may omit additional costs from glyphosate. We combine our estimates from Table 14 on the increased probability of preterm birth and the number of total births in 2012. This calculation implies the economic costs

of the reduced-form policy effect were 800 million USD annually and just over 1 billion USD annually for the 2SLS glyphosate effect (both in 2023 dollars).

Our findings, combined with other recent work (Camacho & Mejia, 2017; Dias et al., 2023; Skidmore et al., 2023), challenge the prevailing regulatory position that GM crops and their associated agricultural practices are safe—and even beneficial—for health. Advocacy historically argued that glyphosate is less toxic than the herbicides it replaced. The mounting evidence of the negative health externalities associated with the rollout GM crops and the ensuing glyphosate intensification warrant new policy discussions about informed, efficient, and equitable regulation of these technologies. Efficient policy must carefully weigh these practices’ economic benefits against the adverse health effects we identify and other human/ecological costs. Further, policymakers must consider the unequal burden glyphosate appears to levy.

Additional work is needed to better understand the benefits and costs of GM crops, glyphosate, and the specific exposure mechanisms underlying their effects. For instance, unlike Dias, Rocha, and Soares Dias et al. (2023) and Skidmore, Sims, and Gibbs Skidmore et al. (2023), we do not find evidence of health effects from upstream glyphosate use (see SI Appendix replaceN). This difference could result from differences in water treatment, measurement error in upstream-glyphosate exposure, or other meteorologic/geologic/hydrologic factors.

The US EPA’s current approval process for agricultural chemicals is rife with opportunities for regulated entities to affect the review’s outcome—including a reliance upon non-peer-reviewed, applicant-generated studies (ignoring conflicts of interest) (C. Benbrook, 2020; C. M. Benbrook, 2018; Sheppard, McGrew, & Fenske, 2020); selective reporting/testing (C. M. Benbrook, 2018, 2019);

flawed research designs (vom Saal et al., 2010), implementation (Mie, Rudén, & Grandjean, 2018), and analyses (Sheppard et al., 2020); limited post-approval monitoring (C. M. Benbrook, 2018); and a general lack of transparency/oversight (Davoren & Schiestl, 2018). (See *Regulatory oversight* in *SI Appendix A* for an extended discussion.) Uncertainty around glyphosate and GM health impacts has not slowed the spread of these technologies. Neither has this uncertainty led to substantive monitoring that would enable regulators or researchers to precisely estimate damages, disentangle exposure mechanisms, or understand heterogeneity in glyphosate’s damage function. Despite the relative dearth of data, consistent evidence is emerging that GM-spurred glyphosate intensification adversely affects health. Nevertheless, without further research, improved monitoring/data, and re-optimized policies, the public will likely continue to bear the health burden of glyphosate inefficiently and unequally.

3.3 Methods

3.3.1 Infant health data. We use the universe of births in the United States between 1990 and 2013 from the National Vital Statistics System (NVSS). The birth data contain information recorded on the birth certificates, including month and year of birth, sex, birthweight, APGAR score, live birth order, total birth order, whether the birth was a C-section, gestation length, and dummies for a battery of birth defects. The birth certificates also contain certain demographics of the mother and father—namely race, ethnicity, age, education, marital status, and residence status. We have access to the restricted (human-subjects exempt (*Data Sets Not Requiring IRB Review*, 2024)) versions of these files, which identify the county of birth and county of mother’s residence for all

births, compared to the publicly available data, which hides geographic identifiers for counties with less than 100,000 residents.

3.3.2 Pesticide use estimates. Our measure of glyphosate use comes from the United States Geological Survey’s Pesticide National Synthesis Project (Baker & Stone, 2015; Thelin & Stone, 2013). The USGS surveyed farmers to calculate pesticide use rates per acre of different crops planted at the crop reporting district level. They then multiply these usage rates by the total acreage of each crop within the crop reporting district to estimate the total amount of each pesticide used, measured in kilograms. Each pesticide has two estimates, high and low, where the high value assumes crop reporting districts with a missing usage rate for a pesticide applied the pesticide at the same rate as their neighbors on each crop. The low value assumes a missing usage rate for a pesticide means that farmers did not apply that pesticide. We use the high value throughout our analysis. Additionally, we normalize by the total area of each county, thus our measure of glyphosate and other pesticides are in kilograms of active ingredient per square kilometer (kg/km^2).

3.3.3 Attainable yield. We use the FAO-GAEZ attainable yield for soybeans, corn, and cotton to measure the suitability of a county for genetically modified crops (FAO & IIASA, 2021). These data assign potential yield values to one square kilometer pixels based on environmental factors such as soil type, slope, and climate. We aggregate the pixels to the county by taking the average of all pixels within a county. We take the difference between the high-input attainable yield and the low-input attainable yield to focus on the counties with the largest incentive to adopt GM crops. The underlying model calculates the high-input scenario assuming that farmers have access to modern technology for

crop management, including GM seeds. Whereas the model calculates the low-input scenario assuming more traditional farming methods. We aggregate the three GM crops by standardizing the attainable yield difference so that each crop has a mean of zero and a standard deviation of one. We then take the maximum across the standardized yield differences, then re-scale this average into national percentiles. SI Appendix B discusses the validity of the GAEZ data in light of new work comparing historical yields to GAEZ predictions (Rhode, 2024).

3.3.4 Additional data. We supplement our analysis with several additional data sources. First, classify counties as *rural* using the 2003 Rural-Urban Continuum codes from the US Department of Agriculture (USDA) (mapped in Fig. S1). A rural county is any non-metro county, where the USDA defines a metro county as, “broad labor-market areas that include central counties with one or more urban areas with populations of 50,000 or more people. They also include outlying counties that are economically tied to the core counties as measured by labor-force commuting” (*Documentation*, 2023).

We use workforce data from the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics. These data report the annual average number of employed and unemployed workers going back to 1990. We supplement these data with annual, county-level income and employment data split by farm and non-farm from the Bureau of Economic Analysis (BEA). They report these data going back to 1969 and source the data from the US Bureau of Labor Statistics and Internal Revenue Service.

Finally, we use annual county-level acreage and yield of various crops—including the main GM crops, corn, soy, and cotton—from the USDA National Agricultural Statistics Service (NASS). Additionally, we collect data on county-level

fertilizer use from the USGS Falcone (2021). We employ an interpolation process described in SI Appendix B to generate annual fertilizer estimates for each county.

3.3.5 Empirical strategy. We aim to estimate the causal effect of glyphosate exposure—induced by the roll-out of GM seeds—on infant health. To isolate exogenous variation in glyphosate use, we leverage temporal variation from the commercial release of GM crops in 1996 and spatial variation in GM adoption due to differences in environmental suitability for growing those crops.

Define a linear model for the effect of local glyphosate (β)

$$\text{Health}_{ijt} = \beta \text{GLY}_{jt} + \Gamma X_{ijt} + \alpha_j + \lambda_t + \varepsilon_{ijt} \quad (3.1)$$

for individual i , in county j , in year t . GLY_{jt} represents local glyphosate exposure, measured as the total mass of glyphosate sprayed per square kilometer of total area in county j in year t (kg/km²). X_{ijt} provides a vector of controls. α_j and λ_t are county and month-of-sample (e.g. January 2012) fixed effects. Estimating (3.1) with OLS is unlikely to identify the true effect of glyphosate on health due to endogeneity and measurement error.

The main endogeneity concern is that the adoption of GM technology and glyphosate may correlate with unobservable factors that affect perinatal health. To rectify the measurement and endogeneity issues, we use instruments—in a reduced-form DID and a 2SLS estimator—that isolate exogenous variation in local glyphosate exposure. Our instruments are the maximum percentile of attainable yield for GM crops for county j , denoted GM_j , interacted with dummy variables for each year (described in *Attainable yield*). The first stage in our 2SLS estimator is

$$\text{GLY}_{ijt} = \sum_{\tau \neq 1995} \gamma_{\tau} \text{GM}_j \times \mathbf{1}(t = \tau) + \Gamma X_{ijt} + \alpha_j + \lambda_t + \varepsilon_{ijt} \quad (3.2)$$

2SLS uses predictions from (3.2) in place of GLY_{jt} in (3.1) to estimate a causal effect of glyphosate on perinatal health.

While we have micro-data (birth-level) on birth outcomes, we do not have precisely measured glyphosate exposure leading to measurement error. We expect some mothers within a county are highly exposed to glyphosate while others are not. However, data limitations force us to assign the same level of exposure to all births within a county. This measurement concern relates to the ecological fallacy and likely leads us to underestimate the magnitude of glyphosate’s damages to an individual’s health. As described by Banzhaf, Ma, and Timmins Banzhaf et al. (2019b) “When measuring the correlation between pollution and demographics, the ‘ecological fallacy’ can arise when inferring relationships between individual units (like households) from larger, more aggregated units (like counties) that contain those units.”

3.3.6 Identification. Throughout the paper, we employ two complementary empirical approaches: a reduced-form DID and 2SLS. Both approaches use similar instruments that combine spatial variation in the suitability for GM crops with temporal variation from the introduction of glyphosate-tolerant seeds. While thus related, the two approaches rely upon different assumptions for causal identification.

3.3.6.1 Reduced-form DID. The reduced-form DID results rely upon a parallel-trends assumption: had GM seeds not been released, then the *relative difference* between more and less GM-crop-suitable counties (high- and low-attainable yield counties) would have remained constant in terms of infant health—conditional on fixed effects and controls.

Differential trends in higher and lower attainable-yield counties, prior to the release of GM crops, would violate this assumption. A single outlier year for higher- or lower-yield counties could also affect our results despite a smooth increase in glyphosate through time. To assess these concerns, we estimate the event-study model

$$\text{Health}_{ijt} = \sum_{\tau \neq 1995} \gamma_{\tau} \text{GM}_j \times \mathbb{1}(t = \tau) + \Gamma X_{ijt} + \alpha_j + \lambda_t + \varepsilon_{ijt}$$

interacting GM_j (continuous percentile of GM crop suitability) with year indicators. The measured effects (γ_{τ}) provide the difference in perinatal health between the highest and lowest attainable yield counties relative to their 1995 difference (the year before the GM-seed rollout and ensuing glyphosate intensification). County fixed effects absorb average differences between higher and lower yield counties. Consequently, an event study with no trend in the γ_{τ} prior to treatment (pre-1996) supports our identifying assumption for the reduced-form DID.

Event studies in Figs. 13 and S5 support this assumption: prior to 1996, the event studies are relatively flat and near zero—bearing no strong evidence of differential trends in infant health or (glyphosate intensity) related to counties’ GM suitability. After the GM-seed rollout higher-suitability areas accelerate in glyphosate intensity and decline in perinatal health outcomes.

The panel variation in our instrument is critical for identification, as it permits a comparison of pre-GM *versus* post-GM outcomes across higher- and lower-suitability areas. Without this panel variation (and the shock of GM seed’s rollout), identification would exclusively come from cross-sectional pesticides differences and could easily conflate health effects from glyphosate exposure with cross-sectional health differences that correlate with the geophysical determinants

of GM suitability. Instead, reduced-form estimates come from comparing the difference between higher- and lower-GM-suitability counties *after the GM rollout* to the difference between higher- and lower-GM-suitability counties *before the GM rollout*. Thus, even if the geophysical determinants of GM-crop suitability correlate with cross-sectional differences in health (e.g. healthcare access), our identification only relies upon more- and less-suitable areas following similar trajectories (i.e. parallel trends). Cross-sectional differences will be absorbed by our fixed effects. Again, the timing of the effects depicted in the event studies support the idea that *differences* between high- and low-GM-suitable counties were stable until to the introduction of GM seed and the glyphosate intensification. These points are also relevant for the 2SLS results discussed next.

3.3.6.2 Two-stage least squares. Our two-stage approach estimates (3.1) with two-stage least squares (2SLS), instrumenting glyphosate exposure with GM-crop suitability interacted with year indicators (Eq. 3.2). Interpreting these 2SLS estimates as the causal effect of glyphosate requires the instrument is relevant to glyphosate intensity, exogenous, and satisfies an exclusion restriction (Angrist & Pischke, 2009).

The *relevance* of the instrument is evident Fig. 13a: on average, after the GM rollout, glyphosate intensity increased more in higher-GM-crop-suitability counties than in lower-suitability counties.

The exogeneity of this instrument—and thus the variation it extracts—comes from the fact that its spatial variation (GM suitability) is a function of environmental factors unaffected by GM/glyphosate adoption and its temporal

variation comes from the arbitrary timing of GM seed’s rollout. The event studies discussed above support this exogeneity.¹

Interpreting 2SLS estimates as the causal effect of glyphosate hinges on the instrument satisfying an exclusion restriction. In our context, this exclusion restriction is that conditional on our controls, GM-crop suitability (interacted with year) only affected infant health through changes in glyphosate use. In other words: GM-crop suitability (interacted with year) should only affect infant health via increased glyphosate exposure, conditional on controls. If the rollout of GM technology induced other, non-glyphosate changes in more-GM-suitable counties *and* these changes had infant-health effects (after conditioning on covariates and fixed effects), then the exclusion restriction would be violated.

Non-glyphosate pesticides replaced by glyphosate potentially violate this exclusion restriction. Consequently, 2SLS without controls for non-glyphosate pesticides could be biased (likely downward) for glyphosate’s actual health effects—failing to exclude health improvements from reduced exposure to non-glyphosate pesticides. Fertilizer, employment, and income changes present similar potential to violate the exclusion restriction. Thus, we emphasize that the *glyphosate effect* is the appropriate estimand for 2SLS—directly controlling for potential exclusion-restriction violations. With these controls and the instrument’s exogeneity, we believe the exclusion-restriction is plausible.

Finally, settings with potentially heterogeneous treatment effects require an additional assumption: monotonicity. In our context, monotonicity requires that increasing attainable yield in a county would result in that county using

¹Table S1 compares summary statistics for high-yield counties, low-yield counties, and urban counties before the release of GM crops. Births in high- and low-yield rural counties were quite similar during this period—as were many other outcomes.

weakly more glyphosate—i.e. increasing suitability for GM crops would not reduce a county’s glyphosate intensity. This assumption appears quite reasonable in our setting, and we have no reason to believe that *decreasing* a county’s suitability for corn, soy, or cotton would increase glyphosate, all else equal.

3.3.7 Birth-outcome index. Because we estimate the impact of glyphosate on multiple birth-health outcomes, we are testing multiple hypotheses. Accordingly, the stated error rate of each hypothesis test may not accurately reflect the joint error rate. Following Currie, Greenstone, and Meckel (2017), to assuage multiple-hypothesis-testing concerns, we construct a single birth-health index from our five variables of interest—birthweight, gestation length and indicators for low birthweight ($<2500\text{g}$), very low birthweight ($<1500\text{g}$), and pre-term (≤ 37 weeks). (Because *C-section* functions as a falsification test, we exclude it from our index.) This single index provides a single test for composite health at birth upon which one may focus without concerns regarding multiple inference (M. L. Anderson, 2008). Following Currie, Greenstone, and Meckel, the index is a weighted average of the variables, where the weights use inverse-covariance weighting—capturing variables’ covariance and differences in their variances (Currie et al., 2017).

3.3.8 Predicting birthweight. Models like Equation (3.1) identify a single, *average* effect—as in Table 14. If individuals respond differently to glyphosate exposure, an average effect obscures this heterogeneity. Further, this heterogeneity is potentially critical to understanding the health impacts and policy remedies of the documented losses of birthweight—for example, if the lost birthweight came from lower-weight infants, as opposed to higher-weight infants. Thus, *who* lost weight could be key. We estimate heterogeneity in the effect of GM seeds and glyphosate on infant birthweight and gestation as a function of the

infant’s birthweight percentile—with two nuances. The first nuance centers on an issue of causal identification; the second is operational.

First, rather than using an infant’s *actual* birthweight percentile, we use the percentile of the infant’s *predicted* birthweight. We train a random forest to predict birthweight using features from the NVSS that describe the infant, mother, father, and birth location.² The training sample for this prediction is the universe births in the contiguous US before the mass introduction of GM crops and the ramp-up of glyphosate application (before 1996)—i.e., before our *treatment* began. We tune³ the model on 80% of the pre-1996 births using five-fold cross-validation. Finally, we train the selected (minimum RMSE) model in a five-fold pattern—ensuring each prediction comes from a model that has not seen the predicted individual. This hold-out approach in the prediction step and our predicted birthweight approach help avoid bias in the heterogeneity regressions. This bias could arise because birthweight—our heterogeneity dimension—is the outcome variable; conditioning on the outcome can introduce endogeneity (Elwert & Winship, 2014). Instead, we condition on *predicted* birthweight, which is a function of (1) the infant’s family’s observable features and (2) *other*, pre-treatment infants’ birthweights. Because our predictions are relatively accurate (predicted birthweight is, on average, quite close to actual birthweight as shown in Fig. S8), we can estimate how glyphosate differentially affects lower- and higher-birthweight infants without introducing bias from conditioning on the outcome.

²These features are the infant’s sex; the parents’ races, ethnicities, ages, and marital status; the mother’s education, residence status, plurality, and tobacco use; the birth location’s state; whether the birth location or mother’s home are in a rural county; whether the birth occurred in a facility; and month of year. When missing, we impute these features’ values.

³*Tuned*: the number of random features selected and the minimum observations in a terminal node.

The second nuance relates to the structure of the heterogeneous treatment effect. Because the shape of the heterogeneity is unknown, we take a semi-parametric approach that allows us to remain relatively agnostic. We split the sample using infants’ predicted birthweight percentiles (e.g., deciles) and then separately estimate the 2SLS model for each group. For example, Fig. 15 contains the estimated effect of glyphosate on birthweight for the first through tenth deciles. This bin-based approach is commonly applied to recover heterogeneous treatment effects, as it allows one to approximate arbitrary nonlinear functions without making stronger assumptions about functional form (Schlenker & Roberts, 2009). This approach still returns an *average* treatment effect estimate *within* each bin/group. Fig. S10 shows robustness to different bins in predicted birthweight across all of our main outcomes.

This approach—using ML-based counterfactuals in a causal-inference setting—relates to recent work by Burlig *et al.* Burlig, Knittel, Rapson, Reguant, and Wolfram (2020). While we use the predictions to estimate heterogeneity in our treatment effect, Burlig *et al.* take advantage of high-frequency electricity-consumption data for each individual (school)—using ML-based counterfactuals to identify the causal effect of energy-efficiency upgrades (and heterogeneity thereof). Another related approach—Athey and Wager’s causal forests Athey and Wager (2019)—provides a more generalized approach for using random forests to estimate conditional average treatment effects in higher dimensions, permitting interactions and nonlinearities. Because the heterogeneity of interest for this paper is only along a single dimension (expected birthweight), our bin-based approach provides sufficient flexibility.

3.3.9 Exposure to glyphosate sprayed upstream through water.

We use a spatial water model to estimate glyphosate exposure based on the amount of glyphosate sprayed upstream of each county using a methodology similar to that of Dias, Rocha, and Soares Dias et al. (2023). The HydroBASINS watershed shapes form the building blocks of this model (Lehner & Grill, 2013). We summarize the process here but leave the details in the SI Appendix N.

The amount of glyphosate that runs off into surface water will be affected by the erodibility of the soil, the slope of the land, and precipitation. We collect soil erodibility and slope data from the USGS gridded soil survey in each watershed, which we aggregate to the watershed level by taking the average over all grid cells within a watershed (Soil Survey Staff, 2021). We take the interaction of soil erodibility and slope and then convert that interaction into a percentile based on the distribution from all watersheds in the US. These data are static and do not change over time.

Next, we use gridded monthly precipitation from the PRISM climate group to capture whether there was potential for glyphosate to run off (PRISM Climate Group, 2022). We aggregate rainfall during the growing season (April through September) for each watershed and again convert it into a percentile from the distribution of all watershed-month-years.

We then map our county-level attainable yield percentile to watersheds and take the interaction between high erodibility, high precipitation, and attainable yield to create an instrument for upstream glyphosate use. We expect there to be effects from upstream spraying only when there is both high soil erodibility and high precipitation. We also estimate the effect of high attainable yield upstream without the interaction with soil erodibility and precipitation. The HydroBASINS

data allows for easy linking of upstream and downstream watersheds. In the linking process, we calculate the distance between two watersheds by tracing along centroids of all watersheds between those two watersheds. This measure allows us to aggregate variables over 50-kilometer distance bins upstream and downstream from each watershed. The downstream variables serve as a nice placebo test since we do not expect glyphosate sprayed downstream to affect infant health.

Once we estimate values upstream of each watershed, we must aggregate to the county level to analyze them with our health metrics. We take the weighted average of the upstream variables for all watersheds in a county, where the weights are the portion of the county’s population that lives within that watershed. We use population estimates for one square kilometer pixels from SEDAC to calculate the population weights for each watershed (CIESIN, 2017).

In addition to the binned approach described above, we use a machine-learning model to predict concentrations of glyphosate in surface water using a limited dataset of samples taken across the US. We utilize the geographic structure and physical characteristics of land, along with spatially disaggregated herbicide use, to predict downstream concentrations of glyphosate in surface water. We then regress perinatal health outcomes on these predictions on glyphosate and AMPA exposure from water. The details are in SI Appendix N.

CHAPTER IV

OPTIMAL SUBSIDIES FOR RESIDENTIAL SOLAR

This chapter is co-authored with Assistant Professor Mark Colas and is forthcoming in the *Journal of Political Economy Microeconomics*. Mark came to me with the original idea for the paper. I contributed to data collection, coding, and an alysis, particularly for the electricity grid part of the paper. I also contributed to the writing and revision of the paper.

4.1 Introduction

Governments around the world are more likely to subsidize technologies that lead to environmental benefits than they are to tax environmental damages (Borenstein, 2012). In the case of solar panels, these environmental benefits vary geographically. All else equal, the environmental benefits of solar panels are likely to be the largest in sunny areas and areas where high-polluting power plants would otherwise produce electricity. This spatial variation in environmental benefits suggests a rationale for spatially differentiated subsidies.

This paper uses a structural model of electricity production and household demand for rooftop solar panels to study how subsidies should optimally vary across space. Heterogeneous households across the country choose the number of solar panels to install, accounting for the installation cost, the lifetime value of the electricity produced and subsidies received, and the nonpecuniary costs and benefits of installation.¹ Households can also purchase electricity produced by a system of power plants. Individual power plants vary in the extent to which their production leads to environmental damages, their production capacity, and their location,

¹See Borenstein (2017) for a detailed discussion of the private benefits of solar installation in the case of California. See Borenstein and Bushnell (2022) for a discussion of how these factors shape the spatial distribution of solar panels.

which dictates how the grid transmits the plant’s electricity across geographic regions.

Residential solar installations reduce environmental damage by decreasing fossil-fuel power plants’ electricity production. Therefore, panels installed in areas with more sunlight have greater environmental benefits because they lead to larger decreases in electricity produced by these plants. The environmental benefits of solar panel installations also vary geographically because of differences in the distribution of technology employed by power plants across space. Panels installed in areas where environmentally unfriendly plants would otherwise produce electricity will be more beneficial than panels installed in areas with cleaner plants.

These environmental benefits are not internalized by the household, thus suggesting a role for government intervention. The primary tool currently employed by US policymakers to deal with this externality is a system of federal and state subsidies for solar panels. We use the model to solve for the nationally-optimal subsidies and quantify the benefits of switching from the current system of subsidies to the optimal subsidies.² Doing so requires understanding how solar panel installation rates and the damages associated with electricity production would change in response to alternative subsidy schemes. Therefore, our approach is to estimate a quantitative version of our model to calculate the optimal policies and the associated benefits.

Our primary data source is the DeepSolar Project (Yu, Wang, Majumdar, & Rajagopal, 2018), a dataset of the universe of residential solar panel installations in the contiguous US. Deepsolar uses a machine-learning framework to identify solar

²We focus on the optimal choice of subsidies for rooftop solar panels and do not allow for other types of government intervention, such as pricing the externality via a carbon tax. See Eichner and Runkel (2014) for an argument for why countries may choose to subsidize green energy production even when they have access to carbon taxes.

panel installations from satellite imagery. We supplement these data with data from Google Project Sunroof, another satellite-imagery-based dataset that provides information on solar irradiance across the US and on the number and size of rooftops suitable for solar panel installations. Combined, these two datasets provide the distribution and size of solar panel installations as well as solar irradiance and space suitable for solar panels across the US. We utilize these novel data sources to estimate the household component of the model via indirect inference, thereby providing the first estimated model of solar panel demand across the US. To facilitate identification, we utilize a *border discontinuity* approach, which exploits variation in subsidies on either side of state borders. Though sparsely parameterized, our household installation model matches the spatial distribution of installations well. We also show that our estimates are consistent with quasi-experimental evidence on the responsiveness of installations to solar rebates.

To model power plants, we develop a novel policy function approach that maps electricity demand and renewable production across the country to plant-level electricity production and emissions. Our approach allows for endogenous changes in power plants’ production profiles in response to electricity demand and renewable production over the day and year. We estimate these policy functions using Open Grid Emissions (OGE) data, which provide hourly production and emissions data covering nearly every power plant in the United States. We show that the estimated model matches the data’s temporal and spatial distribution of electricity generation. We translate these emissions into environmental damages using AP3, a state-of-the-art integrated air pollution model.

Our estimated model of solar panel demand and electricity production provides a framework to calculate the spatial distribution of installations,

environmental benefits of solar panels, and government cost of subsidies under counterfactual subsidy schemes. We first use this framework to solve for the nationally-optimal cost-neutral subsidy reforms and quantify the spatial misallocation caused by the current subsidy system. Our main result is that the current subsidy system leads to a severe misallocation of solar panels across space. Consider Washington, for example, a state where current subsidies are high even though sunlight is low and households receive marginal energy from relatively environmentally friendly power plants. We find that solar panels in Washington are over-subsidized by 100% relative to the nationally-optimal subsidy system, leading to 120% greater installations than optimal. Decreasing subsidies in Washington would lead to large decreases in fiscal costs with small decreases in environmental benefits. On the other hand, in West Virginia, where current subsidies are low and the environmental benefits of solar installations are high, we find that panels are under-subsidized by nearly 70%, leading to installations that are 60% lower than optimal. More generally, panels are under-allocated by roughly 30% in the Midwest and South and over-allocated by nearly 80% in the Northeast.

We find that the misallocation caused by the current system of subsidies leads to substantial environmental costs. Switching from the current subsidies to the welfare-maximizing subsidies leads to a 11.5% increase in aggregate environmental benefits—environmental damages decrease by approximately the same amount as a 11.5% increase in the productivity of every rooftop solar panel in the US. Switching to subsidies set by a planner aiming to minimize environmental damages rather than maximize welfare would lead to a 12.6% increase in aggregate environmental benefits.

Next, we calculate nationally-optimal subsidies when the government does not face an externally set budget constraint. The optimal unconstrained subsidies are significantly lower than current levels in the Northeast and West, but slightly higher than current subsidies in the Midwest and South. Nationally, total installations under the optimal subsidies are roughly 20% less than the current amount, leading to a decrease of 43 million dollars in annual environmental benefits relative to the current level. However, the accompanying 162 million dollar annual decrease in government costs thoroughly outweighs the decrease in environmental benefits. Put another way; the optimal unconstrained subsidies achieve over 80% of current environmental benefits at less than 60% the current cost. Our results suggest rooftop solar subsidies not only deviate from the optimum in how they vary across space but are also excessively generous in general.

Finally, we compare the effects of marginal subsidy changes around the current system of subsidies. We find large differences in the cost-effectiveness of subsidy increases across states. For example, the environmental damages offset per dollar of government funds associated with subsidy increases in West Virginia are over 3.5 times greater than the damages offset per dollar of subsidy increases in Washington. These results highlight that changes around the current system of subsidies could lead to decreases in both environmental damages and fiscal costs.

The remainder of the paper consists of various extensions and robustness checks. We analyze the sensitivity of our results to 1) alternative values of the marginal cost of public funds, 2) alternative specifications of household preferences and discounting, 3) accounting for line losses in transmitting electricity from plants to homes, 4) accounting for transmission constraints, 5) the introduction of improved electricity storage technology, and 6) changes in utility-scale renewable

electricity production. We find that the optimal system of subsidies remains qualitatively the same across these specifications. Quantitatively, our results suggest that optimal unconstrained subsidies will be even lower in the future as utility-scale renewable electricity production continues to expand.

It is important to caveat that residential solar subsidies may be associated with additional benefits that we do not model in this paper. Solar subsidies may increase innovation and technological growth via increased R&D or learning-by-doing (Bentham, Gillingham, & Sweeney, 2008; Bollinger & Gillingham, 2019; Gerarden, 2023; Goulder & Mathai, 2010).³ Further, numerous studies have found that solar panel suppliers hold considerable market power (Gillingham et al., 2016; Pless, Langheim, Machak, Hellow, & Sigrin, 2017; Pless & Van Bentham, 2019). Subsidies may help to alleviate the distortions caused by this market power (Dorsey, 2024). Solar installations may encourage additional solar installation via peer effects (Bollinger & Gillingham, 2012; Bollinger, Gillingham, Kirkpatrick, & Sexton, 2022) and may reduce reliance on the transmission grid, thereby reducing grid congestion and the need for costly investments in the transmission grid. Although not accounted for, these benefits would need to be substantial in order to justify current subsidy spending levels.

Further, the goal of this paper is to quantify *nationally-optimal* subsidies. Specifically, we focus on a social planner who maximizes national welfare, accounting for aggregate environmental damages and national fiscal costs. In reality, many solar subsidies are set by state governments who face state-level budget constraints and may be motivated to create local jobs or stimulate the local

³Learning by firms only constitutes a positive externality if it is non-appropriable and therefore not internalized by the firm. Bollinger and Gillingham (2019) find that non-appropriable learning-by-doing only leads to small learning spillovers for solar panel installers in California.

economy via “green multipliers” (Batini, Di Serio, Fragetta, Melina, & Waldron, 2022; Hasna, 2021; Popp, Vona, Marin, & Chen, 2020). In our analysis, we abstract away from these state-level incentives and constraints. Our results quantify the theoretically optimal distribution of funds across space and the geographic misallocation associated with the suboptimal current system of subsidies.

Our paper is most closely related to several papers which use model-based approaches to quantify the effectiveness of various types of subsidies on inducing solar panel installations (e.g., Burr (2014), De Groote and Verboven (2019), Langer and Lemoine (2022), Feger, Pavanini, and Radulescu (2022)). These papers use rich, dynamic models to study the trade-offs associated with various subsidy schemes. Of these, our paper is closest to Feger et al. (2022), which studies optimal installation subsidies and energy tariffs in a model with household energy consumption and solar panel demand.⁴ While all these papers focus on solar panel installations, we provide a framework that can additionally quantify the environmental benefits of solar panel installations, arguably the main reason these subsidies exist. As such, we are the first paper in this literature to quantify the trade-offs between the environmental benefits and fiscal costs of residential solar subsidies. We additionally contribute by quantifying the spatial misallocation due to current subsidy schemes through our analysis of how these subsidies should optimally vary across space. As discussed in the following paragraph, the reduced-form literature has emphasized the importance of spatial differences in the environmental benefits of solar panels. However, no quantitative research has incorporated these spatial differences in a study of optimal subsidy design.

⁴Feger et al. (2022) quantify the cost-minimizing and welfare-maximizing subsidy and tariff schemes subject to a network financing constraint and a solar energy target. They also allow the government to have a preference for redistribution, something we refrain from doing in our paper.

This paper is also related to a literature estimating the extent to which the marginal benefits of renewable energy investments vary geographically (e.g., Holland and Mansur (2008), Siler-Evans, Azevedo, Morgan, and Apt (2013), Graff Zivin, Kotchen, and Mansur (2014), Holland, Mansur, Muller, and Yates (2016), Millstein, Wiser, Bolinger, and Barbose (2017), Callaway, Fowle, and McCormick (2018), Holland, Mansur, Muller, and Yates (2020), Brown and O’Sullivan (2020)). In particular, this paper is similar to Borenstein and Bushnell (2022), who relate estimates of the marginal social cost of electricity production to the spatial distribution of solar panels in the US, and to Sexton, Kirkpatrick, Harris, and Muller (2021), and Lamp and Samano (2023), who study the marginal benefits of solar panel installations and calculate the benefits of reallocating panels across space.⁵ While these papers establish that the current spatial distribution of solar panel installations does not maximize environmental benefits, they do not model demand for solar panels and, therefore, do not quantify how installations respond to various subsidy schemes. The goal of our paper is to quantify the extent to which government policy causes this misallocation and solve for the system of subsidies that remedies this misallocation. Our contribution is, therefore, to build and estimate a structural model of solar panel demand and electricity production, which we use to quantify the effects of alternative subsidy schemes on the distribution of solar panel installations and calculate the resulting environmental benefits and fiscal costs. Specifically, we provide the first estimated model of solar panel demand across space in the US. We also develop a novel, tractable approach

⁵In the sustainability literature, Tibebe, Hittinger, Miao, and Williams (2021) calculate the subsidies which maximize environmental benefits less government cost at the national and state level. Their analysis does not account for household utility and therefore omits a key component of the social benefit of subsidies. They also do not model the household decision to install solar panels but instead model solar installation rates as following a normal distribution in the net present value of installation.

to modeling power plant production and the associated emissions over space and time. This approach involves directly modeling how individual plants' electricity production and emissions endogenously respond to changes in solar and other renewable production.

Finally, this paper is related to several empirical papers which estimate the elasticity of solar panel installations with respect to subsidies. We discuss these papers in Section 4.5.1.3. We use the estimates from these papers to evaluate the performance of our estimated model.

4.2 Model

We combine a model of household solar panel demand with an electricity production model. Households are distributed geographically across the United States, and states vary in their electricity prices, installation prices, and the set of subsidies for solar panels they offer. Within states, households vary in their local solar irradiance (sunlight), the amount of space they have for potential solar panels, and their preferences over solar panel installation. Households choose the number of solar panels to install, accounting for electricity and installation prices, solar panel subsidies, and their individual preferences for installing solar panels.

In addition to residential solar panels, central generation power plants produce electricity. Power plants differ in the extent to which their electricity production leads to environmental damages and their location, which determines how the electricity they produce is distributed across the country. Further, power plants face capacity and non-negativity constraints and vary in the order in which they are dispatched, implying that some power plants will only operate when demand is sufficiently high while others will operate even when demand is low.

4.2.1 Households. Households, indexed by i , are endowed with income y_i and $\bar{N}_i$ spaces they can potentially use for solar panels. Household i has access to a solar panel technology that can produce a stream of solar energy of $\{A_{it}\}_{t=0}^T$ over the lifetime of the panel for each panel they choose to install. In practice, we will think of t as indexing hours and set the lifespan of a solar panel to 25 years.⁶ We can think of this solar technology as reflecting the intermittent sunlight profile at a given household's residence, accounting for the depreciation of solar panel efficacy over time. Let j index the state in which the household lives.

Households choose whether or not to install solar panels, $m_i \in \{0, 1\}$, the number of panels conditional on installation, $N_i \in (0, \bar{N}_i]$, and how much electricity to consume each period. Specifically, households choose a sequence of electricity usage $\{e_{it}\}_{t=0}^T$, where e_{it} gives household i 's energy consumption in period t . We assume the household pays a constant price of p_j for all electricity purchased.⁷ Let r denote the real interest rate and let $e_i = \sum_{t=0}^T \frac{e_{it}}{(1+r)^t}$ denote the discounted sum of energy consumed, such that $p_j e_i$ gives the present discounted cost of electricity consumed.

If a household chooses to install solar panels, they pay the cost of installation of $p_j^{\text{Ins}}(N_i)$, which is a function of N_i , the number of panels they choose to install. The installation cost function $p_j^{\text{Ins}}(\cdot)$ can vary nonlinearly in

⁶This is a standard value of the average useful life of solar panels (see e.g., Y. Xu, Li, Tan, Peters, and Yang (2018), Chowdhury et al. (2020), or Sodhi, Banaszek, Magee, and Rivero-Hudec (2022)).

⁷We assume that electricity prices are constant over time. While electricity prices change over time, there is evidence that consumers do not correctly forecast the extent to which prices change over time and expect future prices to be similar to current prices (S. T. Anderson, Kellogg, & Sallee, 2013; Hughes & Podolefsky, 2015). Further, we assume these electricity prices are fixed across counterfactuals and, therefore, abstract from the general equilibrium effects of subsidies on electricity prices.

N_i and is allowed to vary by state j .⁸ Households can use electricity generated by solar panels to power their home or can sell it back to the grid. Assume, for now, that households can sell back to the grid at the price of electricity purchased, p_j , as is the case for households in states with net metering. We discuss how we model households without net metering in Online Appendix C.2.1.⁹ Letting $A_i = \sum_{t=0}^T \frac{A_{it}}{(1+r)^t}$ denote the discounted sum of electricity production, we can write the present discounted value of energy produced by each solar panel for household i as $p_j A_i$.

Households receive subsidies for solar installations. We allow for three types of solar panel subsidies that capture the majority of state and federal subsidies in the US. First, households can receive a cost-based subsidy s_j^{Cost} , which pays a percentage of the solar installation cost, similar to the federal investment tax credit. Second, households can receive a production-based subsidy of s_j^{kWh} for each kWh of electricity produced by their solar panels, similar to solar renewable energy certificates. Finally, we allow for a per-panel subsidy s_j^{Panel} , such as subsidies that pay per kilowatt of solar capacity installed.

We can thus write the household's budget constraint as

$$c_i + \underbrace{p_j (e_i - m_i N_i A_i)}_{\text{Net cost of electricity}} + \underbrace{m_i (1 - s_j^{\text{Cost}}) p_j^{\text{Ins}} (N_i)}_{\text{Net cost of installation}} = y_i + \underbrace{m_i N_i A_i s_j^{\text{kWh}}}_{\text{kWh Subsidy}} + \underbrace{m_i N_i s_j^{\text{Panel}}}_{\text{Per-Panel Subsidy}} \quad (4.1)$$

where c_i is consumption of the numeraire good.¹⁰

⁸We assume a nonlinear pricing function to allow for the possibility that there is a fixed cost associated with installing a positive number of panels.

⁹In 2017, 39 states mandated net metering policies. Idaho did not have a state net-metering policy, but each of the state's three investor-owned utilities had a net-metering policy. Five other states in our sample have distributed generation rules other than net metering.

¹⁰We can think of c_i and y_i as the present values of consumption and income over time, respectively.

Households have the following quasilinear utility function

$$c_i + \nu_i \left(\{e_{it}\}_{t=0}^T \right) + m_i \gamma_i (N_i),$$

where $\gamma_i (N_i)$ is a strictly concave function which gives the nonpecuniary benefit of adding N_i solar panels for household, and $\nu_i \left(\{e_{it}\}_{t=0}^T \right)$ is a function which gives the lifetime utility of electricity usage. The function $\gamma_i (\cdot)$ captures inconvenience costs and any other individual preferences for installing solar panels.

Note that the choice of electricity consumption does not depend on the household's choice to install panels. Thus, we can think of household optimization as a two-step process. First, the household chooses electricity use, $\{e_{it}^*\}_{t=0}^T$, then decides whether to install solar panels and the number of panels conditional on installation. In this second stage, we can rewrite the household's optimization problem as a choice of N_i and a discrete choice of m_i :

$$V_i = \max_{N_i, m_i \in \{0,1\}} m_i [\mu_{ij} (N_i) + \gamma_i (N_i)]. \quad (4.2)$$

where

$$\mu_{ij} (N_i) = \underbrace{N_i A_i (p_j + s_j^{\text{kWh}})}_{\text{Total electricity value}} - \underbrace{(1 - s_j^{\text{Cost}}) p_j^{\text{Ins}} (N_i)}_{\text{Net installation cost}} + \underbrace{N_i s_j^{\text{Panel}}}_{\text{Per-panel subsidy}} \quad (4.3)$$

denotes household i 's net monetary benefit of installing solar panels.¹¹ Let m_i^* denote the household's optimal installation choice and let N_i^* denote the optimal number of panels conditional on installation.

From equations (4.2) and (4.3), we can see that different types of subsidies will differ in the distribution of households they induce to install panels.

Households in sunny areas (high A_i) are more likely to respond to the production

¹¹Note that we have dropped a constant representing the household's utility from electricity use and costs but does not affect the decision to install solar panels.

subsidy s_j^{kWh} , while households in areas with high installation costs are more likely to respond to the cost subsidy s_j^{Cost} , for example.¹² As we show in Section 4.2.3, the planner chooses the nationally-optimal set of subsidies accounting for the fact that different households are marginal with respect to each type of subsidy.

Before proceeding, it is important to highlight that several features of the solar panel market are not included in our model. First, our model does not include a role for peer effects, which Bollinger and Gillingham (2012) and Bollinger et al. (2022) find can influence households' decisions to install solar panels.¹³ Second, De Groote and Verboven (2019) find that households heavily discount future benefits associated with solar installations. We instead assume households discount future subsidy payments according to the market discount rate. We consider a version of our model in which households more heavily discount future subsidy payments in Section 4.7.2. Third, we assume a partial equilibrium setting where installation costs are given exogenously and do not change in response to changes in subsidies. This partial equilibrium setting implies that increases in subsidies decrease the effective costs paid by consumers dollar for dollar. If the inverse supply curve for solar installations in our model were upward sloping, then increasing subsidies would also increase the equilibrium price charged by installers, thereby changing the price faced by the consumer.¹⁴

¹²Note that electricity prices and production subsidies have the same effect on the monetary benefits from installation for households in states with net metering.

¹³Our estimation strategy uses tract-level data to estimate the effects of solar subsidies on installations. Tract-level responses to subsidies include both the direct effect of subsidies on installations as well as the effect of within-tract peers on total installations. Therefore, the empirical elasticities we use in our analysis incorporate the effects of peer effects on installations. However, we do not model these peer effects directly.

¹⁴See Reguant (2019) for a model-based exploration of how various renewable energy policies differentially affect the prices faced by consumers.

Finally, our paper assumes prices and subsidies are fixed over time and, therefore, abstracts away from the substantial changes in residential solar subsidies that have occurred over the past 20 years in the United States (Barbose & Darghouth, 2019). We view our model of solar installations as capturing the *long-run* installation decisions of households who face a given set of subsidies. As such, our model is not appropriate for analyzing the short-run effects of policy changes nor for studying the transition path between two policy regimes. We discuss the implications of assuming constant prices for estimation and quantitative results in Section 4.7.8.

4.2.2 Electricity Production.

4.2.2.1 Background. Before proceeding to the model, we give a brief overview of electricity production in the US. The electricity sector in the US is highly regulated and does not operate like a traditional market. Each of the around 10,000 central generation power plants in the US is overseen by a balancing authority, an entity tasked with matching electricity supply and demand by managing production from individual plants and trading with other balancing authorities. Transmission of electricity between balancing authorities disproportionately occurs within larger regions called NERC *regions*, each constituting a relatively closed market of balancing authorities. Transmission across regions does occur, but this inter-regional transmission occurs almost exclusively within *interconnections*, a geographic unit larger than a region. There are three interconnections in the US: Eastern, Western, and Texas.

We can divide power plants into those that are *dispatchable* and those that are *nondispatchable*. Nondispatchable power sources are those whose output cannot be easily controlled in response to fluctuations in electricity demand and generally

produce when available, such as wind and solar. These energy sources are generally intermittent, meaning their productive capacity fluctuates over time in response to environmental factors, e.g., sunlight and wind. Nondispatchable power plants generally do not produce pollutants or greenhouse gases.

On the other hand, balancing authorities can control production by dispatchable power plants to satisfy electricity demand. The production profile of a given dispatchable plant is determined by its position in the dispatch curve—the order at which balancing authorities dispatch power plants to satisfy different electricity demand levels.¹⁵ This implies that a power plant’s production is not simply proportional to demand—some power plants operate continuously throughout the day while others only operate at peak levels of demand. As such, the set of marginal power plants, and therefore the marginal benefits to residential solar installation, vary geographically and within location as a function of demand that must be satisfied by dispatchable plants.

4.2.2.2 Model: Electricity Production. Within the model, three sources supply electricity: 1) residential solar, 2) nondispatchable plants, and 3) dispatchable plants.¹⁶ Nondispatchable units are assumed to operate at full capacity conditional on environmental conditions (e.g., sun and wind) and conditional on total demand exceeding the amount produced by these

¹⁵Power plants’ variable cost of production generally determine the dispatch curve. Power plants with the lowest variable costs (often nuclear and hydroelectric) typically satisfy low demand. Meanwhile, plants with higher variable costs (such as gas-fired plants) begin operating only when electricity demand is sufficiently high.

¹⁶We assume the distribution of power plants and the characteristics of the grid are exogenous. In reality, a large change in residential solar production may lead to the entry and exit of generators and changes in the organization of the electricity grid. In Section 4.7.6, we analyze the robustness of our results to alternative assumptions about the distribution of power plants. See Holland, Mansur, and Yates (2022) for a model which includes endogenous entry and exit of generators and storage capacity. See Arkolakis and Walsh (2022) for a model with endogenous grid formation.

nondispatchable generators.¹⁷ Therefore, as long as demand exceeds the amount produced by nondispatchable sources, the production by these power plants is independent of demand and production by other plants. Alternatively, the production by dispatchable units depends on excess demand remaining after production by residential solar and nondispatchable plants.

Residential Solar and Nondispatchable Plants Let R index NERC regions.¹⁸ Total residential solar production in region R in a given hour t is the sum of energy produced by residential solar panels, $E_{Rt}^{\text{Solar}} = \int_{i \in I_R} m_i^* N_i^* A_{it} di$, where I_R is the set of households who reside in region R . Similarly, total production by nondispatchable plants in region R in time t is given by $E_{Rt}^{\text{NonD}} = \sum_{k \in K_R} y_{kt}^{\text{NonD}}$, where y_{kt}^{NonD} denotes electricity production by nondispatchable power plant k in time t , and K_R is the set of nondispatchable plants in region R . y_{kt}^{NonD} is allowed to vary fully by power plant k and time t , reflecting differences in environmental factors across plants and over time, and we assume it is independent of demand and production from other power plants.

Dispatchable Plants To capture the centralized manner by which balancing authorities dispatch power plants to satisfy electricity demand, we model dispatchable plants' behavior via policy functions that map excess demand to plant-level production. Let Load_{Rt} denote the total electricity demand in region R in time t , and let $\text{ELoad}_{Rt} = \text{Load}_{Rt} - E_{Rt}^{\text{NonD}} - E_{Rt}^{\text{Solar}}$ give the electricity demand

¹⁷In this case, we assume production of nondispatchable plants is curtailed such that supply does not exceed demand.

¹⁸We will assume 7 NERC regions in our quantitative analysis. Officially, the North American Electric Reliability Corporation (NERC) divides the US into 6 regional entities. Following Holland et al. (2016), we separate California from the WECC region, leaving us with 7 regions. We discuss how we define the regions in Online Appendix C.1.7.

in region R that is not satisfied by residential solar and nondispatchable plants.

We write production by dispatchable plants as a reduced-form function of excess demand across regions, subject to non-negativity and capacity constraints. Letting y_{kt}^{Disp} denote production by dispatchable plant k in time t , we specify

$$y_{kt}^{\text{Disp}} = \begin{cases} 0 & \text{if } f_k(\text{ELoad}_t, \varepsilon_{kt}) \leq 0 \\ f_k(\text{ELoad}_t, \varepsilon_{kt}) & \text{if } 0 < f_k(\text{ELoad}_t, \varepsilon_{kt}) < \bar{y}_k, \\ \bar{y}_k & \text{if } f_k(\text{ELoad}_t, \varepsilon_{kt}) \geq \bar{y}_k \end{cases}, \quad (4.4)$$

where $\bar{y}_k$ is power plant k 's nameplate capacity, the maximum productive capacity of the plant, ELoad_t is the vector of excess loads in each region at time t , and $f_k(\text{ELoad}_t, \varepsilon_{kt})$ is a plant-specific function of excess load across regions and a cost shifter ε_{kt} .¹⁹ We allow the function $f_k(\cdot)$ to differ across plants to reflect heterogeneity in the order in which plants are dispatched. We also allow $f_k(\cdot)$ to depend not only on excess load in a power plant's own region but potentially to depend on excess load across other regions as well. This dependence reflects that electricity can be transmitted across regions in response to excess demand.

Intuitively, y_{kt}^{Disp} captures how production by an individual power plant k in a given hour t responds to fluctuations in electricity demand and nondispatchable production across the grid. For example, as the sun goes down and solar production decreases, excess load will increase across the country, particularly in regions heavily reliant on solar energy. y_{kt}^{Disp} tells us how individual power plants are dispatched to match these increases in excess load.

¹⁹We assume that the plant's policy function depends only on the current excess demand levels. Hypothetically, production could also depend on previous electricity demand and production if, for example, the grid can store significant amounts of electricity over time or if plants face significant ramping constraints. We can accommodate this extension by allowing the function $f_k(\cdot)$ to depend on lagged values of excess demand, or on lagged production levels of the individual plant.

4.2.2.3 Damages. Let $d_k(y_{kt}^{\text{Disp}})$ be a function that maps dispatchable power plant k 's electricity production in time t to the total environmental damages associated with the plant's emissions of greenhouse gases and air pollutants. We assume these damage functions are given exogenously and, therefore, rule out the possibility that power plants endogenously switch fuel types in response to policy changes. Let $D_t(\text{ELoad}_t) = \sum_k d_{kt}(y_{kt}^{\text{Disp}})$ denote total damages from all power plants in time t and let $D(\text{ELoad}) = \sum_{t=0}^T \frac{D_t(\text{ELoad}_t)}{(1+r)^t}$ denote the net present value of all damages over time, where ELoad gives the excess load across all region and time periods.

The external benefit of a marginal solar panel installed by household i equals the damages offset over the panel's lifetime. We write this as

$$\Delta D_i(\text{ELoad}) \equiv \left| \frac{\partial D(\text{ELoad})}{\partial N_i} \right| = \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\cdot)}{\partial E_{Rt}^{\text{Solar}}} \right|,$$

the present discounted sum of the product of A_{it} , the electricity produced by the panel in any given period, and the absolute value of $\frac{\partial D_t(\cdot)}{\partial E_{Rt}^{\text{Solar}}}$, the marginal damages associated with nondispatchable plant production.

4.2.3 Government's Problem and Nationally-Optimal Subsidies.

The government chooses subsidies to maximize the sum of total utility minus total environmental damages subject to an externally set budget constraint.²⁰ We consider a government who does not face a budget constraint in Section 4.6.4. To ease up on notation, let $s_{ij} = s_j^{\text{Panel}} N_i^* + s_j^{\text{kWh}} A_i N_i^* + s_j^{\text{Cost}} p_j^{\text{Ins}}(N_i^*)$ denote the total

²⁰We are assuming that these are the only policy instruments the government can access. The government is restricted to not price the externality directly, as in Pigou (1920). Changes in subsidies could also change firm profits. We assume that the government does not value profits of utility companies or solar panel installation companies. In reality, utility companies operate as regulated monopolies, where profits are directly limited. Profits of solar panel installation firms not entering the government's objective is also consistent with a model in which the price of installation is always equal to the marginal cost of an installation. We analyze a planner who minimizes environmental damages in Section 4.6.2.

subsidy paid to household i conditional on installation. Further, let $\frac{\partial N_i^*}{\partial s_j^\theta}$ give the derivative of solar panels installed by household i with respect to a given subsidy type $\theta \in \{\text{kWh, Panel, Cost}\}$, and let $\vec{m}_i^\theta$ indicate the household i is on the margin of installing a positive number of panels with respect to a θ subsidy, meaning the household does not install given the current subsidies but would install in response to a small increase in the given subsidy. Finally, let $M_j = \int_{i \in I_j} m_i di$ denote the total number of households who install solar panels in state j , where I_j is the set of households in state j .

The government maximizes the sum of utility less environmental damages,

$$\underbrace{\int_i V_i di}_{\text{Utility}} - \underbrace{D(\text{ELoad})}_{\text{Damages}}. \quad (4.5)$$

The government faces the constraint that the sum of subsidies cannot exceed an externally set constraint

$$\underbrace{\sum_j \int_{i \in I_j} s_{ij} m_i^* di}_{\text{Government Cost}} \leq G, \quad (4.6)$$

where G is the maximum amount the government can spend on subsidies. We can reformulate the government's objective function as the Lagrangian

$$W = \underbrace{\int_i V_i di}_{\text{Utility}} - \underbrace{D(\text{ELoad})}_{\text{Damages}} - \lambda \left(\underbrace{\sum_j \int_{i \in I_j} s_{ij} m_i^* di}_{\text{Government Cost}} - G \right), \quad (4.7)$$

where λ is equal to the marginal cost of public funds. In practice, we will set G to the present discounted value of the national cost of solar subsidies, given the current system of subsidies.

The nationally-optimal system of subsidies must satisfy $\frac{\partial W}{\partial s_j^\theta} = 0$ for each type of subsidy in each state, which implies

$$\underbrace{\frac{\partial M_j}{\partial s_j^\theta} \times \left(\overrightarrow{\Delta D}_j^{\theta, \text{ext}} - \lambda \overrightarrow{s}_j^{\theta, \text{ext}} \right)}_{\text{Extensive Margin}} + \underbrace{\frac{\partial N_j}{\partial s_j^\theta} \bigg|_{M_j^{\text{st}}} \times \left(\overrightarrow{\Delta D}_j^{\theta, \text{int}} - \lambda \overrightarrow{\frac{\partial s}{\partial N}}^{\theta, \text{int}} \right)}_{\text{Intensive Margin}} + \underbrace{(1 - \lambda) M_j \overrightarrow{\frac{\partial s_{ij}}{\partial s_j^\theta}}}_{\text{Mechanical Effect}} = 0. \quad (4.8)$$

We provide a derivation for equation (4.8) in Online Appendix C.2.3 and provide definitions for each individual object in the upcoming text. The first term (“Extensive Margin”) captures the trade-off between environmental benefits and fiscal costs associated with households who are *additional* with respect to a small subsidy increase: the households who currently do not install any solar panels but would install solar panels in response to a slight increase in a given subsidy s_j^θ . The term

$$\frac{\partial M_j}{\partial s_j^\theta} = \int_{i \in I_j} \overrightarrow{m}_i^\theta di$$

gives the number of households on the margin of installing solar panels with respect to a given subsidy type s_j^θ . These additional installations lead to a societal benefit by reducing environmental damages. The average damages offset across additional installer households is denoted as $\overrightarrow{\Delta D}_j^{\theta, \text{ext}}$ and is formally given by

$$\overrightarrow{\Delta D}_j^{\theta, \text{ext}} = \frac{\int_{i \in I_j} \Delta D_i (\text{ELoad}^{\text{SB}}) N_i^* \overrightarrow{m}_i^\theta di}{\int_{i \in I_j} \overrightarrow{m}_i^\theta di},$$

where ELoad^{SB} is the excess load under the optimal system of subsidies. These additional installations also receive subsidies and thus are associated with a fiscal cost. We denote the average cost associated with a marginal installation household as $\overrightarrow{s}_j^{\theta, \text{ext}}$, formally written as

$$\overrightarrow{s}_j^{\theta, \text{ext}} = \frac{\int_{i \in I_j} s_{ij} \overrightarrow{m}_i^\theta di}{\int_{i \in I_j} \overrightarrow{m}_i^\theta di}.$$

The second term of equation (4.8) (“Intensive Margin”) captures the environmental-fiscal trade-offs associated with intensive margin adjustment: increases in the number of panels purchased for households who already choose to install a positive number of panels. The term $\left. \frac{\partial N_j^{\text{st}}}{\partial s_j^\theta} \right|_{M_j^{\text{st}}}$ gives the total increase in panels associated with an increase in a given subsidy, holding the set of households who install solar panels constant, which we write as

$$\left. \frac{\partial N_j^{\text{st}}}{\partial s_j^\theta} \right|_{M_j^{\text{st}}} = \int_{i \in I_j} m_i^* \frac{\partial N_i^*}{\partial s_j^\theta} di.$$

The terms $\overrightarrow{\Delta D_j^{\theta, \text{int}}}$ and $\overrightarrow{\frac{\partial s}{\partial N} \theta, \text{int}}$ give the average damages offset and the average fiscal cost, respectively, associated with these additional panels.²¹ Taken together, these first two terms show that the government will optimally increase subsidies which induce a greater number of additional installations and additional panels from households associated with significant environmental benefits and for whom fiscal costs are low.

The final term (“Mechanical Effect”) captures the effects of increasing subsidies for the *non-additional* households: the households who already choose to install solar panels and thus receive a larger subsidy from the government. The total size of this transfer is the total number of panels installed in state j , M_j , multiplied by the average increase in subsidy for households who have installations, $\overline{\frac{\partial s_{ij}}{\partial s_j^\theta}}$, holding installations and the number of panels constant.²² Each dollar transferred to these non-additional households increases welfare by $(1 - \lambda)$,

²¹These are formally given by $\overrightarrow{\Delta D_j^{\theta, \text{int}}} = \frac{\int_{i \in I_j} \Delta D_i (\text{ELoad}^{\text{SB}}) m_i^* \frac{\partial N_i^*}{\partial s_j^\theta} di}{\int_{i \in I_j} m_i \frac{\partial N_i^*}{\partial s_j^\theta} di}$ and $\overrightarrow{\frac{\partial s}{\partial N} \theta, \text{int}} = \frac{\int_{i \in I_j} \frac{\partial s_{ij}}{\partial N_i} m_i \frac{\partial N_i^*}{\partial s_j^\theta} di}{\int_{i \in I_j} m_i \frac{\partial N_i^*}{\partial s_j^\theta} di}$.

²²Formally this is $\overline{\frac{\partial s_{ij}}{\partial s_j^\theta}} = \frac{\int_i m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta} di}{\int_i m_i^* di}$.

which reflects the increase in household utility less the decrease in government funds.²³ In summary, equation (4.8) measures the effects of subsidy changes on welfare, accounting for environmental benefits, fiscal cost, and household utility.

However, solving for nationally-optimal subsidies requires more structure on the problem. While it may be possible to calculate the environmental benefits of marginal solar panel installations given the current distribution of solar panel installations, to solve for the optimal subsidies, we need to know how marginal benefits change in response to different subsidy schemes. Given that power plants' production plans are nonlinear, the marginal damages evaluated at current installation levels will differ from those at the optimum. Further, the optimal subsidies characterized by equation (4.8) depend not only on marginal damages, but also on the number of non-additional households and the number of households on the margin of installation with respect to various types of subsidies. Like the marginal damages, both of these objects are a function of the system of subsidies.

Therefore, our approach is to estimate a fully specified version of our model, and then use that model to quantify the system of nationally-optimal subsidies. Further, the structural model allows us to quantify the effects of alternative subsidy schemes on the spatial distribution of rooftop solar and the resulting environmental benefits and fiscal costs.

4.3 Quantitative Model

4.3.1 Household preferences. Let ℓ denote the census tract in which a household lives. We assume that the nonpecuniary utility of installation

²³Note that the utility of additional households does not show up in equation (4.8) since there is no first-order welfare effect for households who are additional with respect to a marginal subsidy increase (i.e., households who choose to install solar panels in response to the increase in subsidies). This result comes from the envelope theorem. See also Colas, Findeisen, and Sachs (2021) for a discussion of the roles played by marginal and inframarginal agents in the first-order effects of targeted subsidy increases.

of installing N_i panels, $\gamma_i(N_i)$, is given by a polynomial term in N_i , a term that captures differences in the benefits of installation across demographic groups, and an idiosyncratic term. Specifically, we parameterize the nonpecuniary value of installation as

$$\gamma_i(N_i) = \underbrace{\gamma_0 + \gamma_{1N}N_i + \gamma_{2N}N_i^2}_{\text{Polynomial in } N_i} + \underbrace{\gamma_{dem}X_\ell}_{\text{Local Demographics}} + \underbrace{\sigma\epsilon_i}_{\text{Idiosyncratic}}$$

where γ_0 , γ_{1N} , and γ_{2N} are parameters, γ_{dem} is a vector of parameters, X_ℓ is a vector of demographic characteristics associated with the tract in which the household lives, and ϵ_i is a logit preference draw with scaling parameter σ . In practice, we specify $\gamma_{dem}X_\ell = \gamma_{Coll}X_\ell^{Coll} + \gamma_{Pol}X_\ell^{Pol}$, where X_ℓ^{Coll} the fraction of individuals in the census tract with a college education and X_ℓ^{Pol} is the percent of voters who voted Democrat in the 2016 presidential election.

Recall that the number of panels installed cannot exceed the space the household has available for panels, denoted by $\bar{N}_i$. The optimal number of panels conditional on installation is therefore given by

$$N_i^* = \min \left[\bar{N}_i, - \left(\frac{\frac{\partial \mu_{ij}}{\partial N_i} + \gamma_{1N}}{2\gamma_{2N}} \right) \right], \quad (4.9)$$

where $\frac{\partial \mu_{ij}}{\partial N_i}$ is the derivative of monetary benefit of installations (μ_{ij}) with respect to installation size (N_i). Loosely speaking, we can see that the ratio $\frac{\gamma_{1N}}{\gamma_{2N}}$ dictates the average size of installations while the parameter γ_{2N} dictates the degree to which N_i^* varies with subsidies.²⁴ For example, a smaller value of γ_{2N} in absolute value would imply that households are more responsive to subsidies along the intensive margin.

²⁴As we explain below, we parameterize $p_j^{\text{Ins}}(N_i)$ as a fixed cost plus a constant per-panel cost. This implies that $\frac{\partial \mu_{ij}}{\partial N_i}$ will be constant for a given household for $N_i > 0$.

Given that draws of ϵ_i are from a logit distribution, the probability that a household installs panels is equal to

$$\pi_i = \frac{\exp\left(\frac{\mu_{ij}(N_i^*) + \gamma_0 + \gamma_{1N}N_i^* + \gamma_{2N}N_i^{*2} + \gamma_{dem}X_\ell}{\sigma}\right)}{1 + \exp\left(\frac{\mu_{ij}(N_i^*) + \gamma_0 + \gamma_{1N}N_i^* + \gamma_{2N}N_i^{*2} + \gamma_{dem}X_\ell}{\sigma}\right)}. \quad (4.10)$$

The partial elasticity of installation probability with respect to monetary benefits is equal to

$$\frac{\partial \log(\pi_i)}{\partial \mu_{ij}(N_i^*)} = \frac{1}{\sigma} (1 - \pi_i). \quad (4.11)$$

Therefore, the parameter σ dictates the extent to which increases in subsidies will lead to increases in installations. A smaller value of σ implies that increases in subsidies will lead to larger increases in the number of installations.

4.3.2 Dispatchable Power Plant Production. Production by dispatchable power plant k in time t is given by equation (4.4). We assume the latent function $f_k(\text{ELoad}_t, \varepsilon_{kt})$ is a quadratic of excess load in each NERC region within plant k 's interconnection with an additive shifter denoted by ε_{kt} . Formally, letting $\mathbf{R}_k$ denote the set of NERC regions within the interconnection that contains plant k , we specify

$$f_k(\text{ELoad}_t, \varepsilon_{kt}) = \psi_k^0 + \sum_{R \in \mathbf{R}_k} (\psi_{Rk}^1 \text{ELoad}_{Rt} + \psi_{Rk}^2 \text{ELoad}_{Rt}^2) + \varepsilon_{kt}, \quad (4.12)$$

where ψ_k^0 is a constant term, ψ_{Rk}^1 is a parameter which dictates how $f_k(\cdot)$ changes in response to excess load in region R , ψ_{Rk}^2 is a parameter which dictates how $f_k(\cdot)$ responds to excess load squared in region R , and ε_{kt} is a normally distributed idiosyncratic term with a mean of 0 and a variance of σ_k^2 . Note that all the ψ_k parameters and σ_k^2 are plant-specific. We allow $f_k(\cdot)$ to depend on excess demand in all regions within an interconnection but not on excess demand in other

interconnections. This dependence reflects that electricity can be transmitted across regions within interconnections but is rarely transmitted across interconnections.²⁵ Our assumed functional form implies that y_{kt} is a Tobit function with latent variable f_k that is right censored at $\bar{y}_k$, plant k 's nameplate capacity, and left censored at 0.

This specification allows for relatively complex production patterns as a function of excess demand. The parameter ψ_k^0 dictates the values of excess demand over which a plant will produce electricity, allowing for the possibility that some plants will operate when excess demand is low while others will operate when excess demand is sufficiently high. For example, all else equal, plants for which ψ_k^0 takes a large negative value will only have positive electricity production when excess demand is very high, as is true of plants that tend to have a late position in the dispatch curve. Conditional on positive production, the parameters ψ_{Rk}^1 and ψ_{Rk}^2 will dictate the intensity at which the balancing authority dispatches a plant. Finally, this specification allows plants to differ in the extent to which their production is transmitted across regions. Some plants may predominantly transmit power within their own region, while others may transmit large amounts of power to other regions within an interconnection.

Further, while the latent function $f_k(\cdot)$ is assumed to be constant across time, we show in Section 4.5.2 that our model can replicate differences in dispatchable production over the day and year in response to fluctuations in nondispatchable production and electricity demand. In particular, our model can generate the ramping pattern of dispatchable generators through the afternoon as

²⁵We constrain $f_k(\cdot)$ such that the function is weakly increasing in excess load for all regions. That is, we set $f_k(\cdot)$ to its value at the inflection point if the function would otherwise be decreasing in excess load.

solar generation decreases and electricity demand increases. An important caveat to our approach is that these plant-specific policy functions are not invariant to changes in factors which may change the order in which plants are dispatched, such as changes in fuel costs or the introduction of a carbon tax. However, we do not expect first-order effects on these factors in the counterfactual subsidy schemes we investigate.

It is worth discussing how the specification of power plant production we develop here differs from the specifications used in Holland et al. (2016) and Sexton et al. (2021). Those papers estimate marginal emissions rates for individual power plants in which time-specific reduced-form coefficients capture all differences in emissions rates across time. These specifications, therefore, do not model how production levels endogenously respond to fluctuations in renewable production. As such, the estimated emissions rates for each power plant are constant conditional on time, and independent of the amount of solar electricity produced.

Since we aim to estimate marginal emissions both under current conditions and under significant changes to the distribution of residential solar panels, we require a different approach to modeling power plants. In our model, production varies flexibly in excess load and therefore is endogenous to both electricity demand and production from solar and other renewable sources. Thus, marginal emissions are not constant as a function of residential solar production. An additional benefit of our approach is that we identify the model’s parameters with excess load, which takes advantage of variation in both demand and production from nondispatchable units. The other models only leverage variation in demand.

4.3.3 Damages. The final piece of the model is determining damages from electricity production at power plant k , as described by the function $d_k(y_{kt})$.

We specify this function in two parts, first mapping electricity generation into emissions and then mapping emissions to damages. Both parts are plant-specific, capturing that damages from electricity production depend on a power plant's technology, location, and stack height. A power plant's technology dictates the extent to which electricity production leads to emissions, while a plant's location and stack height determine the extent to which emissions of local pollutants affect population centers.

Concretely, let $g \in \mathcal{G}$ index pollutants, where we assume the set $\mathcal{G}$ consists of the pollutants NOx, PM2.5, SO2, and CO2 equivalent (CO2e).²⁶ As shown by Holland, Kotchen, Mansur, and Yates (2022), power plants' marginal emissions rates tend to decline as utilization increases, where utilization is a plant's production level relative to its capacity. Therefore, to allow for emissions rates which vary across production levels, we specify emissions of each pollutant as a power-plant-specific linear spline in production with a slope that differs above and below power plant k 's median production. Letting y_k^{50} denote the median amount of power plant k 's production in the data conditional on positive production, we write power plant k 's emissions of pollutant g as

$$\text{Emis}_{gkt}(y_{kt}) = \begin{cases} \kappa_{gk}^1 y_{kt} + e_{gkt} & \text{if } (y_{kt} - y_k^{50}) < 0 \\ \kappa_{gk}^1 y_{kt} + \kappa_{gk}^2 (y_{kt} - y_k^{50}) + e_{gkt} & \text{if } (y_{kt} - y_k^{50}) \geq 0 \end{cases}. \quad (4.13)$$

Power plant k 's damages in time t are then given by $d_{kt}(y_{kt}) = \sum_{g \in \mathcal{G}} \delta_{gk} \text{Emis}_{gkt}(y_{kt})$, where δ_{gk} gives the marginal damages associated with emissions of g by power plant k , accounting for power plant k 's location and stack height.

²⁶CO2 equivalent includes emissions of other greenhouse gasses in addition to carbon dioxide, in particular, methane and nitrous oxide. These other GHGs are converted into a common global warming potential equal to that of one ton of carbon dioxide.

4.4 Data and Estimation

4.4.1 Data Sources. In this section, we give an overview of the main data sources we use in our analysis. Additional details on data sources and cleaning can be found in Online Appendix C.1.

Solar Panel Installations Our primary source for solar panel installations is the Deepsolar database (Yu et al., 2018), a database of solar panel installation in the contiguous US created by applying a deep-learning model for detecting solar panels on satellite imagery from the year 2016.²⁷ From Deepsolar, we use tract-level data on the total number of residential solar systems and on the total panel area covered by residential solar panels. Combining these two measurements gives us the average size of solar installations, which we use to infer the average number of panels per installation in each tract.

We supplement these data on solar installations with data from Google Project Sunroof (GPS), another dataset created by applying a machine-learning framework to satellite imagery. This dataset provides the distribution of rooftop sizes that are suitable for solar panel installation in each tract, which we use as the empirical analog of $\bar{N}_i$ within each tract for 56,940 census tracts in the US.²⁸

Rooftop Solar Production Next, we need data on $\{A_{it}\}_{t=0}^T$, the stream of electricity potentially produced by each panel installed by household i . For this,

²⁷Deepsolar is the first high-fidelity database of solar panel installations in the United States. Other solar panel databases rely on either self-reported data or surveys (e.g., Open Solar Project) or do not cover the entire contiguous US (e.g., Tracking the Sun). The machine-learning algorithm employed by Deepsolar is highly accurate, achieving a precision of 93% and a recall of 89% in residential areas.

²⁸These tracts include 90% of the 33 million square meters of residential solar panels in the Deepsolar database. The GPS data specifically provide the number of buildings in each tract with the potential for various installation size bins. We set $\bar{N}_i$ as the midpoint of the installation size bin for all buildings which fall in a given bin.

we combine data on yearly solar production potential from GPS with county-level time profiles of solar production from the National Renewable Energy Laboratory's System Advisor Model (SAM). Specifically, GPS provides measures of yearly kWh that can be produced by panels in a given tract, accounting for local weather conditions and shading. We set a household's yearly solar potential for newly installed panels as the mean household solar potential in the GPS data for the household's tract. We assume solar panel efficacy depreciates by a constant rate of 0.5% each year.²⁹ Next, we need to determine the distribution of solar production over each hour of the panel's lifetime. For this, we utilize SAM, which provides engineering estimates of electricity production with panel specifications and climate as the model inputs (Blair, Dobos, & Gilman, 2013). For each county, we calculate the fraction of yearly solar production produced at any given hour over the year. See Online Appendix C.1.4 for details. We multiply this fraction of energy produced each hour by a household's annual solar potential to calculate our measure of A_{it} , hourly electricity production for any hour t over the panel's lifetime.

Subsidies and Prices For subsidies, we rely on data from Sexton et al. (2021), who assemble data from the Database of State Incentives for Renewables & Efficiency to calculate state and federal subsidies in 2017. For the price of electricity, we use the average retail price of electricity as reported by the EIA.³⁰ We use a value of $r = 2\%$ for the real interest rate.

²⁹Jordan and Kurtz (2013) review the literature on photovoltaic degradation rates and find a median degradation rate of 0.5%.

³⁰These data can be downloaded at <https://www.eia.gov/electricity/state/>. See Ito (2014) for evidence that consumers respond to average, rather than marginal, electricity prices.

We estimate installation prices using data from Tracking the Sun, a project collecting data on solar panel installations by the Lawrence Berkeley National Lab. As Tracking the Sun only covers 25 states, we assume that all states within a given Census region share the same installation pricing function. Specifically, we assume that installation prices take the form $p_{R(j)}^{\text{Ins}}(N_i) = p_{R(j)}^{0,\text{Ins}} + N_i p_{R(j)}^{1,\text{Ins}}$, where $p_{R(j)}^{0,\text{Ins}}$ is a fixed cost and $p_{R(j)}^{1,\text{Ins}}$ is a per-panel cost, and $R(j)$ is the Census region containing state j . We present our estimates of the installation price functions and provide evidence that this linear pricing function is a good approximation of prices in the data in Online Appendix C.3.1.

Power Plants Our electricity generation data come from Open Grid Emissions (OGE), an open-source project aimed at creating high-quality electricity emissions data that is publicly available (Miller, Pease, Knowles, & Shi, 2022). These data combine commonly used electricity data sources, namely hourly electricity generation and emissions for generating units from the EPA’s Clean Air Markets Division (CAMD), monthly production and emissions from EIA form 923, and hourly balancing authority by fuel type electricity generation from EIA form 930.

We use their power-systems-level and plant-level data products from 2019.³¹ The power-systems-level data gives hourly electricity production for each balancing authority, broken out by fuel category, enabling us to calculate each region’s total hourly load. The plant-level data gives hourly electricity production and emissions for nearly 10,000 power plants. This coverage is the main innovation of the OGE data, as previously hourly emissions and production were only available

³¹We use data from 2019 as it is the first year available from OGE and thus closest to the Deepsolar data while also reflecting the modern electricity grid. The OGE methodology relies on the EIA form 930, which is only available starting in mid-2018.

for sufficiently large fossil-fuel plants included in the EPA’s CAMD data. The plants excluded from CAMD data account for nearly 30% of NO_x emissions, 8% of SO₂ emissions, and 7% of CO₂ emissions. We use the 4,625 dispatchable plants with positive, non-constant production in estimation, yielding over 40 million plant-hour observations after the cleaning process described in detail in Online Appendix C.1.7.

Damages To calculate damages associated with emissions, we utilize AP3, a state-of-the-art integrated assessment model that translates emissions from locations across the US into physical and economic damages. Specifically, AP3 uses a reduced-complexity air quality model to map emissions of local pollutants to an ambient concentration of air pollutants in each county in the US. The model then translates these ambient concentrations into damages, using estimates of the physical effects of pollution exposure from the literature and considering population distribution and vital statistics across counties.³² AP3 and its predecessors, APEEP and AP2, have been employed extensively in the environmental economics literature.³³ In addition to AP3, we use the social cost of carbon to quantify damages from greenhouse gas emissions.

4.4.2 Descriptive Patterns. Figure 17a shows how the generosity of subsidies varies across states under the current system of subsidies. We measure subsidy generosity as the present discounted value of subsidies an average household in each state would receive if they installed a 15-panel system, roughly

³²AP3 calculates damages as increased mortality risk from pollution exposure. For the value of mortality risk reduction, we use the EPA’s suggested value of \$7.4 million translated into 2014 dollars.

³³See, e.g., N. Z. Muller, Mendelsohn, and Nordhaus (2011), Holland et al. (2016), Shapiro and Walker (2020), Holland et al. (2020), Sexton et al. (2021), Cicala, Holland, Mansur, Muller, and Yates (2021), Holland, Mansur, Muller, and Yates (2021).




Figure 17. Expected subsidies and monetary benefit for a 15-panel system in each state.

Colors are scaled by the percentile of their respective value. Subsidies and monetary benefits are measured in 2014 dollars. See text for details.

the average size of installations in the data. There is considerable variation across states in the generosity of these subsidies. New Jersey delivers nearly 29 thousand dollars in subsidies, compared to seven states providing no additional funding, leading to under 6 thousand dollars in subsidies from the federal government. For comparison, the average cost of a solar installation in 2017 was roughly 24.5 thousand dollars. Therefore, the present discounted value of subsidies ranges from roughly one-quarter to more than the entire cost of an average installation. Figure C.4 in Online Appendix C.1.6 shows the state-level subsidy generosity separately for each of the three subsidy types. The majority of the value of subsidies comes from cost-based, rather than production-based or panel-based, subsidies.

In addition to subsidies, the monetary incentives to install panels vary geographically because of spatial differences in prices and sunlight. Figure 17b shows the monetary benefits associated with solar panel installations. Specifically, for every household within the model, we calculate the net present value of




Figure 18. Installed solar systems per 1000 individuals.

monetary benefits of installation, $\mu_{ij}(N_i^*)$, evaluated at $N_i^* = 15$. We then take the average monetary benefit over all households within a state. This total monetary benefit therefore measures the net present value of installing solar panels in a given state for the average household, taking into account local differences in solar irradiance, electricity and installation prices, and the set of local subsidies. The states with the highest monetary benefits are located in the Northeast, a region with high electricity prices and subsidies. Additionally, California has a high monetary value of installation, combining high electricity prices with high levels of solar irradiance. Meanwhile, several states in the Midwest and Mountain West have negative values, driven by lower subsidies, electricity prices, and solar radiation. Figure 18 shows installations per capita at the census tract level. We can see that installations are generally higher in areas with larger monetary benefits, such as most of the Northeast and California. Meanwhile, households in the Midwest, where there are relatively low subsidies, less sunlight, and low electricity prices, install few solar panels.




Figure 19. Border Discontinuities in Monetary Benefits and Installations.

Each graph plots estimated location-bin fixed effects from a regression of the variable in question on border and location-bin fixed effects. Monetary benefits are measured in 2014 dollars. Positive values on the X-axis represent households on the side of the border with more generous subsidies, and negative values on the X-axis indicate the side of the border with less generous subsidies.

In Appendix C.3.3, we investigate how the monetary benefits of solar installations correlate with the average size of solar installations. Across a number of specifications, we find that increased monetary benefits are associated with a statistically significant, but small in magnitude, increase in the average size of installations.

Border Discontinuities When we estimate our structural model, we use a border discontinuity approach, which compares installation rates on either side of state borders. Here, we present descriptive evidence on how the monetary benefits of installation and installation rates change as we cross the border from states with relatively less generous subsidies to states with more generous subsidies.

For this exercise, we limit our sample to tracts within 50 miles of state borders. Given the average solar irradiance in the border region, we calculate the

net present value of subsidies a household would receive for a 15-panel installation. Then, we classify the side of the border for which this hypothetical subsidy is higher as the “generous” side. Finally, we divide households into 5-mile-wide bins in locations relative to the state border. To examine how a given variable changes as we approach and cross state borders, we follow Bayer, Ferreira, and McMillan (2007) and regress the variable in question on state-border fixed effects and location-bin fixed effects. We plot these estimated location-bin fixed effects, which give the conditional average of the variable in question in a given location bin relative to the bin nearest to the border on the less generous subsidy side (the omitted location-bin fixed effect).

Figure 19a shows how the expected monetary benefit varies by the location relative to the border. As before, we measure monetary benefit as the average net present value of installing a 15-panel system in a given tract. Positive values on the X-axis represent households on the side of the border with more generous subsidies, and negative values on the X-axis indicate the side with less generous subsidies. By construction, tracts on the negative (left) side of the border have lower monetary benefits, with a jump of roughly \$8,000 as we move from the less generous side of the border to the more generous side. Figure 19b shows the results for log installation rates. There is a sharp increase in installation rates across the border—moving from the less generous side of state borders to the more generous side is associated with roughly a 50% increase in installations per capita.

As argued by Black (1999), Bayer et al. (2007) and D. S. Lee and Lemieux (2010), these discontinuities in installation rates are informative about the causal effect of subsidies on installations if other variables which may influence installation rates are continuous at state borders. One particular concern is that

household preferences for solar panels may be discontinuous at state borders. This discontinuity would occur if, for example, households with stronger preferences for solar panels sorted onto the side of the state border with more generous solar subsidies. Though we cannot measure preferences for solar panels directly, in Online Appendix C.3.4, we look for suggestive evidence for this type of sorting by plotting 1) the percent of households with college degrees, 2) the percent of voters who voted Democrat in the 2016 presidential election, and 3) average household income as a function of distance to state border. We find no evidence of sorting around state borders among these characteristics.

Another concern is that differences in other state policies, such as state tax rates, could drive these discontinuities. To assess this concern, we run border discontinuity regressions of log installations per capita on subsidies for solar panels, both with and without controls for various state taxes. Table 6 shows the results. For each column, we limit our sample to tracts within 10 miles of state borders and regress tract-level log installations per capita on the net present value of subsidies for a 15-panel system, tract-level household demographics, and state-border fixed effects.³⁴ Column (1) gives regression results without any controls for state and local taxes. Column (2) adds the state income tax rate, measured as the average state tax rate evaluated at a household income of 60,000 dollars, roughly the average income in the data. Column (3) adds sales tax, measured as the average state and local sales tax in a given state, and Column (4) adds the average property tax rate in the state.³⁵ The estimates from these border

³⁴We include fixed effects for all state pairs which share a border. For example, California-Oregon, California-Nevada, and California-Arizona are all included as separate fixed effects.

³⁵Sales tax data are taken from Walczak and Drenkard (2017). Property tax rates are from Sexton et al. (2021), who estimate property tax rates using data on real estate tax payments and property values from the American Community Survey.

	Dependent Variable:			
	Log Installations per Capita			
	(1)	(2)	(3)	(4)
NPV Subsidies (\$1000s)	0.0629*** (0.00920)	0.0755*** (0.00438)	0.0751*** (0.00570)	0.0826*** (0.00761)
Income Tax	NO	YES	YES	YES
Sales Tax	NO	NO	YES	YES
Property Tax	NO	NO	NO	YES
Distance Bandwidth	10	10	10	10
Observations	6,052	6,052	6,052	6,052
*** p<0.01, ** p<0.05, * p<0.1				

Table 6. Regression of log installations per capita on the net present value of subsidies for a 15-panel installation within 10 miles of state borders.

Subsidies are measured in thousands of 2014 dollars. State-clustered standard errors in parentheses. All regressions contain border fixed effects and tract-level demographic controls.

discontinuity regressions are similar across all specifications, but the parameter estimate slightly increases as we add state tax controls. The estimate in Column (4) suggests a \$1,000 increase in subsidies is associated with roughly an 8% increase in installations. We will control for these state tax variables in the border discontinuity regressions we use in our structural estimation procedure.

Moreover, apart from the policies categorized as subsidies, there exist additional state measures designed to promote solar adoption. If the prevalence of these policies correlates with subsidy generosity, this may lead us to overestimate the responsiveness of installations to subsidies. In Appendix C.3.5, we rerun these border discontinuity regressions with additional controls for other state-level programs aimed at increasing solar installations that our subsidy measures do not account for. Our results are robust to the addition of these additional controls.

4.4.3 Estimation.

Households We estimate the household installation component of the model via indirect inference. In essence, we first compute a set of “auxiliary models” that describe installation behavior in the data, and then simulate the structural model and calculate the auxiliary models with simulated data. We choose the six structural parameters $\sigma, \gamma_0, \gamma_{Coll}, \gamma_{Pol}, \gamma_{1N}$ and γ_{2N} such that the auxiliary models computed from the model are as close as possible to those from the data.

Our first auxiliary model is a border discontinuity regression which measures how installation rates change as we cross state borders.³⁶ Formally, we run the following regression using only tracts within 10 miles of state borders:

$$\log M_\ell = \alpha \bar{\mu}_\ell + \beta Z_\ell + \theta_{border(\ell)} + \varepsilon_\ell,$$

where M_ℓ is the total number of solar installations in tract ℓ , $\bar{\mu}_\ell$ is the average monetary benefit of installing a 15-panel system in tract ℓ , β is a vector of regression coefficients, Z_ℓ is a vector of tract-level demographic and tax rate controls, and $\theta_{border(\ell)}$ is a state-border fixed effect.³⁷ The parameter α measures the relationship between installation rates and monetary benefits in narrow bandwidths around state borders, controlling for border fixed effects. We target the coefficient α as an auxiliary model parameter.³⁸ As argued above, α is informative of the causal effect of monetary benefits on installation, given that household

³⁶Similar border discontinuity approaches have been used frequently in the environmental literature (see, e.g., Ito (2014), Feger et al. (2022), Hughes and Podolefsky (2015), or Rubin and Auffhammer (2023)).

³⁷We again include fixed effects for all state pairs which share a border. The vector of controls Z_ℓ includes the tract-level college completion percentage and percent of voters who voted Democrat in the 2016 presidential election, the state income tax rate, the sales tax rate, and the average property tax rate.

³⁸Our estimates of α are not sensitive to the size of the bandwidth around state borders, nor the inclusion of demographic controls. We obtain similar estimates of structural parameters when we instead include a regression of log installations on subsidy levels rather than total monetary benefits.

characteristics potentially correlated with preferences for solar do not exhibit a discontinuity at state borders. Additionally, we also target 1) log installations per household in each census tract, and 2) the average number of panels per array in each census tract.

The six structural parameters of interest are well identified. As we can see from (4.11), the parameter σ dictates the extent to which installations increase with monetary benefits. This parameter is thus identified by the coefficient α from the border discontinuity regression. Variation in demographics across tracts then jointly identifies γ_0 , γ_{Coll} , and γ_{Pol} . Finally, the average number of panels in each array and how the size of arrays varies across cities identify γ_{1N} and γ_{2N} .

Dispatchable Power Plants We estimate the power-plant-specific policy functions described by equations (4.4) and (4.12) via maximum likelihood. We provide the likelihood function and additional details in Online Appendix C.2.2. Variation over time in both electricity demand and production by nondispatchable plants creates variation in excess loads across regions that identifies the parameters of the plant-specific policy functions.

Damages We estimate damages by combining power-plant level emissions data from EPA’s Clean Air Markets Division with estimates of marginal damages from AP3. We estimate the damages given by equation (4.13) via ordinary least squares using power-plant level emissions data from OGE. To translate these emissions into damages, we need an estimate of δ_{gk} , the marginal damages associated with emissions of pollutant g by power plant k . The AP3 model calculates the marginal damages associated with local pollutants emitted from every county in the United

		Estimate	Standard Error
Dispersion of Idiosyncratic Utility	σ	11.6	0.5
Percent College	γ_{Coll}	7.5	0.6
Percent Democrat	γ_{Pol}	17.6	1.7
Constant	γ_0	-308.4	13.2
Number of Panels	γ_{1N}	34.1	1.6
Number of Panels Squared	γ_{2N}	-1.1	0.1

Table 7. Parameter estimates for household utility function.

Standard errors calculated via bootstrapping. “Percent Democrat” refers to the percent of voters who voted Democrat in the 2016 presidential election.

States for varying stack heights. We, therefore, calculate δ_{gk} by matching power plants to their corresponding county and stack height in the AP3 model.

We assume a social cost of carbon of 185 dollars per ton of CO22, based on the mean estimate from Rennert et al. (2022).³⁹

4.5 Estimation Results and Model Fit

4.5.1 Households.

4.5.1.1 Parameter Estimates. Table 7 displays the estimates of parameters governing the household utility function. The nonpecuniary value of installations is increasing in average local education and in the percent of voters who voted Democrat in the 2016 presidential election. The final two parameter estimates, which dictate utility as a function of installation size, imply that the optimal size of an installation is increasing in monetary benefits, but only marginally so: a \$1000 increase in the monetary benefit associated with installing

³⁹Rennert et al. (2022) is the same methodology and results on which US EPA based its December 2023 guidance for the social cost of greenhouse gasses. They report estimates in 2020 US Dollars. We convert this to 2014 dollars for consistency with the rest of our analysis. We do not account for the environmental damages associated with producing and disposing of solar panels. These costs are small relative to the environmental benefits of power produced by a solar panel (Heath & Mann, 2012).

an additional panel leads to only a $\left| \frac{1}{2 \times (-1.1)} \right| \approx 0.45$ increase in the optimal number of panels.⁴⁰

To get a better sense of what the parameter estimates imply for installation probabilities, recall that the partial elasticity of installation probability with respect to monetary benefits is approximately equal to $\frac{1}{\sigma}$.⁴¹ Given that we measure monetary values in thousands of dollars, our estimate of $\sigma = 11.6$ implies that a thousand dollar increase in the monetary value of installation leads to approximately a $\frac{1}{11.6} \approx 9$ percent increase in the number of installations.

4.5.1.2 Model Fit (Installations). Figure 20 assesses model fit with regard to solar installations. Figure 20a shows the relationship between tract-level log installations per household and the monetary benefits of installation in the data and simulation. We calculate the lifetime monetary benefits of installation as the net present value of installing a 15-panel array in each census tract. We can see that in both the data and simulations, installations are strongly increasing in monetary incentives.⁴² Subfigures 20b and 20c show the relationship between installations and the percentage of households with a college education, and the percent of voters who voted Democrat in the 2016 presidential election. The fit is quite good in both dimensions.

⁴⁰Recall from equation (4.9) that the optimal number of panels is given by $N_i^* = \min \left[\bar{N}_i, - \left(\frac{\frac{\partial \mu_{ij}}{\partial N_i} + \gamma_{1N}}{2\gamma_{2N}} \right) \right]$. We estimate $\gamma_{2N} = -1.1$ and increasing the per-panel subsidy by \$1000 increases $\frac{\partial \mu_{ij}}{\partial N_i}$ by 1.

⁴¹Differentiating equation (4.10) yields $\frac{\partial \log(\pi_i)}{\partial \mu_{ij}(N_i^*)} = \frac{1}{\sigma} (1 - \pi_i)$. The average value of π_i in our dataset is about 0.02.

⁴²The slight decrease in installation rates for the highest monetary values in the data reflects that Massachusetts and New Jersey have very generous subsidies and high electricity prices, but installation rates are lower than in states such as California and Arizona, which have less generous subsidies.




Figure 20. Household solar installation model fit.

Panels (a), (b), and (c) show local nonparametric fit of tract-level log installation per household in the data (red dotted line) and simulations (solid blue line) on (a) the total monetary benefit of installing 15 solar panels, (b) the percent of households with a college degree, and (c) the percent of voters who voted Democrat in the 2016 presidential election. Monetary benefits are measured in 2014 dollars. Panel (d) plots estimated location-bin fixed effects from a regression of log panels per household on state tax rates and location-bin and state-border fixed effects. Positive (negative) values on the X-axis represent households on the side of the border with more (less) generous subsidies. The dots represent estimates of location-bin fixed effects from the data (red hollow dots) and the simulations (blue dots).

Subfigure 20d examines how log panels per household change as we cross state borders. We first divide households into 10-mile-wide bins in location relative to the border using the procedure described in Section 4.4.2. We then regress tract-level log panels per household on location-bin fixed effects, controlling for border fixed effects and state tax rates. Finally, we plot the estimated distance-bin fixed effects for both the data (blue dots) and our simulated model (red dots). Positive values on the X-axis represent households on the side of the border with more generous subsidies, and negative values on the X-axis indicate the side with less generous subsidies. Both data series show a similar “jump” as we move from the side of the border with less generous benefits to the side with more generous benefits.

4.5.1.3 Comparison to Existing Estimates. Gillingham and Tsvetanov (2019) estimate the price elasticity of demand for solar panel installations using an approach that accounts for excess zeroes, unobserved heterogeneity, and the endogeneity of installation prices. Their estimates imply a price elasticity of demand evaluated at the mean installation price equal to -0.65. We simulate a marginal increase in installation prices and calculate the implied price elasticity evaluated at the mean installation price. This yields an estimate of -0.74, close to the elasticity estimated by Gillingham and Tsvetanov (2019).

Crago and Chernyakhovskiy (2017) analyze the effects of policy incentives on residential solar panel installations using county-level panel data from 12 states in the US Northeast. They find that increasing rebates by \$1 per watt increases solar panel installations by 47%. We replicate this experiment using our structural model and find that increasing rebates by \$1 per watt in the same 12 states

increases installations by 36%, of a similar magnitude to the estimates in Crago and Chernyakhovskiy (2017).⁴³

Hughes and Podolefsky (2015) estimate the effects of subsidies on solar panel installations by examining the introduction of a solar rebate in California. In their preferred estimate, they find that a \$470 increase in total rebate leads to a 10% increase in installations. From our estimates above, we can see that a \$470 increase in subsidies would lead to approximately a $.47 \times \frac{1}{\sigma} \approx 4$ percent increase in the number of installations. Thus, our result is smaller than the estimate in Hughes and Podolefsky (2015) but of a similar magnitude.

4.5.2 Power Plants. We now evaluate the performance of the power plant portion of the model. We include additional model fit results for power plants in Online Appendix C.3.6, including out-of-sample predictions using 2020 data. Figure 21 shows predicted electricity production of dispatchable plants against actual electricity production.⁴⁴ Each dot represents an hour of aggregate production by dispatchable plants for each interconnection in the data (X-axis) and predicted by the model (Y-axis). The model fits well in all three interconnections, producing R-squared values of 0.99, 0.96, and 0.96 in the East, West, and Texas, respectively.

We assess fit over hours and seasons in Figure 22. Each panel shows predicted and actual production of dispatchable plants for the average day for

⁴³Specifically, we increase the per-panel subsidy, s_j^{Panel} , in these same 12 states. We convert the per-watt subsidy into a per-panel subsidy by assuming 250 watts per panel.

⁴⁴While generally performing very well, we note that these predictions are lower than actual production values by about 2 percent on average. The Tobit model relies on normal and homoskedastic errors to produce consistent estimates (Cameron & Trivedi, 2005, p. 538). As these models are power plant specific, this assumption applies to the errors *within* power plants, some of which may violate the normality assumption. We do not think this issue meaningfully affects our results as the discrepancy between actual and predicted production is small.




Figure 21. Model fit at the interconnection level.

The Y-axis gives the production of dispatchable plants predicted by the model, measured in GWh while the X-axis gives the production in the data. Dots represent an hour of production for each interconnection, smoothed lines show the fit of a generalized additive model.

each of the three interconnections across four seasons. This figure shows that the model matches daily peaks and troughs of production in response to changes in demand and differences in intraday timing of those peaks and troughs between seasons. In particular, our model is able to generate the pattern of increasing dispatchable production through the afternoon, the time where solar power generation decreases and electricity demand increases. This is especially true for seasons and interconnections when solar makes up a larger share of electricity production. A region-level breakdown of these plots is available in Online Appendix C.3.6.

Not only does the model match total production, but it also replicates changes in the fuel mix at varying demand levels, reflecting that plants differ in how balancing authorities dispatch them as a function of excess demand. Figure 23 shows the percentage of total production in each interconnection produced by plants of each fuel type in the model and the data. The X-axis of each panel varies




Figure 22. Model fit at the interconnection level by hour and season.

Hour reflects Eastern Standard Time (EST). Each panel shows predicted and actual production of dispatchable plants over the course of the average day, for each of the three intersections and across four seasons. The dashed green line gives electricity production in the data while the solid orange line gives predicted production.




Figure 23. Fuel mix of production by interconnection.

The X-axis gives excess load at the interconnection level and the Y-axis gives the percent of electricity production that is produced by each of the fuel types. The dashed lines show the fuel mix in the data while the solid lines show the simulated fuel mix.

the interconnection-level excess load—the total amount of electricity demand that must be satisfied by dispatchable plants. Across all interconnections, our model’s predictions match the observed fuel mix very well. In all interconnections in the model and data, natural gas as a share of production increases in excess load.

Meanwhile, production levels of clean, low-marginal-cost nuclear and hydroelectric plants generally decrease as a percentage of total production. An important difference between the Eastern and Western Interconnections is that coal increases its share of production in the East, whereas coal’s production share declines except at the lowest levels of excess load in the West.

These changes in the fuel mix imply that the marginal damages of electricity production may vary not only spatially but also as a function of electricity demand. To illustrate this, Figure 24 plots simulated marginal damages of energy production

in each region within each interconnection as a function of excess load.⁴⁵ Overall, marginal damages are highest in regions within the Eastern interconnection, reflecting, in part, the interconnection’s reliance on production from coal-fired power plants. However, there is significant heterogeneity in the marginal damages across regions within this interconnection. Marginal damages are highest from power plants in the RFC region, which spans much of the Mid-Atlantic and lower Great Lakes. Regions also vary in the extent to which their marginal damages of production change in excess load. In the Western and Texas interconnections, marginal damages are relatively flat as a function of excess load while marginal damages in several regions in the Eastern interconnection are strongly increasing in excess load. For example, in the NPCC, the region covering the Northeast, marginal damages increase by over 30% between the 10th to 90th percentile of excess load, going from \$67/MWh to \$88/MWh.

4.6 Counterfactuals and Nationally-Optimal Subsidies

4.6.1 Welfare-Maximizing Subsidies. We now use the estimated structural model to quantify the welfare-maximizing solar subsidies characterized by equation (4.8). We outline the algorithm we use to numerically solve for welfare-maximizing subsidies in Online Appendix C.2.4. The results are displayed in Figure 25 and in Table 8. In all tables, “Baseline” refers to simulated outcomes under the current system of subsidies.

⁴⁵To calculate this, we simulate increasing excess load by a small amount in the region in question. We then divide the resulting change in total damages associated with power plants in the region by the change in total production by these power plants. Note that this is the marginal damage with respect to electricity production within a given region, not electricity demand from a given region. To the extent that a region imports electricity from other dirtier or cleaner regions, the marginal damage of electricity demanded may be higher or lower. For example, the NPCC imports electricity from the relatively dirty RFC, making marginal damage of electricity demand in NPCC higher than the electricity produced there.




Figure 24. Estimated marginal damage of electricity production by region.

The X-axis of each panel varies the total excess load in each of the three interconnections, and the Y-axis gives the simulated marginal damages per MWh of electricity produced in each region. In the Western interconnection, CAL covers most of California and WECC the rest of the western states. TRE is the only region in Texas. In the eastern interconnection, MRO covers the Midwest, NPCC the Northeast, RFC the mid-Atlantic, and SERC the Southeast. See Appendix C.1.7 for a detailed description of the regions and Figure C.5 for service areas of each region. See text for additional details on marginal damage calculations.

	(1)	(2)	(3)	(4)	(5)
	Baseline	State-Specific Subsidies		Tract-Specific Subsidies	
		Welfare Max	Damage Min	Welfare Max	Damage Min
I. Average Subsidy (\$Thousands)					
Midwest	10.0	15.4	15.9	15.3	15.8
Northeast	17.1	13.1	11.2	13.1	11.1
South	10.6	14.6	14.5	14.6	14.4
West	11.7	12.6	10.5	12.6	10.3
II. Installations per 1000HHs					
Midwest	5.4	8.1	8.9	8.1	8.9
Northeast	20.3	11.3	11.0	11.3	10.9
South	6.6	9.0	9.3	9.0	9.3
West	11.9	13.1	11.3	13.1	11.2
National	9.9	10.3	10.0	10.3	10.0
III. Annual Damages Offset (\$Millions)					
CO2e	193.4	216.9	219.0	217.1	219.4
NOx	12.0	13.3	13.3	13.3	13.3
PM2.5	33.2	36.8	34.7	36.8	34.6
SO2	23.6	25.5	28.5	25.7	28.9
Total	262.3	292.6	295.4	292.9	296.1

Table 8. Welfare maximizing results

Panel I shows the average present discounted value of subsidies received for a 15-panel installation for each census region. Panel II gives the simulated number of solar installations per 1000 households in the model for each Census region. Panel III gives the total damages offset by rooftop solar. All monetary values are measured in 2014 dollars. See text for details on each simulation.

Figure 25a and the first panel of Table 8 show how total subsidy generosity varies across states under the nationally-optimal system of subsidies.⁴⁶ We measure subsidy generosity as the present discounted value of subsidies an average household in each state would receive if they installed a 15-panel system. Several states in the Northeast and Northwest, areas with relatively environmentally friendly power plants and little sunlight, have the lowest optimal subsidies, at under 10 thousand dollars in present value. On the other end of the spectrum, eight states have optimal subsidies valued at over 17 thousand dollars. In West Virginia, one of these states, current subsidy levels are some of the least generous in the country at under 6 thousand dollars. More generally, nationally-optimal subsidies are highest in the Mid-Atlantic and lower Great Lakes and are lowest in the Northwest and Northeast.

Figure 25b and the second panel of Table 8 quantify the misallocation caused by the current system of subsidies on the spatial distribution of solar panel installations.⁴⁷ Current installations in the Midwest and South are roughly 30% lower than under nationally-optimal subsidies, while installations in the Northeast are 80% higher than the optimal level. These results suggest that the current system of subsidies leads to a substantial misallocation of solar panels across states.

Panel III of Table 8 summarizes the environmental cost of this misallocation. Switching from the current subsidy scheme to the nationally-optimal scheme would increase annual damages offset by rooftop solar from \$262.3 million to \$292.6

⁴⁶In Online Appendix C.3.7, we show how the government should optimally allocate subsidies across the three subsidy types: cost-based subsidies, per-panel subsidies, and production-based (per-kWh) subsidies.

⁴⁷We find that almost all of the adjustment comes via the extensive margin, rather than the intensive margin (number of panels per installation). In Online Appendix C.3.8, we show how the average installation size changes across counterfactuals.




Figure 25. State-level nationally-optimal subsidies and misallocation for welfare-maximizing reforms.

Panel (a) gives the nationally-optimal state subsidies. Subsidies are measured as the present discounted value associated with a 15-panel installation in 2014 dollars, averaged across all households in the state. Panel B shows state-level installations under the current system as a percentage of installations under the nationally-optimal system. These results are shown in table form in Online Appendix C.3.9.

million, equal to a 11.5% increase in the aggregate environmental benefits of solar panels.⁴⁸ A decrease in CO2 equivalent emissions drives most of the environmental gains, with relatively minor effects on damages from other pollutants.

4.6.2 Damage-Minimizing Reforms. An alternative social objective is to choose the system of subsidies that minimizes environmental damages. In this section, we consider a government who chooses subsidies to

⁴⁸These environmental benefits are considerably smaller than the environmental benefits of reallocating panels found in Sexton et al. (2021). There are two main reasons for this difference. First, we consider a government with much more limited policy instruments. Sexton et al. (2021) consider a planner who can directly allocate panels across states subject to local capacity constraints. Here we consider a government which can only influence installations through subsidies. Second, we utilize emissions data from 2019 rather than data from 2007-2016. Holland et al. (2020) find that power plant emissions decreased dramatically between 2010 and 2017. This decline was especially large in the Eastern interconnection, where emissions are historically higher than in the Western interconnection. This decrease in variation of damages across locations lowered the environmental benefits of reallocating panels across space.




Figure 26. State-level nationally-optimal subsidies and misallocation for damage-minimizing reforms.

Panel (a) gives the nationally-optimal state subsidies. Subsidies are measured as the present discounted value associated with a 15-panel installation in 2014 dollars, averaged across all households in the state. The color scale censors subsidy levels below \$5K and above \$20K. Panel B shows state-level installations under the current system as a percentage of installations under the nationally-optimal system. These results are shown in table form in Online Appendix C.3.9.

minimize the net present value of environmental damages, $D(\text{ELoad})$, subject to the government budget constraint. We formalize the government's problem and present the first-order conditions in Online Appendix C.2.5.

Figure 26 and the third columns of Tables C.8 and Table 8 show the results. Like the welfare-maximizing policies, the damage-minimizing policies are most generous in the Mid-Atlantic and lower Great Lakes, and are least generous in the Northwest and Northeast. However, the variation across states in subsidy generosity is greater than under the welfare-maximizing subsidies: optimal damage-minimizing subsidies range from under 4 thousand dollars to over 21 thousand dollars. The reallocation of solar panels induced by the damage-minimizing subsidies would lead to a 12.6% increase in aggregate environmental benefits.

4.6.3 Tract-level Subsidies. In the results above, we found that optimally set state-level subsidies lead to large environmental benefits relative to the current system of subsidies. Could subsidies set at a more granular geographic level lead to even larger gains? To answer this, we solve for the welfare-maximizing and damages-minimizing subsidies when subsidy levels are allowed to vary nonparametrically across census tracts.⁴⁹

Columns (4) and (5) of Table 8 display the results. In both the welfare-maximizing and damage-minimizing cases, the average subsidies across regions and the distribution of installations with optimal tract-level subsidies are similar to those under optimal state-level subsidies, and the damages offset with optimal tract-specific subsidies are only slightly larger than those with optimal state-specific subsidies. We conclude that the nationally-optimal system of state-level subsidies can capture most of the gains of more geographically granular subsidies.

4.6.4 Unconstrained Reforms. Our previous counterfactuals have focused on budget-neutral reforms. Here we analyze the case where the government does not face an externally set budget constraint and maximizes utility less environmental damages and government cost. In this case, the government's problem is to maximize

$$\underbrace{\int_i V_i di}_{\text{Utility}} - \underbrace{D(\text{ELoad})}_{\text{Damages}} - \lambda \underbrace{\sum_j \int_{i \in I_j} s_{ij} m_i^* di}_{\text{Government Cost}}, \quad (4.14)$$

⁴⁹ A_i is constant within census tracts in our quantitative model, and therefore solar production within census tract is simply proportional the number of panels installed. Thus, there is no unique nationally-optimal system of subsidies when the planner can use both panel-based and production-based subsidies. Therefore, we set per-panel subsidies to 0 in this exercise.

	(1)	(2)	(3)	(4)
	Baseline	Unconstrained Optimal		
		$\lambda = 1$	$\lambda = 1.25$	$\lambda = 1.5$
I. Average Subsidy (\$Thousands)				
Midwest	10.0	11.9	7.2	4.1
Northeast	17.1	9.9	5.6	2.7
South	10.6	11.2	6.7	3.6
West	11.7	9.5	5.3	2.5
II. Installations per 1000HHs				
Midwest	5.4	6.0	4.1	3.1
Northeast	20.3	8.6	5.9	4.6
South	6.6	6.7	4.6	3.5
West	11.9	10.1	7.0	5.5
National	9.9	7.8	5.4	4.2
III. Annual Damages Offset (\$Millions)				
Total	262.3	218.6	147.4	113.4
IV. Annuitized Total Subsidies Paid (\$Millions)				
National	380.8	218.5	86.2	34.7

Table 9. Unconstrained Nationally-Optimal Subsidies.

The first panel shows the average present discounted value of subsidies received for a 15-panel installation for each census region. The second panel gives the simulated number of solar installations per 1000 households. The third panel gives the annual environmental benefits generated by residential solar panels. The final panel gives the total government spending on subsidies converted into an annuity value. All monetary values are measured in 2014 dollars.

where λ represents the marginal cost of public funds. We present results over a range of λ 's from 1 to 1.5—for $\lambda = 1$, the government weights fiscal costs and household utility equally. We can think of this as a setting in which there are no distortionary costs of raising public funds, e.g. the government has access to lump-sum taxation.⁵⁰ We present the first-order conditions of the government's problem in Online Appendix C.3.10.

⁵⁰Jacobs (2018) shows that the marginal cost of funds is also equal to one in a Mirrleesian framework with the optimal tax system. A value of 1.3 is a commonly-used “middle of the road” estimate of the marginal cost of public funds in the literature (Ballard, Shoven, & Whalley, 1985; Finkelstein & McKnight, 2008; Newhouse, 1992; Olken, 2007; Poterba, 1996) while 1.5 is considered a “conservative” estimate (Finkelstein & Hendren, 2020; Heckman, Moon, Pinto, Savelyev, & Yavitz, 2010).

The results are summarized in Table 9. The first two columns show results under the current subsidy system and unconstrained nationally-optimal subsidies when we use a marginal cost of funds (λ) equal to 1. Optimal subsidies slightly exceed current subsidies in the Midwest and South, but fall below current levels in the Northeast and West. Table C.12 in Online Appendix C.3.10 shows the state-level optimal subsidies. These optimal subsidies result in fewer installations nationally, with Panel II showing that installations under the optimal subsidies are 80% of current levels nationally.

Panels III and IV show the annual environmental benefits of rooftop solar and the annuitized government spending on subsidies. Switching to unconstrained nationally-optimal subsidies decreases environmental benefits by 43 million dollars annually. However, the accompanying 162 million dollar decrease in fiscal costs dwarfs this decrease in environmental benefits. The optimal subsidy scheme achieves 83% of the environmental benefits at 57% the current cost.

To account for the additional costs of raising funds for solar subsidies from the use of distortionary taxes, we now calculate the optimal unconstrained subsidy with alternative values of λ , the marginal costs of public funds. In Columns 3 and 4 of Table 9, we recalculate the optimal subsidies with $\lambda = 1.25$ and with $\lambda = 1.5$, respectively. With these larger values of the marginal cost of public funds, the optimal amount of spending on solar subsidies decreases, as raising money for subsidies entails an efficiency cost. With $\lambda = 1.25$, optimal spending on subsidies is only 23% of current levels, and with $\lambda = 1.5$, optimal spending is less than one-tenth of current levels.

As discussed in the introduction, residential solar subsidies may be associated with additional external benefits that we do not capture here. However,




Figure 27. Damages offset per additional dollar of government funds

Damages associated with marginal increases in production-based subsidies, s_j^{kWh} , around the current system of subsidies. All monetary values are measured in 2014 dollars.

our results suggest that these additional benefits would have to be quite large to justify the current spending levels on these subsidies nationally.

4.6.5 Marginal Subsidy Increases. Relative to the current system of subsidies, what marginal subsidy increases are the most cost-effective way to decrease damages? To answer this, we calculate the damages offset per additional dollar of government cost associated with marginal subsidy increases around the current system of subsidies. Specifically, we first simulate the model 1) under the current system of subsidies and 2) with marginally more generous subsidies of a given type in a given state. We calculate the damages offset per dollar of this particular subsidy as the difference in damages between the two simulations divided by the difference in the fiscal cost. We repeat this process for each subsidy type in each state.

Figure 27 shows the marginal damages offset per dollar for production-based subsidies.⁵¹ There are large differences in damages offset across states. For example, a small subsidy increase in Washington only leads to 22 cents less

⁵¹Within each state, there are only small differences in the damages offset per dollar across the three subsidy types.

environmental damages per dollar of government funds. On the other hand, subsidy increases in West Virginia are highly cost-effective—for an additional dollar of government spending, environmental damages decrease by 80 cents. Subsidy increases in Ohio and Pennsylvania are also associated with damages offset per dollar of over 70 cents. These results imply that small, cost-neutral shifts in subsidies across states could lead to substantial decreases in environmental damages. For example, a cost-neutral shift from subsidies in Washington to subsidies in West Virginia would lead to decreases in environmental damages of $80 - 22 = 58$ cents for each dollar reallocated. Put another way: if the goal of Washington policymakers were to reduce total environmental damages, they would be significantly better off subsidizing solar installations in West Virginia, rather than in Washington.

4.7 Extensions, Robustness and Further Issues

4.7.1 Alternative Specifications of Household Utility. In our baseline specification, we specified $\gamma_i(\cdot)$, the function which dictates a household’s nonpecuniary benefits of solar installation, as a function of the number of panels installed and the local average education level and political leaning. In Online Appendix C.4.1, we assess the sensitivity of our main results to this specification of the utility function by changing the specification of $\gamma_i(\cdot)$. In each specification, we re-estimate the model given the alternative specification of utility and then solve for the optimal cost-neutral policy given the new estimates of the household utility function. Across all specifications, we find similar optimal subsidies, similar changes in installation rates, and similar environmental benefits.

4.7.2 Alternative Discounting of Future Subsidy Payments.

In our baseline model, we assumed that households value production subsidies

at the present discount value of subsidy payments, calculated using the market interest rate. However, De Groote and Verboven (2019) find that households use a much higher implicit interest rate when evaluating the future subsidy payments associated with solar panel installations.

In Online Appendix C.4.2, we reestimate our model and recalculate the optimal subsidies when households calculate the net present value of future subsidy payments using an implicit discount rate of 15%, based on the estimates from De Groote and Verboven (2019). We have two main takeaways from this extension. First, the change in household discounting of future subsidy payments has important implications for the *type* of subsidy the planner chooses. When households discount future subsidies more heavily than upfront subsidies, the social planner relies less on production subsidies, as they are valued less than upfront subsidies.⁵² Therefore, in this extension, the social planner relies mostly on investment subsidies, whereas the planner in our baseline results uses production subsidies most heavily. Our results echo those from De Groote and Verboven (2019), who find that upfront investment subsidies are more cost-effective than production subsidies at inducing installations. Second, the optimal *geographic* distribution of subsidies and the optimal overall level of subsidy spending are similar to our baseline results. We conclude that this alternative assumption about household discounting does not change our main results about spatial misallocation of panels nor about optimal levels of spending across space.

4.7.3 Line Losses. Our baseline model does not account for line loss: the electricity that is lost as electricity is transmitted over the grid from a power

⁵²We do not allow the government to set negative production subsidies. Without this constraint, the government will heavily tax solar production.

plant to a consumer. Rooftop solar reduces line loss by reducing the amount of electricity that needs to be transported across the grid.

In Online Appendix C.4.3, we re-calculate our main results in a model which accounts for line loss, where we base our model of line loss on the model and estimates from Borenstein and Bushnell (2022). In the extension, line losses are determined endogenously as a function of the amount of electricity in each region that must be transmitted between central generation plants and households. Therefore, residential solar offsets damages not only by directly producing power that would otherwise be produced by fossil-fuel plants, but also by reducing transmission across the grid and the resulting line losses.

The takeaways are qualitatively the same as our main results. As expected, the environmental benefits of solar panels increase. As a result, the total government spending under the optimal unconstrained subsidies is higher than in the baseline case but still less than under current policies. The welfare-maximizing and damage-minimizing cost-neutral reforms generate increases in aggregate environmental benefits of 11.4% and 12.3%, respectively.

4.7.4 Transmission Constraints. In times when transmission lines are under heavy use, grid operators will limit the extent to which electricity is moved over long distances to prevent transmission lines from overheating. Our baseline model does not account for these transmission constraints because we assume that a given power plant's production level can always depend on excess load in other regions within the plant's interconnection. If transmission constraints are binding, however, the plant's power production is unlikely to depend on excess load in other regions because grid operators will avoid transmitting electricity across regions.

We now consider a model extension that includes a stylized model of potentially binding transmission constraints. We assume that if the excess load in a given region R exceeds a certain threshold $\overline{\text{ELoad}}_R$, then the policy functions of power plants will depend only on the excess load in their own region and not on excess load in other regions. Thus, we can express the latent function for power plant k in region R as:

$$f_k(\text{ELoad}_t, \varepsilon_{kt}) = \begin{cases} \psi_k^0 + \sum_{R' \in \mathbf{R}_k} (\psi_{R'k}^1 \text{ELoad}_{R't} + \psi_{R'k}^2 \text{ELoad}_{R't}^2) + \varepsilon_{kt} & \text{if } \text{ELoad}_{Rt} \leq \overline{\text{ELoad}}_R \\ \psi_k^{0,\text{Cons}} + \psi_k^{1,\text{Cons}} \text{ELoad}_{Rt} + \psi_k^{2,\text{Cons}} \text{ELoad}_{Rt}^2 + \varepsilon_{kt} & \text{if } \text{ELoad}_{Rt} > \overline{\text{ELoad}}_R \end{cases}. \quad (4.15)$$

In practice, we set $\overline{\text{ELoad}}_R$ as the 75th percentile of excess load in a given region in our data. We re-estimate the power plant policy functions via maximum likelihood and recalculate the optimal policies with the updated policy functions. We present the results in Appendix C.4.4. As in our baseline specification, we find that optimal unconstrained subsidies are lower than current levels in the Northeast and the West, but are similar to or slightly above current levels in the Midwest and South. We find that current spending on subsidies exceeds the optimal level by about 45% in this specification.

4.7.5 Improved Storage of Nondispatchable Technology.

A significant issue facing the expansion of renewable electricity generation is that solar and wind are nondispatchable. Thus, these sources can only produce electricity when environmental conditions are suitable—when the sun is shining, or the wind is blowing. One of the leading solutions to this problem is an expanded

capacity of electricity storage in the form of batteries. In Online Appendix C.4.5, we consider a stylized way to incorporate storage technology into our model. We allow nondispatchable electricity to be stored and used proportionately to the total load. Effectively this means we reallocate solar and wind production from their exogenous time profile of production to match the time profile of demand, which loosely matches the optimal behavior of storage owners arbitraging electricity across time to maximize profits.

We find that the storage technology itself generates considerable environmental benefits. However, storage technology does not substantially change the environmental benefits of solar panels across locations. Therefore, we find that adding storage technology does not qualitatively change the optimal cost-neutral or unconstrained reforms, the distribution of installations under the optimal subsidies, or the environmental benefits of switching to optimal subsidies. See Butters, Dorsey, and Gowrisankaran (2021) or Holland, Mansur, and Yates (2022) for a detailed treatment of storage technology.

4.7.6 Cleaner Electricity Production. Electricity production in the United States has become considerably cleaner over the past few decades. Our baseline results quantify the value of optimizing solar panel subsidies given current electricity production technology. Here, we are interested in determining what would happen to our main results if the grid were considerably cleaner than it is presently.

Increased production of utility-scale renewables and fuel switching (from dirty to clean coal and from coal to natural gas) are the two primary drivers of the reduction in emissions from electricity generation. We perform four additional simulations to assess how further clean-up of electricity production would affect

our results. First, we find the optimal subsidies under expanded production from utility-scale solar and wind based on three scenarios of projected renewable expansion by 2030 from the EIA (Nalley & LaRose, 2022). Second, we recalculate results considering each coal plant to have “cleaned up” production. Our method of cleaning up coal plants is to adjust marginal damages from coal plants so that the mean and standard deviation of marginal damages from coal plants match that of natural gas plants.

We present the results in Online Appendix C.4.6. Intuitively, we find that if electricity production becomes cleaner, the environmental damages offset by solar panels will decrease. Therefore, the optimal unconstrained subsidies for residential solar with cleaner electricity production are lower than in our baseline case, suggesting optimal unconstrained subsidies will be even lower in the future if electricity production continues to become cleaner. However, there remains considerable heterogeneity in the marginal damage of electricity production across space. Thus, we still find substantial benefits in switching from the current subsidies to optimal cost-neutral subsidies.

4.7.7 Distributional Effects. The proposed switch in the system of subsidies could have distributional impacts through two channels—directly through a change in subsidies received by households and indirectly through the induced change in pollutant damages caused by electricity generation. Households who install solar panels, and therefore receive subsidies, tend to be wealthy (Borenstein & Davis, 2016). Our proposed nationally-optimal subsidies will likely be progressive relative to the current subsidies since switching from current to optimal generally involves decreasing subsidies in high-income states such as Massachusetts and increasing subsidies in low-income states such as West Virginia. For this same

	Dependent Variable:			
	Log Installations per Capita			
	(1)	(2)	(3)	(4)
Historically-Adjusted Subsidies	0.059*** (0.0059)	0.060*** (0.0047)	0.060*** (0.0055)	0.071*** (0.0088)
Income Tax	NO	YES	YES	YES
Sales Tax	NO	NO	YES	YES
Property Tax	NO	NO	NO	YES
Distance Bandwidth	10	10	10	10
Observations	6,052	6,052	6,052	6,052

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

Table 10. Regression of log installations per capita on “historically-adjusted” measures of subsidy generosity.

Measures from Sexton et al. (2021), subsidies are in thousands of dollars. Sample is limited to tracts within 10 miles of state borders. Historically-adjusted subsidy generosity measures the net present value of subsidies a household would expect to receive given subsidies in place in the year of installation, measured in thousands of 2014 dollars.

reason, switching to optimal subsidies will likely improve the distribution of damages caused by electricity generation. Similarly, Dauwalter and Harris (2023) find that shifting solar capacity to locations where the environmental benefits are greatest would lead to environmental benefits for disadvantaged groups.

We have refrained from accounting for distributional effects when calculating optimal subsidies, as this paper is primarily concerned with the spatial misallocation of solar due to differences in the generosity of solar subsidies across states. Seriously tackling the distributional effects of solar subsidies requires a different set of policy instruments than those analyzed here, such as means-tested subsidies for solar installations.⁵³

⁵³Colas and Reynier (2024) solve for the optimal income-contingent subsidies for residential solar panels.

4.7.8 Installation Elasticities with Historical Subsidy Measures.

As described earlier, our model of solar installation captures the long-run installation behavior of households facing a given menu of subsidies. One potential issue is that we estimate the model using solar subsidies from 2017 when, in fact, many of the installations observed in our data were made before 2017 by households facing pre-2017 subsidy levels.

To get a sense of how accounting for pre-2017 subsidies would affect our estimates of the household installation model, we perform a simple exercise using a measure of “historically-adjusted” subsidy generosity from Sexton et al. (2021). To calculate this measure, the authors first calculate the net present value of subsidies a household in a given state would expect to receive, given the subsidies in place from each year from 2000 to 2017. They then take the average over these yearly measures, weighing them by the number of installations in each year.⁵⁴ This is a measure of the net present value of subsidies from a solar installation, adjusting for the fact that households have faced different subsidies based at the time when they installed.

We then replicate the border discontinuity regressions from Table 6 using this historically adjusted measure of the net present value of subsidies. Specifically, we regress tract-level log installations per capita on this historically adjusted measure of the net present value of subsidies, limiting our sample to tracts within 10 miles of state borders. Table 10 shows the results. Across all specifications, we estimate that a 1000 dollar increase in the historically-adjusted net present value

⁵⁴To generate the subsidy measures reported in Sexton et al. (2021), the authors then divide the net present value of subsidies by the total kWh of electricity an installation is expected to produce, such that the reported measure represents the subsidy received per kWh. We do not perform this conversion so that our measure represents the total net present value of subsidies for a given installation, measured in thousands of dollars. This makes our estimates here easier to compare to our regressions in Section 4.4.2.

of subsidies is associated with a 5.9% to 7.1% increase in installations per capita, similar to the results we found in Table 6, when we used our baseline subsidy measures. These results suggest that not accounting for the profile of subsidies households have faced over time may not introduce significant bias when estimating the long-run responsiveness of installation to subsidies.

4.8 Conclusion

We have used a structural model of solar panel demand and electricity production to calculate the nationally-optimal system of subsidies for residential solar panels and to quantify the benefits of switching to such a system. Our main conclusions are that the current system of subsidies leads to a spatial misallocation of panels, and nationally spending on subsidies is too high. However, our results do not necessarily imply that the US should lower funding for renewable energy programs in general, rather that government funds spent on subsidies for residential solar subsidies would be better spent on other programs. These alternative programs could include other investments in renewable energy, such as subsidies for utility-scale solar or wind power, both of which provide energy at lower cost than residential solar (Lazard, 2023).

Future work can extend our model to incorporate endogenous entry and exit of electricity generators, as in Holland, Mansur, and Yates (2022). In that case, residential solar subsidies could disincentivize entry of new generators, which could be costly from an environmental perspective if the new generators employ cleaner technology than incumbents. It would also be interesting to utilize similar frameworks to analyze other consumer subsidies for energy-related products, such as subsidies for home insulation, small wind systems, and geothermal heat pumps. We leave these questions for future work.

APPENDIX A

APPENDIX FOR THE DISTRIBUTIONAL IMPACTS OF CLIMATE CHANGE ACROSS US LOCAL LABOR MARKETS

A.1 Data Appendix

A.1.1 Sample Selection. We use 5% samples of the 1990 and 2000 censuses in addition to the 2010 and 2019 5-year aggregated American Community Survey. In all years, of data, we drop individuals living in group quarters, veterans, and any households younger than 16 and older than 64. We consider the household head the decision maker—unless the head is unemployed and the spouse is employed, in which case the spouse is the decision maker. We consider a household college educated if they have at least four years of college.

A.1.2 City-Level Indices. In this section, we detail how we construct-city level estimates of energy consumption, wages, and rents. In each section, we use individual data from the Census and American Community Survey.

A.1.2.1 Energy Use. Our model features electricity and natural gas use in the household utility function. The main concern is that some renters may not pay utility bills separately from rent and thus falsely report zero energy expenditure in the census and ACS data. We follow Glaeser and Kahn (2010) and use the EIA’s Residential Energy Consumption Survey (RECS) to correct this feature of the census data. RECS directly reports energy consumption, but we cannot use it as our primary data source since it does not report the location of each household and has a relatively small sample. We proceed in two steps, first estimating city-demographic group-year level energy consumption for single-family homeowners in the census data, then adjusting those estimates by the difference between single-family homeowners and renters in the RECS data.

Specifically, we estimate the following regression for single-family homeowners in the census,

$$e_{idjt}^m = \gamma_{djt}^m + X'_{idjt}\beta^m + \varepsilon_{idjt}, \quad (\text{A.1})$$

where e_{idjt}^m is use of energy type $m \in \{elec, gas\}$ for household i , in demographic group d , city j , and year t . Our vector of controls, X'_{idjt} , includes the number of people in the household, the number of children, and the age of the household head. We estimate this for single-family homeowners since they are most likely to report accurate energy expenditures in the census. We then take the national average of the controls for each year-demographic group, $\bar{X}'_{dt}$, and use the coefficients estimated in equation (A.1) to predict energy use for single-family homeowners for each city, demographic group, and year,

$$E_{djt}^{SFO} = \hat{\gamma}_{djt}^{elec} + \bar{X}'_{dt}\hat{\beta}^{elec} \quad (\text{A.2})$$

$$G_{djt}^{SFO} = \hat{\gamma}_{djt}^{gas} + \bar{X}'_{dt}\hat{\beta}^{gas}. \quad (\text{A.3})$$

Then, we turn to the RECS, which is administered every three to six years. We match our census years to the closest RECS survey available, using the 1993, 2001, 2009, and 2015 surveys.¹ We estimate the following regression,

$$e_{it}^m = \alpha_{mt}^1 MultiFam_{it} + \alpha_{mt}^2 Rent_{it} + \alpha_{mt}^3 MultiFam_{it} \times Rent_{it} + X'_{it}\beta_{mt} + \varepsilon_{it}, \quad (\text{A.4})$$

¹The 1997 and 2001 RECS surveys did not ask about the education status of surveyed households. We use division, race, home type, rental status, and household size from 1993 and 2009 to predict the probability of each observation in the 2001 sample being college educated using logistic regression. After predicting the probability of being college educated, we draw actual outcomes according to the predicted probability.

where e_{it}^m is use of energy type $m \in \{elec, gas\}$, $MultiFam_{it}$ is an indicator for whether household i lives in a multifamily home in year t , $Rent_{it}$ is an indicator for whether household i rents their home, X'_{it} is a vector of controls. The controls include indicators for the census division, number of children, household size, age of the household head, and an indicator for whether the household head is white. To estimate energy usage for each demographic group, we use the estimated coefficients $\hat{\alpha}_{mt}^1$, $\hat{\alpha}_{mt}^2$, and $\hat{\alpha}_{mt}^3$ to adjust the single-family owners estimates, E_{djt}^{SFO} and G_{djt}^{SFO} , weighting by the proportion of single-family and multifamily renters as follows,

$$E_{djt} = \pi_{djt}^{SFO} E_{djt}^{SFO} + \pi_{djt}^{MFO} E_{djt}^{SFO} e^{\hat{\alpha}_{elec,t}^1} + \pi_{djt}^{SFR} E_{djt}^{SFO} e^{\hat{\alpha}_{elec,t}^2} + \pi_{djt}^{MFR} E_{djt}^{SFO} e^{\sum_k \hat{\alpha}_{elec,t}^k} \quad (A.5)$$

$$G_{djt} = \pi_{djt}^{SFO} G_{djt}^{SFO} + \pi_{djt}^{MFO} G_{djt}^{SFO} e^{\hat{\alpha}_{gas,t}^1} + \pi_{djt}^{SFR} G_{djt}^{SFO} e^{\hat{\alpha}_{gas,t}^2} + \pi_{djt}^{MFR} G_{djt}^{SFO} e^{\sum_k \hat{\alpha}_{gas,t}^k}, \quad (A.6)$$

where π_{djt}^{SFO} , π_{djt}^{MFO} , π_{djt}^{SFR} , and π_{djt}^{MFR} are the proportion of single/multifamily owners/renters in demographic group d , city j , and year t .

A.1.2.2 Wages. We can write total income for a household of a demographic group d , in the city j and year t , as $I_{djt} = W_{ejt} \times l_{dt}$, where l_{dt} is the efficiency units of labor supplied by a worker in the demographic group d in year t . Efficiency units account for the probability that a worker is employed, the hours worked, and the productivity of that work conditional on being employed. Specifically, we have $l_{dt} = E_{dt} \times \tilde{l}_{dt}$, where E_{dt} is the probability that a household in the demographic group d is employed in year t , and $\tilde{l}_{dt}$ is productivity conditional on working. We consider a household employed if the household head or their spouse reports they are employed, work at least 48 weeks per year, and at least

35 hours per week.² We parameterize the productivity term as $\log \tilde{l}_{dt} = X'_{dt}\beta_{et}$ for a vector of controls X'_d , where $e \in \{\text{No College}, \text{College}\}$ indexes education level.

We first calculate E_{dt} as the proportion of employed households in the demographic group d and year t . Next, we estimate the following regression for employed households,

$$\log I_{djt} = X'_{idjt}\beta_{et} + \nu_{ejt} + \varepsilon_{idjt}, \quad (\text{A.7})$$

where ν_{ejt} is a education-city-year fixed effect that estimates $\log(W_{ejt})$. Our vector of controls, X'_{idjt} includes whether the household is married, whether the household has over 25 years of potential experience, and whether the household head or spouse is white. We assume that no unobservable factors affect income after adding the vector of controls, X'_{idjt} , and the education-city-year fixed effect. Baum-Snow and Pavan (2012) and Roca and Puga (2017) provide evidence that after conditioning on education, there is little selection on unobservable factors that affect income. We then calculate demographic-city-year income, I_{djt} , using estimates of $\hat{\nu}_{ejt}$ and $\hat{\beta}_e$, as $I_{djt} = E_{dt} \times \exp(\hat{\nu}_{ejt} \times \bar{X}'_{dt}\hat{\beta}_{et})$. Here, $\bar{X}_{dt}$ is the demographic group-year average value for each control variable. Using demographic group-year averages holds these observable characteristics constant across all cities.

A.1.2.3 Housing Quantity and Rent. Our model requires a quantity of housing units, H and the price for those housing units, P_j^H , but our data from the Census reports total housing expenditure. Specifically, the variable *RENT* gives monthly contract rent payments in nominal dollars, which we convert to annual, real dollars. Note that an individuals' housing expenditures can be

²We assume that the probability of employment within a demographic group does not vary across cities.

written as $\mathbb{E}_{ijt}^H = P_{jt}^H H_{ijt}$ where H_{ijt} is the quantity of housing the individual consumes. We parameterize housing quantity as $H_{ijt} = X'_{ijt}\beta_t$, where X'_{ijt} is a vector of observable characteristics, including the number of rooms, the number of units in the building, the number of bedrooms, the number of people living in the house per room, and the decade the house was built. Plugging in our parameterization and taking logs yields our estimating equation:

$$\log(\mathbb{E}_{ijt}^H) = \mu_{jt} + X'_{ijt}\beta_t + \varepsilon_{ijt}. \quad (\text{A.8})$$

where μ_{jt} is a city-year fixed effect that estimates P_{jt}^H . The coefficients on all of the observable housing characteristics are allowed to vary by year. We then define the per-unit cost of housing using the city-year fixed effect, $P_{jt}^H = e^{\mu_{jt}}$, and quantity of housing as housing expenditure divided by the per-unit cost of housing, $H_{ij} = \mathbb{E}_{ijt}^H / P_{jt}^H$. We limit the sample to renters aged 16 to 64 in this estimation. We refer to the per-unit cost of housing as rent throughout the paper.

A.2 Additional Descriptive Results

Figure A.1 shows mean fire and flood risk scores by census tract from the First Street Foundation data. The arid West generally has the highest fire risk scores, while the gulf coast has the highest flood risk scores.

Figure A.2 shows how the number of no-rain days and annual precipitation have changed between 1990 and 2019. Generally, the eastern half of the US has gotten wetter.

Figures A.3 and A.4 show the relationship between 1990 share of each city's population that is low-income or non-white and changes in precipitation variables between 1990 and 2019. Similar to temperature, there is a positive relationship between the share of a city that is low-income or non-white and increases in total




Figure A.1. Mean fire and flood risk scores from the First Street Foundation data by census tract.

The property-level scores range between 1 and 10, with 10 representing the most at-risk property. White tracts are those not included in the public FSF data.




Figure A.2. Change in the no rain days and annual precipitation.

No rain days are days with less than 1mm of precipitation. Difference taken between 5 year moving averages and are censored at the 5th and 95th percentiles of grid cells.




Figure A.3. Relationship between share non-white and climate.

Each point is a city, where the size of the point represents the city's 1990 population. Low income defined as households in the bottom quintile of the national income distribution in 1990. Regression lines are weighted by population.

precipitation. There is no clear relationship between the number of rain free days and either of the demographic variables.

A.2.1 Energy demand and climate. A robust energy justice literature explores differences in energy expenditures and demand between demographic groups. For example, Lyubich (2020) demonstrates that black households spend more on energy than observably similar white households. We begin by replicating this result using our census and ACS data. We regress total expenditure on electricity and natural gas on indicators for race interacted with income decile, adding fixed effects for city, year, education, and several other characteristics that may affect energy expenditures.³ Figure A.5 shows the

³Specifically, these are characteristics of the physical house: number of rooms, kitchen present, plumbing present, decade built, units in the structure, and number of bedrooms, as well as characteristics about the household: the number of children and marital status.




Figure A.4. Relationship between share non-white and climate.

Each point is a city, where the size of the point represents the city's 1990 population. Low income defined as households in the bottom quintile of the national income distribution in 1990. Regression lines are weighted by population.

gap in energy expenditures between black and white households throughout the income distribution. Black households spend more on energy than observably similar households throughout the income distribution—aside from the highest income decile, where black households spend slightly less on electricity. These results suggest that either (1) black households have stronger preferences for the services provided by this energy, or (2) black households live in structures with unobservably lower energy efficiency. Research demonstrates that black households have faced, and continue to face, more significant barriers in the housing market than white households (Christensen, Sarmiento-Barbieri, & Timmins, 2021; Christensen & Timmins, 2021).

The gap in expenditures between white and black households does not tell us whether they respond differently to changes in weather. Thus, we run similar




Figure A.5. Black-white energy expenditure gap by income decile.

The coefficients plotted are those from a regression of energy expenditures on indicator variables for income decile and income decile interacted with whether the household is black, with fixed effects for city, year, education level, number of children, marital status, number of rooms, kitchen present, plumbing present, decade built, units in the structure, and number of bedrooms. We randomly downsampled to 1 million observations to ease computation. Standard errors are clustered by city.

regressions but now include CDD and HDD as regressors and look for heterogeneity in the effects of CDD and HDD on energy expenditure by race and income. We add the same battery of fixed effects as before. Table A.1 shows heterogeneity in the effect of degree days by race. Black households increase spending by more than white households for an increase in HDD and by less than white households for an increase in CDD. Differences in the predominant heating fuel may explain this—black households increase electricity expenditure by more than white in response to more HDDs and decrease their natural gas use by less than white households in response to CDDs.

Next, we explore heterogeneity in the effect of degree days across the income distribution. Figure A.6 shows how the effect of degree days on energy expenditures differs from the fifth income decile for the rest of the income distribution. The

Table A.1. Effect of CDD and HDD on Energy Expenditures.

Dep. Var. Model:	Total (1)	Electricity (2)	Natural Gas (3)
<i>Variables</i>			
Black	155.0* (87.76)	-27.22 (68.10)	143.3*** (49.16)
HDD $\times$ Black	0.0580** (0.0251)	0.0493** (0.0191)	0.0235 (0.0154)
HDD	0.0434 (0.0801)	0.0989* (0.0503)	-0.0569 (0.0500)
CDD $\times$ Black	-0.0740 (0.0519)	-0.0126 (0.0400)	-0.0440* (0.0255)
CDD	0.3588** (0.1370)	0.1204 (0.1295)	0.2531** (0.1263)
<i>Fit statistics</i>			
Observations	820,930	872,960	829,054
R ²	0.30894	0.31055	0.18668
Within R ²	0.00411	0.00117	0.00528

Clustered (City) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Regressions include city, year, education, number of children, marital status, kitchen, number of rooms, plumbing, decade built, units in structure, number of bedroom fixed effects. Data from downsampled census and ACS data from 1990, 2000, 2010, and 2019 matched to PRISM data from that year.




Figure A.6. Heterogeneity in the effect of degree days on energy expenditures by income decile.

The coefficients plotted are those from a regression of energy expenditures on CDD and HDD interacted with indicator variables for income decile, with fixed effects for city, year, education level, number of children, marital status, number of rooms, kitchen present, plumbing present, decade built, units in the structure, and number of bedrooms. Thus, the estimates are difference in effects of degree days for an income decile relative to the fifth income decile. We randomly downsampled to 1 million observations to ease computation. Standard errors are clustered by city.

wealthiest households have the strongest increase in electricity expenditure in response to more CDDs and in natural gas expenditure in response to HDDs. Generally, the electricity response to HDDs is decreasing in income—suggesting that low-income households are more likely to use electricity as their heating fuel source. Since the highest income decile is least likely to be income-constrained with respect to their electricity spending, these results may suggest that their stronger response in CDD is the result of always maintaining their preferred indoor temperature. Other households may have a lower expenditure response if they tolerate a higher indoor temperature when it is hotter outside instead of spending more money on electricity to run air conditioning more intensely.




Figure A.7. Heterogeneity in the effect of degree days on energy expenditures by income decile and race.

The coefficients plotted are those from a regression of energy expenditures on CDD and HDD interacted with indicator variables for income decile and race, with fixed effects for city, year, education level, number of children, marital status, number of rooms, kitchen present, plumbing present, decade built, units in the structure, and number of bedrooms. Thus, the estimates are difference in effects of degree days for an income decile relative to white households in the fifth income decile. We randomly downsampled to 1 million observations to ease computation. Standard errors are clustered by city.

Figure A.7 shows heterogeneity in degree days’ effect by race and income decile. All estimates are relative to white households in the fifth income decile. Black and white households respond similarly in electricity expenditure to CDDs, except for the highest income decile—where white households have a much stronger response. Black households tend to have stronger responses to HDDs, though these differences are generally not statistically significant.

We can also use the Residential Energy Use Survey (RECS) to evaluate how households respond to climate. Though the surveys here do not report the respondents’ exact locations, the EIA does report HDD and CDD for those

locations.⁴ We are interested in heterogeneity in household energy responses to changes in climate. We regress electricity and natural gas usage on HDD and CDD, as well as interactions between climate variables and race. We add year, census division, and race fixed effects. The results, shown in Table A.2, are descriptive rather than causal—demonstrating differences in energy demand responses to climate by race. Households increase their electricity usage in response to higher HDD and CDDs. Black households are less responsive than white households on average, though the coefficients are not statistically significant. The results for natural gas usage are surprising—the coefficient on HDD is not significant, and the coefficient on CDD is negative and significant. This unintuitive result may be because households in warmer climates (high CDDs and low HDDs) are less likely to have natural gas heaters. We find that black households have a stronger demand response for natural gas to HDDs.

A.3 Model

A.3.1 Derivation of Comfort and Energy Prices. We can solve the households utility maximization problem conditional on choosing city j in three steps. First, households will consume gas and electricity in a cost-minimizing manner to produce a given amount of energy. The first-order conditions from minimizing the cost of energy using equation (2.2) give us the following conditional demand functions for electricity and gas.

$$E^*(\mathcal{E}|d, j) = M_{dj}^{\mathcal{E}} (P_j^G \theta_{dj}^E)^{\sigma_{\mathcal{E}}} \mathcal{E}, \quad (\text{A.9})$$

$$G^*(\mathcal{E}|d, j) = M_{dj}^{\mathcal{E}} (P_j^E \theta_{dj}^G)^{\sigma_{\mathcal{E}}} \mathcal{E} \quad (\text{A.10})$$

⁴The EIA adds noise to this measure to prevent identifying the respondent's exact location. This noise will slightly attenuate the results presented here.

Table A.2. Effect of CDD and HDD on Energy Demand.

Dep. Var. Model:	Electricity Usage (kWh) (1)	Natural Gas Usage (cu ft) (2)
<i>Variables</i>		
HDD	0.6920*** (0.0644)	0.5345 (0.7423)
CDD	1.274** (0.2925)	-5.781** (1.760)
HDD $\times$ Race = Black	-0.1726 (0.1237)	7.200** (2.134)
HDD $\times$ Race = Rest	0.1121 (0.0809)	0.9732 (0.8979)
CDD $\times$ Race = Black	-0.1448 (0.3853)	3.088 (1.452)
CDD $\times$ Race = Rest	-0.0347 (0.0833)	-0.2569 (0.9812)
<i>Fit statistics</i>		
Observations	29,702	29,702
R ²	0.15988	0.18400
Within R ²	0.01079	0.02157

Clustered (year) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Regressions include year, census division, and race fixed effects. Data from the 1993, 2001, 2009, and 2015 RECS surveys.

where $M_{dj}^{\mathcal{E}} = \left(\theta_{dj}^{E^{\sigma_{\mathcal{E}}}} P_{dj}^{G^{\sigma_{\mathcal{E}}-1}} + \theta_{dj}^{G^{\sigma_{\mathcal{E}}}} P_{dj}^{E^{\sigma_{\mathcal{E}}-1}} \right)^{\frac{\sigma_{\mathcal{E}}}{1-\sigma_{\mathcal{E}}}}$. These conditional demand functions allow us to derive a unit cost function for energy that does not depend on the quantity of energy. We will use this unit cost function as the price of energy,

$$P_{dj}^{\mathcal{E}} = \left(\theta_{dj}^{E^{\sigma_{\mathcal{E}}}} P_j^{E^{1-\sigma_{\mathcal{E}}}} + \theta_{dj}^{G^{\sigma_{\mathcal{E}}}} P_j^{G^{1-\sigma_{\mathcal{E}}}} \right)^{\frac{1}{1-\sigma_{\mathcal{E}}}} \quad (\text{A.11})$$

Similarly, Households will produce comfort using energy and housing in a cost-minimizing manner. The first-order conditions from minimizing the cost of producing comfort using (2.1) give us the following conditional demand functions for energy and housing.

$$H^*(\mathcal{C}|d, j) = M_{dj}^c P_{dj}^{\mathcal{E}^{\sigma_c}} \mathcal{C}_{dj}, \quad (\text{A.12})$$

$$\mathcal{E}^*(\mathcal{C}|d, j) = M_{dj}^c P_j^{H^{\sigma_c}} \mathcal{C}_{dj} \quad (\text{A.13})$$

where $M_{dj}^c = \left(P_{dj}^{H^{\sigma_c-1}} + P_j^{\mathcal{E}^{\sigma_c-1}} \right)^{\frac{\sigma_c}{1-\sigma_c}}$. We can use the unit cost function derived from these conditional demand functions to calculate the price of comfort

$$P_{dj}^{\mathcal{C}} = \left(P_{dj}^{H^{1-\sigma_c}} + P_{dj}^{\mathcal{E}^{1-\sigma_c}} \right)^{\frac{1}{1-\sigma_c}} \quad (\text{A.14})$$

We also use the derived demand for comfort to yield the unconditional housing and energy demand functions:

$$H_{dj}^* = M_{dj}^c P_{dj}^{\mathcal{E}^{\sigma_c}} \times \frac{\alpha_{dj}^{\mathcal{C}}(I_{dj} + \Upsilon_d)}{P_{dj}^{\mathcal{C}}} \quad (\text{A.15})$$

$$E_{dj}^* = M_{dj}^c P_j^{H^{\sigma_c}} M_{dj}^{\mathcal{E}} (P_j^G \theta_{dj}^E)^{\sigma_{\mathcal{E}}} \times \frac{\alpha_{dj}^{\mathcal{C}}(I_{dj} + \Upsilon_d)}{P_{dj}^{\mathcal{C}}} \quad (\text{A.16})$$

$$G_{dj}^* = M_{dj}^c P_j^{H^{\sigma_c}} M_{dj}^{\mathcal{E}} (P_j^E \theta_{dj}^G)^{\sigma_{\mathcal{E}}} \times \frac{\alpha_{dj}^{\mathcal{C}}(I_{dj} + \Upsilon_d)}{P_{dj}^{\mathcal{C}}}. \quad (\text{A.17})$$

A.3.2 Marginal effect of climate on comfort demand. Here we

see how changes in climate affect demand for electricity, gas, and housing. The derivative of electricity demand with respect to climate variable z^l is

$$\frac{\partial E}{\partial z^l} = \frac{\kappa_E}{\rho_{\mathcal{E}}(1 - \rho_{\mathcal{E}})} E \left[\rho_{\mathcal{E}} + \frac{P^E E}{P^{\mathcal{E}} \mathcal{E}} \left(\frac{1 - \rho_{\mathcal{E}}}{1 - \rho_{\mathcal{C}}} \left(1 - \rho_{\mathcal{C}} \frac{P^{\mathcal{E}} \mathcal{E}}{P^{\mathcal{C}} \mathcal{C}} \right) - 1 \right) \right] \quad (\text{A.18})$$

This effect has three main pieces, the first term in brackets comes from the direct effect of increases in z^l increasing θ_E , this will always be positive. The second effect comes from the change in energy demand, which moves in the opposite direction as the price of energy (which moves in the same direction as the price of comfort). The final effect comes from increasing the denominator of the “share” function, $M_{\mathcal{E}}$, which will always be the opposite sign as the effect from changes in energy demand.

Changes in gas demand are similar, as the conditional demand function for gas has the same denominator of the share function and energy, but those are the only two effects. Thus, if energy price rises, decreasing the demand for energy, this will have a negative effect on gas demand. Additionally, the increases in the denominator of the share function for energy will further decrease demand for gas

$$\frac{\partial G}{\partial z^l} = \frac{\kappa_E}{\rho_{\mathcal{E}}(1 - \rho_{\mathcal{E}})} G \frac{P^E E}{P^{\mathcal{E}} \mathcal{E}} \left[\frac{1 - \rho_{\mathcal{E}}}{1 - \rho_{\mathcal{C}}} \left(1 - \rho_{\mathcal{C}} \frac{P^{\mathcal{E}} \mathcal{E}}{P^{\mathcal{C}} \mathcal{C}} \right) - 1 \right]. \quad (\text{A.19})$$

Changes in housing demand depend on the change in the denominator of the share function for comfort and changes in comfort demand. If the price of comfort increases, then comfort demand will decrease, thereby decreasing housing demand. Additionally, if the price of energy increases, then this will increase (decrease)

the denominator of the share function if housing and energy are complements (substitutes), causing housing demand to decrease (increase).

$$\frac{\partial H}{\partial z^l} = \frac{\kappa_E}{\rho_E} H \frac{P^E E}{P^C C} \left[\frac{1}{1 - \rho_C} \frac{P^H H}{P^E E} + 1 \right] \quad (\text{A.20})$$

A.3.3 Firm FOC derivation. In this section, we give the details on solving the firm's problem. Plugging the labor aggregator (2.13) into the production function (2.12), the firm's profit function is

$$\pi_j = B_j K_j^\alpha (\lambda_j S_j^{\rho_l} + (1 - \lambda_j) L_j^{\rho_l})^{\frac{1-\alpha}{\rho_l}} - W_j^S S_j - W_j^L L_j - r K_j.$$

Each firm chooses capital K_j , college educated labor S_j , and non-college educated labor L_j to maximize profit. Taking the derivative with respect to each choice, we have the following first order conditions,

$$\begin{aligned} r &= \alpha \frac{Y_j}{K_j} \\ W_j^S &= (1 - \alpha) \left(\frac{Y_j}{\mathcal{L}_j} \right) \mathcal{L}_j^{1-\rho_l} \lambda_j S_j^{\rho_l-1} \\ W_j^L &= (1 - \alpha) \left(\frac{Y_j}{\mathcal{L}_j} \right) \mathcal{L}_j^{1-\rho_l} (1 - \lambda_j) L_j^{\rho_l-1}. \end{aligned}$$

We assume capital is supplied perfectly elastically on the international market at rate $\bar{r}$. Thus, capital demand is $K_j^* = \frac{\alpha Y_j}{\bar{r}}$. Plugging this into the production function yields $Y_j = B_j^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{\bar{r}} \right)^{\frac{\alpha}{1-\alpha}} \mathcal{L}_j$. Letting $\tilde{B}_j = (1 - \alpha) B_j^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{\bar{r}} \right)^{\frac{\alpha}{1-\alpha}}$, we have the first order conditions in equation (2.14).

A.3.4 Labor Demand Parameters. In this section, we describe our calibration and estimation of the labor demand parameters: labor elasticity of substitution σ_l and the labor input use intensities, λ_j . We suppress t -subscripts as

with the above sections for exposition, but use the following procedure for each of our sample years. First, we calibrate $\sigma_l = 2$ based on Card (2009). Then we use the relative labor demand curves to identify λ_j . The relative labor demand curves are given by:⁵

$$\underbrace{\log\left(\frac{W_j^S}{W_j^L}\right)}_{\text{Estimated}} = \underbrace{-\frac{1}{\sigma_l}}_{\text{Calibrated}} \underbrace{\log\left(\frac{S_j}{L_j}\right)}_{\text{Data}} + \underbrace{\log\left(\frac{\lambda_j}{1 - \lambda_j}\right)}_{\text{Unknown}}. \quad (\text{A.21})$$

Note that the only unknowns in equation A.21 are the λ_j values. We can solve for these as:

$$\lambda_j = \frac{K_j}{1 + K_j}, \quad (\text{A.22})$$

where

$$K_j = \left(\frac{W_j^S}{W_j^L}\right) \left(\frac{S_j}{L_j}\right)^{1/\sigma_l}.$$

Next, we can solve for the firm's TFP using either first order condition for college or non-college labor as there is only one unknown.

A.3.5 Housing Supply Calibration. In this section, we describe the calibration of the rental supply parameters. First, we calibrate the inverse supply elasticities to those in Saiz (2010), specifically the inverse of the elasticities reported in Table VI. Some of our cities, Sacramento, Honolulu, and the census division aggregates, are not included in this table—we set the elasticity for Sacramento and the census division aggregates equal to the mean elasticity and calculate an elasticity for Honolulu using estimates from column (4) of Table III, setting land-

⁵Note that $\rho_l - 1 = -\frac{1}{\sigma_l}$.

unavailable equal to the maximum and the WRI to the mean. The remaining parameters from the housing supply curve, equation (2.15), are the intercepts, a_j . Since housing supply equals housing demand in equilibrium, we back these out from the data as

$$\exp a_{jt} = \frac{P_{jt}^H}{(\sum_d N_{djt} H_{djt})^{\zeta_j}},$$

where N_{djt} is the count of households in demographic group d and city j in year t , and H_{djt} is housing demand.

A.3.6 Equilibrium Definition. An equilibrium is characterized by households maximizing utility, firms maximizing profits, and the labor and housing markets clearing. Specifically,

- (1) **Utility Maximization.** Households make optimal location and comfort demand choices given wages, comfort prices, and amenities. These household choices give us aggregate housing, electricity, and gas demand, and labor supply. Household location choices give the population of each demographic group in each city, $N_{dj}^* = s_{dj}^* \times N$ where s_{dj}^* are the choice shares constructed from equation (2.7), evaluated at equilibrium prices.

$$s_{dj}^* = \frac{1}{N} \sum_{i \in d} P_{ij}^*$$

Aggregate housing demand in each city will be the sum over each demographic group's housing demand, H_{dj}^* determined according to equation (A.15), multiplied by the number of households in that city-demographic group,

$$H_j^D = \sum_d N_{dj}^* H_{dj}^*$$

Similarly, energy and gas demand in each city are the sum of each demographic group's electricity and gas demand, E_{dj}^* from equation (A.16) and G_{dj}^* from equation (A.17), multiplied by the number of households in that city-demographic group,

$$E_j^D = \sum_d N_{dj}^* E_{dj}^*$$

$$G_j^D = \sum_d N_{dj}^* G_{dj}^*$$

Finally, labor supply for college and non-college educated labor in each city are the sum of efficiency units for households who live in that city,

$$S_j^S = \sum_{d \in d^S} N_{dj}^* l_d$$

$$L_j^S = \sum_{d \in d^L} N_{dj}^* l_d$$

where d^S are the college-educated demographic groups and d^L are the not-college-educated demographic groups.

(2) **Profit Maximization.** Firms maximize profits according to the first order conditions of their profit maximization problem in equation (2.14). This determines labor demand in each city for college and non-college-educated workers.

(3) **Market clearing.** In the model, labor and housing markets must clear. Labor demand must equal labor supply for both college and non-college educated workers, $S_j^D = S_j^S$ and $L_j^D = L_j^S$, in all cities j . Additionally, housing supply given by equation (2.15) must equal housing demand, $H_j^S = H_j^D$, in all cities j .

A.3.7 Equilibrium Outline. In this section, we provide an overview of how we solve for the counterfactual in which all margins can adjust. We denote the counterfactual vector of climate variables as $\tilde{\mathbf{Z}}$. Amenities are exogenous (to household location choices), so we first recalculate amenities under the new vector of climate variables $\tilde{\mathbf{Z}}$. The algorithm for obtaining the new spatial equilibrium is:

- (1) Guess a vector of choice shares for each demographic group, s_{dj}^0 where 0 denotes iteration number (not time). Use this guess to calculate demographic city population levels as $N_{dj}^0 = s_{dj}^0 \times N$. Additionally, we guess a vector of housing prices, $P_j^{H,0}$.
- (2) Use our guesses and the firms first order conditions (2.14) to calculate city-education wages. Convert those wages to income by adjusting for efficiency units and calculating unearned income from rental profits in equation (2.16).
- (3) We then use our guesses and the counterfactual climate $\tilde{\mathbf{Z}}$ with household first order conditions (2.5) and (2.4) to calculate the price of comfort and price of energy.
- (4) Since housing demand is affected by the climate as well, we then clear the housing market. Given $P_j^{H,0}$ and our implied comfort prices, energy prices, and income, we use the household first order condition from equation (A.15) to calculate housing demand, denoted as H_{dj}^0 . We aggregate housing demand to the city level, $H_j^{0,D}$ and plug this into the rental supply curve in equation 2.15, noting that in equilibrium, the quantity of housing supplied must be equal to the quantity of housing demanded. This yields a new housing price, which we denote $\tilde{P}_j^{H,0}$. We check if $\tilde{P}_j^{H,0} = P_j^{H,0}$. If they are not equal, we update the price of housing as $P_j^{H,1} = \alpha \tilde{P}_j^{H,0} + (1 - \alpha) P_j^{H,0}$ for $\alpha \in (0, 1)$

and start again at the beginning of this step. Once the housing market clears, we have a new vector of housing prices, which we can use to calculate comfort prices.

- (5) Re-calculate the city choice shares given the new income and comfort prices.

We denote these new choice shares as $\tilde{s}_{dj}^0$. Check if $\tilde{s}_{dj}^0 = s_{dj}^0$. If not, we update our guess of the shares as $s_{dj}^1 = \alpha \tilde{s}_{dj}^0 + (1 - \alpha)s_{dj}^0$ for $\alpha \in (0, 1)$ and return to step (2).

A.4 Estimation Appendix

A.4.1 Alternative specifications for effect of climate on

comfort. We could allow for heterogeneity in the effect of climate on preferences for electricity and gas by demographic group by estimating different κ 's for each demographic group. Figure A.8 estimates our main specification separately for each of the six demographic groups we have in our model. While this estimation does not provide clear evidence that there are differences by race or college education, black households may have a slightly more muted gas response to cold days. This is consistent with Doremus et al. (2022), who find similar responses in energy expenditures between high- and low-income households except among the most extreme temperature days.




Figure A.8. Estimation results for the effect of climate on comfort production

Using cubic B-splines of degree 7 with heterogeneity by demographic group. Standard errors are bootstrapped over 10,000 draws clustering by city and year.

However, there could still be heterogeneity in the effect of climate on comfort price since other parameters in the comfort price function are heterogeneous by demographic group. Figure A.9 shows the marginal effect of a day in each temperature bin on comfort price by demographic group. There are




Figure A.9. The effect of an additional day at temperature on comfort price by race.


some differences between races, though they are small in magnitude and within the bootstrapped confidence intervals.

A.4.2 Energy and comfort prices. Figure A.10 shows how log comfort price is correlated with log rent and log energy prices by demographic group. We expect comfort prices to be determined largely by rent, but also be affected by energy prices, since rent makes up a much larger share of expenditure than energy does. This is exactly how it turns out in our model, with comfort prices more tightly correlated with rent than energy prices, but comfort prices are increasing in both rent and energy prices for all demographic groups.


Table A.3 shows average comfort and housing expenditure, as well as degree days for selected cities in 2019, ordered by the percent difference between comfort and housing expenditure. We can see that generally, the places with a larger difference between comfort and housing expenditure are cities with more extreme climates, as measured by heating and cooling degree days. Thus, climate is an

Figure A.10. Rent and Energy Prices versus Comfort Price by demographic group


(a) Comfort Price vs Rent



(b) Comfort Price vs Electricity Price



(c) Comfort Price vs Natural Gas Price



Each observation is a city, size represents the 1990 population of each city.

important part determining the cost of comfort for households beyond just the rent that they pay. We show in the appendix, figure A.10, that comfort price is positively correlated with both rent and energy prices, as we would expect.

City	Rank	Comfort	Housing	% Diff	HDD	CDD
Highest						
Birmingham, AL	1	9,640	6,954	28	1,347	1,244
Memphis, TN	2	10,409	7,553	27	1,684	1,255
Youngstown, OH	3	7,847	5,744	27	3,364	407
Tulsa, OK	4	9,217	6,910	25	1,976	1,129
Toledo, OH	5	9,646	7,239	25	3,291	540
Middle						
Atlanta, GA	33	13,325	11,003	17	1,493	1,121
Allentown, PA	34	13,270	10,964	17	3,083	515
New Orleans, LA	35	11,968	9,912	17	596	1,873
Worcester, MA	36	14,439	11,966	17	3,561	366
Virginia Beach, VA	37	13,316	11,061	17	1,727	1,094
Lowest						
Los Angeles, CA	66	19,887	18,243	8	586	788
San Diego, CA	67	21,218	19,806	7	661	571
Oxnard, CA	68	22,177	20,802	6	824	436
San Francisco, CA	69	24,792	23,256	6	1,351	151
San Jose, CA	70	26,357	24,824	6	1,206	340

Table A.3. Cities ranked by the percent difference between their average comfort expenditure and average housing expenditure in 2019.

Tables A.4 shows an alternative version of A.3, but using comfort prices and rent instead of expenditures. Results are largely similar, places with more extreme temperatures see a larger difference between their comfort price and rent. Note that rent here is the per-unit cost of housing, as described in section A.1.2.3.

City	Rank	Comfort Price	Rent	% Diff	HDD	CDD
Highest						
Youngstown, OH	1	7,589	4,399	42	3,514	475
Birmingham, AL	2	8,945	5,222	42	1,468	1,200
Buffalo, NY	3	10,246	6,044	41	3,840	365
Springfield, MA	4	10,950	6,483	41	3,581	356
Toledo, OH	5	8,205	4,876	41	3,483	529
Middle						
Albany, NY	33	12,563	7,940	37	3,988	283
Charlotte, NC	34	11,205	7,106	37	1,926	891
Knoxville, TN	35	9,013	5,721	37	2,423	634
Kansas City, MO	36	10,170	6,458	36	2,684	908
Baltimore, MD	37	14,342	9,124	36	2,854	645
Lowest						
San Francisco, CA	66	22,345	15,790	29	1,230	465
Miami, FL	67	15,053	10,701	29	113	2,268
Seattle, WA	68	16,778	11,952	29	3,899	38
Austin, TX	69	12,680	9,063	29	1,014	1,856
San Jose, CA	70	24,121	17,426	28	1,276	569

Table A.4. Cities ranked by the percent difference between their average comfort price and average housing price in 2019.

Additionally, A.5 shows the top and bottom cities by energy $\mathcal{E}$ and comfort $\mathcal{C}$ demand. This is comparable to Table 6 from Quigley and Rubinfeld (1989), who estimate "produced comfort" and "dwelling comfort", which are similar to our estimates for energy and comfort respectively. The results qualitatively match, showing that cities with mild climates spend much less on energy than those with more extreme climates.


Rank	City	Energy	Rank	City	Comfort
1	Memphis, TN	1,344.23	1	Austin, TX	1.32
2	Phoenix, AZ	1,153.55	2	Orlando, FL	1.31
3	Oklahoma City, OK	1,146.01	3	San Antonio, TX	1.31
4	Tulsa, OK	1,138.82	4	Bridgeport, CT	1.30
5	San Antonio, TX	1,114.08	5	Phoenix, AZ	1.30
6	New Orleans, LA	1,089.29	6	Tucson, AZ	1.26
7	Dallas, TX	1,073.29	7	Tampa, FL	1.25
8	Houston, TX	1,064.20	8	Miami, FL	1.24
9	Orlando, FL	1,053.08	9	Seattle, WA	1.23
10	Miami, FL	1,040.23	10	Washington, DC	1.23
61	Providence, RI	479.21	61	Oxnard, CA	1.08
62	Sacramento, CA	468.82	62	Los Angeles, CA	1.07
63	Boston, MA	461.41	63	Kansas City, MO	1.06
64	Albany, NY	461.41	64	Memphis, TN	1.06
65	Los Angeles, CA	453.56	65	Buffalo, NY	1.06
66	Honolulu, HI	438.53	66	Syracuse, NY	1.04
67	San Diego, CA	401.64	67	Albany, NY	1.03
68	Oxnard, CA	381.63	68	Youngstown, OH	1.02
69	San Jose, CA	378.38	69	Rochester, NY	1.02
70	San Francisco, CA	359.54	70	Grand Rapids, MI	0.99

Table A.5. Cities ranked by their average aggregate energy demand $\mathcal{E}$ and comfort demand $\mathcal{C}$ in 2019.

A.4.3 Energy Efficiency: Evidence. In this section, we provide evidence of heterogeneity in energy-efficiency by demographic group and across time. Specifically, we plot city-level densities of aggregate energy use $\mathcal{E}$ divided by housing, H . In each figure, $\mathcal{E}$ is calculated using equation 2.2 evaluated at estimates from column 1 of Table 1. Housing is calculated as in Section A.1.2.3. We interpret a larger value of this ratio as loosely indicative that homes in a particular city/demographic group are less energy-efficient.⁶


⁶However, it could also be reflect differences in the strength of preferences for housing relative to energy.

Figure A.11. Energy-efficiency by education group across years and CBSAs




Notes: Energy-efficiency by education group across years and CBSAs.

Figure A.12. Energy-efficiency by education-group across years and CBSAs



Notes: Energy-efficiency by education group across years and CBSAs.

Figure A.13. Energy-efficiency by demographic-group across years and CBSAs



Notes: Energy-efficiency by demographic group across years and CBSAs.

Table A.6. Household moving costs estimated from Equation (2.31) for 1990.

Standard errors are in parenthesis. We estimate standard errors using numerical derivatives. Distance is measured in thousands of miles.

	Non-College			College		
	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$
White	1.433 (1.9e-5)	-6.351 (0.3e-5)	1.933 (0.1e-5)	1.336 (6.4e-5)	-4.573 (1.0e-5)	1.399 (0.2e-5)
Black	1.758 (0.2e-3)	-6.246 (0.5e-3)	2.144 (0.6e-3)	1.699 (2.5e-3)	-4.675 (0.6e-3)	1.592 (0.3e-3)
Rest	0.898 (0.8e-2)	-3.225 (0.3e-2)	0.471 (0.2e-2)	1.256 (3.8e-2)	-1.755 (3.6e-2)	0.192 (1.0e-2)

Table A.7. Household moving costs estimated from Equation (2.31) for 2000.

Standard errors are in parenthesis. We estimate standard errors using numerical derivatives. Distance is measured in thousands of miles.

	Non-College			College		
	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$
White	1.378 (1.5e-5)	-6.556 (0.3e-5)	2.054 (0.1e-5)	1.332 (4.1e-5)	-4.693 (0.7e-5)	1.461 (0.1e-5)
Black	1.774 (0.01e-2)	-6.416 (0.04e-2)	2.223 (0.06e-2)	1.679 (0.11e-2)	-4.858 (0.04e-2)	1.681 (0.017e-2)
Rest	0.984 (0.14e-2)	-2.993 (0.06e-2)	0.451 (0.03e-2)	1.444 (0.47e-2)	-1.414 (1.34e-2)	0.157 (1.83e-2)

Table A.8. Household moving costs estimated from Equation (2.31) for 2010.

Standard errors are in parenthesis. We estimate standard errors using numerical derivatives. Distance is measured in thousands of miles.

	Non-College			College		
	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$
White	1.329 (1.6e-5)	-6.263 (0.3e-5)	1.903 (0.1e-5)	1.352 (3.0e-5)	-4.702 (0.5e-5)	1.45 (0.1e-5)
Black	1.81 (1.1e-4)	-6.301 (3.3e-4)	2.156 (6.6e-4)	1.729 (4.1e-4)	-4.89 (1.7e-4)	1.684 (0.8e-4)
Rest	1.073 (0.1e-2)	-2.964 (0.1e-2)	0.435 (0.1e-2)	1.528 (1.2e-2)	-1.546 (1.7e-2)	0.188 (1.1e-2)

Table A.9. Household moving costs estimated from Equation (2.31) for 2019.

Standard errors are in parenthesis. We estimate standard errors using numerical derivatives. Distance is measured in thousands of miles.

	Non-College			College		
	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$	$\tilde{\gamma}^{st}$	$\tilde{\gamma}^{dist}$	$\tilde{\gamma}^{dist2}$
White	1.271 (1.8e-5)	-6.263 (0.4e-5)	1.899 (0.1e-5)	1.346 (2.7e-5)	-4.794 (0.5e-5)	1.492 (0.1e-5)
Black	1.77 (1.2e-4)	-6.656 (3.6e-4)	2.284 (6.5e-4)	1.659 (3.9e-4)	-5.074 (1.6e-4)	1.75 (0.7e-4)
Rest	1.194 (0.1e-2)	-2.863 (0.3e-2)	0.428 (1.2e-2)	1.627 (0.2e-2)	-1.429 (0.4e-2)	0.184 (0.3e-2)

A.4.4 Moving Cost Parameters for All Years.

A.4.5 Climate amenities. Figure A.14 shows the effect of temperature on climate amenities using different measures of temperature—splines, temperature bins, and degree days.

A.4.6 Model fit. Figure A.15 shows model fit for rents and wages when comparing baseline simulations to actual data.




Figure A.14. Estimate of MWTP as a percent of income for an additional day at the given temperature

normalized to a day at 65 degrees F. This is calculated by rescaling $\hat{\beta}_e^Z(\tau)$ by β_e^w . Standard errors are calculated with clustered bootstrapping using 10,000 draws, clustered by city-year. We have limited the range of temperature to between the 1st and 99th percentile of average daily temperature, weighted by 1990 population.




Figure A.15. Model fit when using observed climate to simulate the model.

Each observation is a demographic group for a particular city and a particular baseline year (1990, 2000, 2010, or 2019). The size of dots represents the number of households in that demographic group-cbsa in the actual data.

A.5 Additional Results

Figure A.17 shows the full distribution of CV across all households.

Table A.10 has the welfare effects across different decompositions in levels for both white and black households.




Figure A.16. Change in prices and demand across states.

Averages are weighted by population. For states without any CBSAs in them, we use the value for the census division in which the state is located.




Figure A.17. Compensating Variation by race in 2019 if the climate was the same as in 1990.

Table A.10. Decomposition of welfare effects. Values are CV as a percent of income.


	Climate Effect From		
	Both	Amenities	Comfort
Fixed Location			
White	5.09	4.70	7.63
Black	-0.52	-0.64	2.57
Difference	-5.61	-5.34	-5.05
Fixed Location, adjust comfort			
White	-2.16	4.70	0.38
Black	-2.97	-0.64	0.13
Difference	-0.81	-5.34	-0.25
Sorting, fixed prices			
White	-2.59	4.36	0.38
Black	-3.45	-1.00	0.12
Difference	-0.87	-5.35	-0.26
Full Effects			
White	-2.48	-2.82	0.37
Black	-2.95	-3.17	0.25
Difference	-0.47	-0.35	-0.12

Table A.11 shows the top and bottom 10 cities by change in population under this historical climate today counterfactual. We see that people generally leave the places that are getting hotter—such as Florida and California, and move to the places that become more pleasant, in the upper Midwest.

Table A.11. Changes in population using 1990 climate as a counterfactual in 2019

City	Change in Pop (%)
Growing Cities	
Grand Rapids-Wyoming, MI	23.4
Hartford-West Hartford-East Hartford, CT	17.0
San Antonio, TX	16.2
Cincinnati-Middletown, OH-KY-IN	15.2
Denver-Aurora, CO	14.8
Tulsa, OK	14.5
Oklahoma City, OK	14.5
Kansas City, MO-KS	14.4
Albany-Schenectady-Troy, NY	14.2
New Haven-Milford, CT	14.1
Declining Cities	
Richmond, VA	-2.9
Youngstown-Warren-Boardman, OH-PA	-4.5
Columbus, OH	-4.7
South Atlantic Division	-4.7
Omaha-Council Bluffs, NE-IA	-5.0
West South Central Division	-7.0
St. Louis, MO-IL	-7.1
Dayton, OH	-8.4
Houston-Baytown-Sugar Land, TX	-9.5
Louisville, KY-IN	-9.6

Climate change has contributed to the increase in the college wage gap. Figure A.18a shows the distribution of wage changes across cities. We find that wages increase for college households and decrease for non-college households. To see why, recall that the firm's FOC (2.14) implies that the ratio of college to non-



(a) Change in wages by education level. (b) Effect of worker ratio on wage ratio.

Figure A.18. College and non-college wages comparing present day climate to that of 1990.

The worker ratio is college workers divided by non-college workers. The size of the points reflect a city's baseline population.

college workers pins down the wage ratio, and these ratios are inversely related since college and non-college workers are imperfectly substitutable ($\rho_l < 1$). Since college households are more mobile than non-college households, the college-noncollege worker ratio tends to decrease, as shown in Figure A.18b—more so in cities where the climate is deteriorating. Thus, college wages increase in those deteriorating cities, and non-college wages decrease.

A.5.1 IRA Subsidies. Figure A.19 shows the share of each state's population that lives in a disadvantaged community according to census block group level population from the 2020 Census.




Figure A.19. Share of the population living in a disadvantaged community.

We use the EPA’s Disadvantaged Communities Map to define census tracts as disadvantaged. The total population in each census tract comes from the 2020 Decennial Census (Walker & Herman, 2024).




Figure A.20. Distribution of baseline CV by income group.

The X-axis is a household’s percentile in CV from the baseline simulation comparing 2019 and 1990 climates. We split the figure into three income bins, below 10th percentile, 10th-50th percentile, and above 50th percentile based on household expected income in the baseline simulation.




Figure A.21. CV for 2019 climate with and without subsidies relative to the 1990 climate without subsidies.

The X-axis is a household's decile in CV from the baseline simulation comparing 2019 and 1990 climates. The Y-axis is average CV under 2019 climate relative to without subsidies under the 1990 climate. We split the figure into three income bins, below 10th percentile, 10th-50th percentile, and above 50th percentile based on household expected income in the baseline simulation. See description in the text on the subsidy distributions




Figure A.22. CV for 2019 climate with and without subsidies relative to the 1990 climate without subsidies for different magnitudes, G .

The X-axis is a household's decile in CV from the baseline simulation comparing 2019 and 1990 climates. The Y-axis is average CV under 2019 climate relative to without subsidies under the 1990 climate. We split the figure into three income bins, below 10th percentile, 10th-50th percentile, and above 50th percentile based on household expected income in the baseline simulation. See description in the text on the subsidy distributions




Figure A.23. Percent change in population under different subsidy program magnitudes and distributions.

The underlying data are city-level, and the plotted lines are smoothed fit from a generalized additive model weighted by city population. The X-axis represents a city's rank in change in climate amenities from 1990 to 2019. The Y-axis is the percent change in population divided by the total size of the program, and the different sizes are differentiated by color.




Figure A.24. Percent change in rents under different subsidy program magnitudes and distributions.

The underlying data are city-level, and the plotted lines are smoothed fit from a generalized additive model weighted by city population. The X-axis represents a city's rank in change in climate amenities from 1990 to 2019. The Y-axis is the percent change in rents divided by the total size of the program, and the different sizes are differentiated by color.




Figure A.25. Percent change in wages under different subsidy program magnitudes and distributions.

The underlying data are city-education-level, and the plotted lines are smoothed fit from a generalized additive model weighted by city population. The X-axis represents a city's rank in change in climate amenities from 1990 to 2019. The Y-axis is the percent change in wages divided by the total size of the program, and the different sizes are differentiated by color. We differentiate between college and non-college educated wages.

APPENDIX B

APPENDIX FOR GLYPHOSATE EXPOSURE AND GM SEED ROLLOUT UNEQUALLY REDUCED PERINATAL HEALTH

B.1 Background

Genetically modified crops Monsanto developed the first genetically modified crops, releasing GM soy, corn, and cotton in 1996 in the United States. These plants are resistant to glyphosate, allowing farmers to spray their fields with glyphosate to kill weeds but not harm their crops. The pairing of GM seeds with glyphosate provides a simple and effective method for controlling weeds—previously, farmers had to use different pesticides, each effective on a smaller subset of weeds at different points in the cultivation process. This herbicide portfolio was supplemented by mechanical tilling. glyphosate previously had to be used sparingly since it would also kill the crops themselves. Farmers adopted GM seeds rapidly in the United States. The USDA provides data on national GM adoption rates starting in 1996 and for specific states beginning in 2000. Figure B.2a shows the time series of adoption rates for the entire country (dark, bold line) and particular states (light, gray lines). In 2000, just four years after their release, GM seeds constituted 54 percent of soy acres, 61 percent of cotton acres, and 25 percent of corn acres. By 2010, adoption rates were around 90 percent for all three crops. Adoption of GM corn was generally slower and more heterogeneous across states than for either soy or cotton. Figure B.2b shows spatial variation in adoption rates in 2000, 2005, and 2010. For corn and soy, states further west adopted slightly faster than states further east. Meanwhile, California and Texas adopted GM cotton slower than the Southeast.

Glyphosate and health Glyphosate is a broad-spectrum herbicide discovered and commercialized by Monsanto in the 1970s. Its popularity grew over the next twenty years because of its relatively favorable properties. glyphosate has a low toxicity relative to other chemicals used on farms. It breaks down fairly quickly and binds to the soil, decreasing runoff (Duke & Powles, 2008). However, it is water-soluble, which means that the part that does not bind to soil enters the water supply (Baer & Marcel, 2014). It is an effective weed killer, working on a broad spectrum of plants. However, glyphosate does not just kill weeds, it also kills fungi and microorganisms in the soil, which can lead to the crops being susceptible to disease (B. Muller, 2020). It also breaks the nutrient cycle, forcing farmers to increase their dependence on fertilizer to feed their crops (Barrows, Sexton, & Zilberman, 2014). Farmers in the US spend nearly \$8 billion on pesticides each year (Grube, Donaldson, Kiely, & Wu, 2011), applying glyphosate to 298 million acres of crops annually (U.S. EPA, 2019).

Glyphosate and health The US EPA’s current approval process for pesticides provides ample opportunities for applicants to steer the process toward approval.

A central critique in the pesticide-regulation literature is the EPA’s tendency to rely upon regulated entities to design, run, and analyze non-peer-reviewed experiments to test chemicals’ safety—ignoring clear conflicts of interest (C. Benbrook, 2020; C. M. Benbrook, 2018; Sheppard et al., 2020). This approach stands in stark contrast to the International Agency for Research on Cancer (IARC)’s reliance on published, peer-reviewed research C. Benbrook (2020); C. M. Benbrook (2018).

EPA’s reliance on applicant-generated tests grants applicants the opportunity to influence evidence in the review process in several ways. First,

this process places regulated entities in a position to effectively selectively report their way to approval by running multiple tests and only reporting studies that show no harm (C. Benbrook, 2020). Benbrook (2019) illustrates considerable differences between regulatory assays versus assays from peer-reviewed journals (C. M. Benbrook, 2019)—leading the EPA and IARC to conflicting conclusions. Several recent reports highlights similar concerns around regulation in the EU—especially in differing conclusions between the IARC and the European Food Safety Authority (EFSA) Clausing, Robinson, and Burtscher-Schaden (2018); Portier et al. (2016).

Among reported tests, regulated entities may test exposure levels lower than (1) levels encountered in occupational settings and (2) levels where adverse health effects may be observed—especially when chemicals may have nonlinear dose-response relationships (Davoren & Schiestl, 2018). Benbrook (2020) also highlights that the EPA’s approval process focuses more on dietary consumption than on occupation exposure—a potentially large shortcoming for Roundup’s high dermal penetration and possible wand application (C. Benbrook, 2020). Benbrook suggests that the regulatory process surrounding glyphosate has largely failed “to add common-sense worker protection provisions” (C. Benbrook, 2020).

Similarly, the review process may conclude with *no harm* due to flawed research designs (vom Saal et al., 2010), implementation (Mie et al., 2018), and/or analyses (Sheppard et al., 2020). Despite its potential impacts on public health, the data involved in these studies is often withheld (Davoren & Schiestl, 2018)—further restricting review and oversight. Compounding the issue: scant post-approval monitoring of chemicals exposure in populations or the environment

(C. M. Benbrook, 2018). Finally, there is the issue of regulatory capture, which we leave for future research.

While the conflict of interest is clear, these mistakes may also follow from the studies' lack of legitimate peer review. Ultimately, the EPA's reliance on industry-conducted studies, in conjunction with the effective null hypothesis of *no harm*, opens the door for avoidable public health risks.

B.2 Data

Fertilizer data We get fertilizer data from the USGS (Falcone, 2021), which estimates county-level nitrogen and phosphorous every five years between 1950 and 2017 using data from the USDA Census of Agriculture. They report separate estimates for farm commercial fertilizer applications, non-farm commercial fertilizer applications, and nutrient loads from manure. Since the rest of our data are annual, we interpolate for the years between the years of the Ag Census. We fit splines separately for each county and fertilizer type using all available years of data and then generate annual predictions for each county-year-fertilizer type. Figure B.3 shows the results for a few example counties, where the points are the actual data and the lines are the fitted splines. The raw data only separates farm and non-farm commercial fertilizer use from 1987. Non-farm commercial fertilizer use is small, only making less than three percent of commercial fertilizer use. Therefore, we use four variables as controls: phosphorous from commercial uses, nitrogen from commercial uses, phosphorous from manure, and nitrogen from manure. We normalize all of these variables by the total size of the county.

Crop Acreage and Yield Data We get crop acreage and yield data from the USDA NASS. Specifically, we get acres planted and yield per acre by county from 1985 to 2017 from their annual survey. These data include barley, several types

of beans, canola, chickpeas, corn, upland cotton, flaxseed, lentils, mustard, oats, peanuts, peas, peppers, potatoes, rice, safflower, sorghum, soybeans, sugarbeets, sweet corn, sweet potatoes, tomatoes, and wheat. We aggregate all crops besides corn, soy, and cotton into an “other” category. The USDA masks data for counties where there are small amounts of farms producing a particular crop. They report acres and yield from those counties as an aggregate for the agricultural statistical district. To account for this, we allocate any production reported at the aggregate district level to counties in that district with missing data proportional to the size of the county.

Accuracy of FAO-GAEZ data New work highlights several issues with the GAEZ suitability data, calling into question the ability of the GAEZ suitability measures to predict historical yield for different crops (Rhode, 2024). While we share many of the concerns raised—notably the lack of clarity around data sources and model validation—we think that the use of the GAEZ data is still appropriate in our setting. We are not using the GAEZ data to predict historical yields but rather as an instrument for glyphosate use. Thus, we only need the GAEZ suitability to be correlated with increases in glyphosate and for the exclusion restriction to hold. See *Methods* for an in-depth discussion of our identifying assumptions. Table B.2 demonstrates that the GAEZ suitability measures correlate with observed acreage and yield of corn, soy, and cotton prior to the introduction of GM crops.

Other Data We collect data on county-level, annual estimates of total population by age and race from SEER (Surveillance, Epidemiology, and End Results (SEER) Program, 2024). Age shares are done by decade from ages 0 to 70,

with everyone older than 70 grouped into the final bin. Race shares are calculated for the white, Black, and hispanic populations.

B.3 OLS results

Panels A and B of Table B.7 contain results for our main specifications, but estimated with OLS rather than 2SLS. We find precise null effects across all outcomes, demonstrating the importance of isolating exogenous variation in glyphosate using our instruments.

B.4 Comparing rescaled reduced-form DID results with 2SLS results

Our reduced-form DID results effectively provide estimates of 2SLS's first stage and reduced form—just in a DID framework without interacting the instrument with indicators for year. By dividing the reduced-form DID results for health outcomes by the results, we can get comparable estimates for the health effects due to glyphosate exposure. Table B.4 provides these calculations—first repeating the results from the main text in Figure 14 and then rescaling the reduced-form DID health effects by the reduced-form DID glyphosate effect. The resulting estimates are very close to the 2SLS estimates.

B.5 Shift-share specification

We can recast our identifying variation to be used similarly to that of a traditional “shift-share” specification, where the “shift” is national glyphosate use and the “share” is attainable yield in each county. Thus, the identifying variation is very similar to our main results—we get temporal variation driven by the nation-wide increase in glyphosate use after the release of GM crops, and we get spatial variation from the suitability of the land in each county for corn, soy, and cotton. In the shift-share specification, our instruments are national glyphosate and national glyphosate interacted with the attainable yield percentile

for corn, soy, and cotton. Thus, the difference between this specification and our main specification is that the first stage uses the national glyphosate trend directly, rather than interacting attainable yield with year dummies. When calculating the national glyphosate for each county, we exclude glyphosate applied within 100km of the county and any glyphosate applied upstream of the county to ensure that the national glyphosate instrument satisfies the exclusion restriction—that national glyphosate only affects perinatal health through its effect on local glyphosate. Table B.8 shows the results, which are generally similar, but smaller in magnitude than our main results.

B.6 Robustness of first-stage and reduced form results

Figure B.18 estimates our model for births to mothers with rural and non-rural residences separately. There is a small but largely insignificant decrease in birthweight after the release of GM seeds in 1996 in high GM attainable yield counties relative to low GM attainable yield counties. However, this effect is gone by 2010. We attribute this difference largely to measurement error in exposure—we do not think that exposure is very high for urban mothers, who are unlikely to be in contact with drift, dust, or water contaminated with glyphosate applied within that county. The lack of direct measurement of exposure to glyphosate is a weakness of our study, as all we know is the amount of glyphosate used in a county each year.

B.7 Robustness of 2SLS results for other outcomes

We provide specification charts for non-birthweight outcomes in Figures B.19 and B.20. Additionally, we test robustness to various spatial subsets in Figure B.21.

B.8 Prediction performance

Figure B.8 depicts how predicted birthweight percentiles maps to predicted (dark blue) and actual (light orange) birthweight. Given a predicted percentile, the mean predicted birthweight is, on average, quite close to the true mean birthweight—suggesting the ML approach indeed captures informative variation in birthweight.

Table B.6 describes the performance of the random-forest model for predicting birthweight (the outcome for which the model trained) and birthweight percentile, decile, and quintile. Panel A evaluates predictions' performances on two metrics—mean absolute error (MAE: $|y - \hat{y}|$) and mean absolute percent error (MAPE: MAE/y). All results in the table focus on the set of infants born to rural-residence mothers, as this group matches the population of interest in the paper. The three rows of Panel A use this full sample of rural-residence mothers and then drop 1% and 5% of the top and bottom birthweights. Finally, Panel B provides two null models for comparison. The *Means* model always predicts the sample mean; the *At random* model predicts at random.

For predicting birthweight, the ML predictions offer improvements over both null models on both metrics (error level and percentage), and this dominance is even stronger when not evaluating the tails of the distribution. While the *Means* null model slightly edges out the full-sample model in terms of MAE for the percentile-based outcomes, the ML-based model is *much* stronger in terms of percent error. Percentage error is an important metric in our application, as we want match births to expected birthweight groups (as in Fig B.8) rather than simply a low-error on-average prediction.

B.9 2SLS results and the shape of the glyphosate damage function

Our estimates for the effect of glyphosate exposure on perinatal health make comparisons between communities with higher and lower levels of glyphosate exposure (instrumented communities' GM-crop suitability, compared to pre-GM-rollout comparisons, and conditional on fixed effects). Following the GM rollout, many low-suitability communities also experienced increased glyphosate exposure (for example, as shown in Figure 12c). As a result, our estimates for glyphosate's health damages compare infants with higher glyphosate exposure to infants with lower glyphosate exposure. Generally, infants with *lower* levels of glyphosate exposure are generally still exposed to some non-zero amount glyphosate. How this non-zero exposure for the lower group affects our 2SLS estimates depends upon two items: the shape of glyphosate's damage function and the interpretation of the 2SLS estimates.

For the moment, ignore the temporal variation in our instrument and estimation. The reduced-form results from our 2SLS estimates effectively compare glyphosate's damages between in high and low GM-suitability areas, i.e. $\mathcal{D}(g_h) - \mathcal{D}(g_\ell)$. The first stage estimates the difference in glyphosate intensity between these areas, i.e. $g_h - g_\ell$. Finally, second-stage scales these differences by the differences in glyphosate intensity.

$$\textit{Reduced form: Damages from high vs. low glyphosate} = \mathcal{D}(g_h) - \mathcal{D}(g_\ell)$$

$$\textit{First stage: Differences in glyphosate intensity} = g_h - g_\ell$$

$$\textit{Second stage: Scaling damage difference by glyphosate difference} = \frac{\mathcal{D}(g_h) - \mathcal{D}(g_\ell)}{g_h - g_\ell}$$

There are three main cases to consider for the shape of the damage function (with respect to its second derivative).

- **Linear damages** If glyphosate’s damages for perinatal health are approximately linear/affine in glyphosate (i.e. $\alpha + \beta g$), then our second-stage estimates are unaffected by whether low-suitability areas have zero or non-zero levels of glyphosate—as long as low-suitability areas apply less glyphosate than high suitability areas (our first-stage estimates and event studies confirm this requirement is satisfied).

$$\frac{\mathcal{D}(g_h) - \mathcal{D}(g_\ell)}{g_h - g_\ell} = \frac{\alpha + \beta g_h - \alpha - \beta g_\ell}{g_h - g_\ell} = \frac{\beta(g_h - g_\ell)}{g_h - g_\ell} = \beta$$

Accordingly, if the damage function is linear, our second-stage will recover the marginal damages from a one-unit increase regardless of the level of glyphosate in low-suitability counties.

- **Concave damages** For concave damage functions, as g_ℓ increases from zero and approaches g_h from below, the reduced-form difference $\mathcal{D}(g_h) - \mathcal{D}(g_\ell)$ shrinks faster relative to the first-stage difference $g_h - g_\ell$. Consequently, higher levels of g_ℓ will generate lower second-stage (and reduced-form) estimates. If the parameter of interest is the average per-unit health damages due to moving from $g = 0$ to $g = g_h$, then the 2SLS will understate the actual damages. However, 2SLS estimates the (weighted) average marginal damage of increasing glyphosate from concentrations found in the US. More formally,

$$\frac{d}{dg_\ell} \frac{\mathcal{D}(g_h) - \mathcal{D}(g_\ell)}{g_h - g_\ell} = \left[\frac{\mathcal{D}(g_h) - \mathcal{D}(g_\ell)}{g_h - g_\ell} - \mathcal{D}'(g_\ell) \right] \frac{1}{(g_h - g_\ell)^2} \leq 0$$

due to $\mathcal{D}$ being concave (strict inequality will follow from strict concavity).

- **Convex damages** Convex damage functions simply ‘reverse’ the results of concave functions: 2SLS will overstate the average damage of moving from $g = 0$ to $g = g_h$ when $g_\ell > 0$ but will provide a (weighted) average marginal damage for glyphosate concentrations common in US counties.

Because we are not in a position to take a strong stand on the shape of glyphosate’s damage function, it likely makes more sense to consider our estimates as the marginal damages relative to glyphosate concentrations commonly encountered in the United States—a policy-relevant parameter requiring weaker assumptions.

B.10 Demographic trends

One concern for identification in our model is that the underlying composition of the population is changing in high vs. low GM attainable yield counties during the period of our study. Figure B.22 shows event studies where we use demographics of the mother as outcomes with county and year-by-month fixed effects and no other controls to test whether demographics are changing over time. We find that births in high-yield counties are less likely to come from black mothers after the release of GM crops—this would otherwise be concerning for our main estimates, however, we (1) control for race and other demographics in our main estimation and they do not meaningfully impact the results, (2) *predicted* birthweight does not change over the time period of the study, and (3) we find significant effects of glyphosate on birthweight for babies with both white and non-white mothers.

B.11 Other forms of heterogeneity

Mother’s race Based on heterogeneity in predicted birthweight, we expect there to be differences in effect by mother’s race. Fig B.23 shows reduced from event studies for different outcomes by mother’s race. Births to non-white mothers have a noisy, but generally larger effect than briths to white mothers.

Heterogeneity by month of birth These results do not exhibit consistent heterogeneity by month of birth, as seen in Figure B.25. There are slightly higher

effects during the first months of the year—which means that their gestational period began in the spring and early summer the time when the most glyphosate is applied.

Rural vs urban We compare results for rural and non-rural counties in *SI* Figure B.18. As expected, the first stage is much weaker in non-rural than for rural counties since non-rural counties have more land uses competing with agriculture. We note that non-rural counties still grow GM crops and apply glyphosate—30% of corn, soy, and cotton acres and 29% of glyphosate applications are in non-rural counties. However, the mismeasurement of glyphosate exposure for infants in non-rural counties is likely to be considerably worse than in rural counties. Mothers residing in urban portions of a non-rural county will have lower glyphosate exposure than mothers residing near that same county’s agricultural production. Yet, data constraints force us to assign all infants in a county the same level of glyphosate exposure. We estimate the effect of GM suitability in non-rural counties to be attenuated relative to rural counties, consistent with the non-rural counties having more measurement error in exposure.

B.12 Effect of GM on acreage and yield

Changes in agricultural activity unrelated to glyphosate that result from GM seed adoption could also affect infant health, threatening our identified effect of glyphosate on birth weight. For example, GM technology could lead farmers to bring marginal, not previously farmed land into agricultural production. This additional production could be associated with increased runoff into water or air pollution from dust or drift. Additionally, if yield increased with GM seeds, farmers could see an economic boost that could affect infant health. In order to rule out

these as mechanisms for the observed effect of GM attainable yield on birth weight, we explore the effect of GM attainable yield on crop acreage and actual yield.

We use USDA NASS data on annual, county-level crop acreage and yield, regressing these variables on the max GM attainable yield percentile interacted with year. Figure B.26a shows reduced-form event study results regressing our suitability measure on total crop acreage as a share of the county area. Unfortunately, 1995 seems to be a low outlier year, making the event study more challenging to interpret—however, total acreage appears to stay around the same level after 1995 as it was prior to 1995. We estimate a difference-in-difference model comparing before vs after 1995 that results in a small and not statistically significant difference in total acreage.

Figure B.26b shows corn acreage as a percent of the county area. As with total acreage, the effect of GM suitability is noisy in the pre-period and seems unchanged until around 2007, at which point there does seem to be an increase in corn acreage for high-suitability counties relative to low-suitability counties. This timing coincides with when the renewable fuel standard increased incentives for farmers to plant corn (Lark et al., 2022). Figure B.26c shows the results for soy acres as a percent of the county area. There was an initial bump in soy acreage after 1995, followed by a return to pre-1995 averages, consistent with the fact that GM seed varieties were available for soy before corn.

Meanwhile, Figure B.26d shows the effect of GM suitability on cotton acreage. Cotton seems to have had a high outlier year in 1995 but remains consistent with the other pre-period years until 2006, after which it decreases in high relative to low suitability counties. In summary, we find that there does

not seem to be an effect of GM suitability on total acreage, but this masks some substitution between crops.

We estimated models adding both fertilizers and acreage as controls. Figure B.13 shows both the first stage and reduced form event studies under various iterations of controls—adding fertilizer and/or acreage controls does not result in meaningful or statistically significant differences of estimates in either the first stage or reduced form. Notably, adding acreage controls to all of the other controls does flatten out the upward pre-period trend in the first stage. The spec charts in Figures B.6, B.19, and B.20 show our estimated coefficient on glyphosate from 2SLS on birthweight, gestation length, and the health index. Adding fertilizer controls slightly increases the magnitude of glyphosate’s effect on birthweight, but by considerably less than when we add pesticide controls. This is true when adding just fertilizers as controls relative to no additional controls and when adding fertilizers to a specification with pesticides and unemployment as controls.

B.13 Other socioeconomic outcomes

Here, we explore the relationship between our attainable yield instrument and some socioeconomic outcomes in order to rule them out as mechanisms for the measured birth weight effect. We regress the socioeconomic variables on GM attainable yield interacted with year dummies with county and year fixed effects. The sample is a county-year panel of rural counties in the US between 1990 and 2013. Figure B.28 shows the results. There is no change in farm or non-farm income, however there do appear to be changes in employment. The unemployment rate jumps after 2000—thus, we control for unemployment in our main regression, but note that this is four years after the release of GM seeds, thus the timing does not align to have been caused by GM. Meanwhile, farm employment is also

declining, however there is a clear pre-trend. The release of GM seeds does not appear to affect this trend.

B.14 Effects of upstream glyphosate in water

B.14.1 Predicting glyphosate in water with machine learning.

To measure spillover effects from glyphosate applied upstream, we must have some measure of glyphosate exposure in water. Ideally, this would come from extensive monitoring, which consistently reports pesticide concentrations in water for a comprehensive set of water sources. Unfortunately, such a monitoring network does not exist, so we must create an alternative methodology to estimate glyphosate exposure from upstream spraying. We train a machine learning model to predict glyphosate concentrations using the limited glyphosate monitoring in water, along with water flow and other environmental characteristics.

Data preparation Our training data come from recent work by Medalie *et al.* Medalie et al. (2020), who took 3204 samples of glyphosate and its main degradate AMPA from 70 sites in the National Water Quality Network (NWQN), a nationally representative set of water bodies, between 2015 and 2017. Both chemicals are nearly omnipresent, with glyphosate detected in 75 percent of samples and AMPA detected in 90 percent. We link these measurements to data on glyphosate use, soil type, slope, and rainfall upstream from the sampling location.

We use a spatial water model to aggregate the amount of glyphosate applied upstream and downstream of each sampling location. Specifically, we use the level 8 HydroBASINS product from HydroSHEDS (Lehner & Grill, 2013). These data are watershed polygons that delineate water basins across the globe in a standardized way. Importantly, they are assigned codes in a way that makes it possible to find all watersheds upstream and downstream from any given watershed.

We begin with the pesticide data. As in our local analysis, one may be concerned with the endogeneity of glyphosate use. Our estimates will be biased if spraying upstream of a sampling location correlates with other factors affecting health outcomes. We deal with this issue by using only exogenous variation in glyphosate use driven by the same instruments from our local analysis, namely that driven by the timing of the release of GM seeds and the suitability of a county for corn, soy, and cotton. We regress glyphosate on the GM attainable yield percentile interacted with year dummies, with year and county fixed effects to generate county-year level predictions of glyphosate. To disaggregate these county-level predictions into watersheds, we assume that spraying is uniform across the county and multiply the glyphosate prediction for each county by the portion of the county's total area covered by the watershed. Figure B.29 shows the spatial distribution of predicted glyphosate by watershed across the United States in 2004.

Additionally, we collect several other variables that affect the runoff of glyphosate in a method loosely following the commonly used universal soil loss equation (USLE). This soil loss equation multiplies the erodibility of the soil, the slope of the land, rainfall, and two measures associated with land use. We aggregate soil erodibility and slope from the gridded soil survey to the watershed level by taking the average over all 30-meter cells in each watershed (Soil Survey Staff, 2021). Similarly, we use gridded, monthly precipitation from PRISM to help inform the potential for glyphosate to run into water (PRISM Climate Group, 2022). We aggregate the 4-kilometer cells to the watershed level by taking the simple average of cells within a watershed. Additionally, we aggregate to the annual level by taking the sum over the growing season, April through September,

when most glyphosate is applied. Figure B.30 shows national percentiles of soil erodibility, slope, and precipitation by watershed.

We then utilize the “Pfafstetter” watershed coding system used by the HydroBASINS data to find all watersheds upstream from each watershed. We have selected an example watershed in Washington County, Illinois, just east of St. Louis, for demonstration purposes. Figure B.31 shows the example watershed in red and then highlights all of the watersheds upstream, which reach further north into Illinois, and all of the watersheds downstream, which follow the Mississippi River to the Gulf of Mexico.

When linking upstream and downstream watersheds, we calculate the distance between any two watersheds by summing the distance between centroids of each watershed that lies along the water flow between the two watersheds. We then aggregate the variables described above into 50-kilometer distance bins from -100 to 350 , where negative values denote values for downstream watersheds. Figure B.31 demonstrates the distance bins for our example watershed. The final dataset contains 2,142 water samples, where we removed 1064 samples from sites with no upstream watersheds entirely outside the site’s county. We remove these to ensure that our measure of upstream spraying does not capture non-water mechanisms of glyphosate exposure, such as dust, drift, or direct contact.

Training the water concentration ML model We train LASSO and Random Forest (RF) models using the above mentioned dataset. We generate a fully saturated set of interaction terms between glyphosate, soil erodibility, slope, and rainfall as predictors in the LASSO model. The month of the sample is the only other predictor variable. Since the model’s primary goal is to predict glyphosate concentrations back in time, we train the model on 1,385 observations

from after October 2015 and validate performance with 757 observations from before October 2015. Within the training set, we tune parameters using 4-fold cross-validation, where each fold trains on 15 months of data and then tests performance on the preceding six months of data. Then, we select the parameter with the lowest average RMSE across folds to estimate the model on the entire training set. Figure B.32 shows the cross-validation results.

We then assess performance of the tuned models using the 757 held out observations. Figure B.33 shows the out-of-sample predictions versus their actual values. Both models predict AMPA concentrations much better than glyphosate concentrations, with an R-squared of 0.59 and 0.31 for the random forest and LASSO models respectively. Figure B.34 shows the density of the out-of-sample predictions for each model, as well as actual values. Generally, the models slightly over-predict at low values, more so for glyphosate than AMPA.

Generating predictions We use the model to predict county-month-level glyphosate and AMPA concentrations. We do this by making predictions for every watershed for each month between January of 1992 and December of 2017. We then take the weighted average of the predictions, where the weights are the proportion of the county’s population that lives in the watershed. Our population estimates come from SEDAC’s 2010 population grid (CIESIN, 2017). This grid estimates the population for one square kilometer pixels across the United States. We add the population counts for pixels within each watershed and then divide by the total population count for cells within the county to obtain the population weights. Figure B.35 shows predicted AMPA in July of 2004 from the LASSO model from each watershed touching Washington County on the right and the population weight for those watersheds on the left. Figure B.36 shows predicted

AMPA in water for each county in July of 2004. We can then link the county-month-level predictions of glyphosate and AMPA to the birth certificate data.

B.14.2 Results: Effect from upstream glyphosate in water.

Before using the machine learning predictions of glyphosate and AMPA in water, we first regress perinatal health outcomes on aggregate suitability over distance bins upstream or downstream from the mother's county of residence. These are of the form,

$$\text{Health}_{ijt} = \sum_{\tau \neq 1995} \gamma_{\tau}^l \text{GM}_j^l \times \mathbb{1}(t = \tau) + \sum_{\tau \neq 1995} \sum_d \gamma_{\tau d}^u \text{GM}_{jd}^u \times \mathbb{1}(t = \tau) + \Gamma X_{ijt} + \alpha_j + \lambda_t + \varepsilon_{ijt},$$

where GM_{jd}^u is the average GM suitability percentile in distance bin d upstream (or downstream) from county j . We now use GM_j^l to denote local GM suitability.

Figure B.37 displays event study plots illustrating the effect of max GM attainable yield in upstream watersheds on birthweight, categorized into 50-kilometer distance bins. These results suggest that having land more suitable for GM crops *upstream* of a county does not lead to a change in birthweight after the release of GM seeds in 1996.

As Dias, Rocha, and Soares Dias et al. (2023) emphasize, the potential effects of upstream glyphosate spraying would be strongest in places where there is more runoff from farms. We estimate the event study allowing for heterogeneity by high-soil-erodibility and high-precipitation, two factors that could increase runoff of glyphosate into surface water. Figure B.38 shows the results for both high- and low-erodibility and precipitation. Neither demonstrate a consistent effect on birthweight.

Finally, we estimate the effect of predicted glyphosate and AMPA in water. We do this by running regressions of the form,

$$\text{Health}_{ijt} = \beta^l \widehat{GLY}_{jt}^l + \beta^u \widehat{GLY}_{jt}^u + \Gamma X_{ijt} + \alpha_j + \lambda_t + \varepsilon_{ijt}, \quad (\text{B.1})$$

where $\widehat{GLY}_{jt}^l$ represents local glyphosate exposure, predicted from the first stage (3.2). $\widehat{GLY}_{jt}^u$ denotes predicted exposure to glyphosate or AMPA from glyphosate applied upstream of county j in year t , where we generate predictions from the machine learning model described above. These predictions are plausibly exogenous, as the models are trained only on exogenous data. Table B.9 shows the results of regressing these predictions of glyphosate or AMPA in water on birthweight. All four estimates, coming from either a LASSO or random forest model predicting either AMPA or glyphosate concentrations demonstrate a null effect of glyphosate or AMPA on birthweight.

We approach these findings cautiously; however, they suggest the absence of substantial downstream health spillovers resulting from glyphosate runoff. The lack of effect may be reasonably expected in the US relative to Brazil, as drinking water treatment in the US is more robust than that in Brazil (United Nations, 2024). However, we cannot definitively exclude water exposure as a potential mechanism driving the local results. glyphosate runoff into the water could be causing issues within a county but not downstream of a county if the chemicals degrade quickly enough. Additionally, given the inherent measurement error in this process and the absence of a more refined chemical transport model, we refrain from making definitive claims about the existence of downstream spillovers from glyphosate use.

Variable	Rural counties				Non-rural counties	
	High GM yield		Low GM yield		Mean	Std. Dev.
	Mean	Std. Dev.	Mean	Std. Dev.		
<i>Infant characteristics</i>						
Female infant	0.488	0.500	0.488	0.500	0.488	0.500
Birthweight (g)	3,327.0	610.1	3,360.1	583.8	3,333.5	604.5
Gestation (wk)	39.050	2.788	39.180	2.565	39.015	2.660
Preterm	0.185	0.388	0.166	0.373	0.178	0.383
C-section	0.243	0.429	0.223	0.416	0.217	0.412
Health index	0.029	0.251	0.031	0.244	0.021	0.242
Low birthweight	0.075	0.263	0.063	0.243	0.072	0.258
Very low birthweight	0.013	0.115	0.010	0.099	0.013	0.114
<i>Maternal characteristics</i>						
Black	0.199	0.400	0.024	0.152	0.172	0.378
Non-White	0.217	0.412	0.074	0.262	0.217	0.412
Hispanic	0.030	0.171	0.103	0.304	0.179	0.383
Married	0.683	0.465	0.746	0.435	0.692	0.462
High-school grad.	0.749	0.433	0.763	0.425	0.769	0.421
College grad.	0.115	0.319	0.131	0.337	0.207	0.405
<i>Local agriculture</i>						
Glyphosate (kg/km ²)	0.003	0.004	0.001	0.002	0.002	0.004
Max. GM crop suit. pctl.	0.759	0.144	0.248	0.138	0.494	0.285
Corn acres/km ²	24.303	31.789	8.586	20.261	7.774	17.128
Soy acres/km ²	25.314	30.369	4.019	12.280	6.105	14.826
Cotton acres/km ²	3.872	11.910	1.258	8.937	1.199	6.268
Other acres/km ²	15.015	22.160	11.386	23.691	6.156	13.409
<i>Local economy</i>						
Unemployment rate	0.070	0.026	0.079	0.041	0.066	0.027
Pct. farm employment	0.078	0.057	0.077	0.071	0.014	0.027
Farm employment per capita	0.036	0.027	0.037	0.036	0.006	0.011
Total population	43,742	28,666	41,164	31,005	1,347,402	2,122,955
Farm income per capita	0.491	0.592	0.600	1.236	0.104	0.257
Non-farm income per capita	15.843	2.580	15.884	3.313	22.165	5.312
Employment rate	0.480	0.092	0.486	0.117	0.559	0.167
<i>Totals</i>						
Number of births	1,981,619		1,734,005		20,031,928	
Number of counties	898		1,125		1,085	

Table B.1. Summary statistics for high- and low-yield rural counties and urban counties.

Means and standard deviations are calculated at the birth level for the county group in the years 1992–1995. GM yield grouping is based upon being above or below the 50th percentile of maximum attainable yield for GM crops. Rural/non-rural split uses USDA rural-urban continuum codes from 2003.

Table B.2. Correlation between GAEZ suitability measures and pre-period acreage for GM crops.

Dep Var: Model:	Acreage Percentiles				Yield Percentiles			
	GM (1)	Corn (2)	Soy (3)	Cotton (4)	GM (5)	Corn (6)	Soy (7)	Cotton (8)
Constant	0.227 (0.009)	0.240 (0.009)	0.141 (0.007)	0.263 (0.008)	0.214 (0.009)	0.326 (0.010)	0.178 (0.008)	0.264 (0.008)
GM GAEZ Yield Percentile	0.546 (0.015)				0.571 (0.015)			
Corn GAEZ Yield Percentile		0.521 (0.015)				0.348 (0.017)		
Soy GAEZ Yield Percentile			0.719 (0.012)				0.643 (0.014)	
Cotton GAEZ Yield Percentile				0.474 (0.013)				0.472 (0.013)
<i>Fit statistics</i>								
Observations	3,109	3,109	3,109	3,109	3,109	3,109	3,109	3,109
R ²	0.30	0.27	0.52	0.29	0.33	0.12	0.42	0.28

IID standard-errors in parentheses

We first calculate the county-level 1990 to 1995 average planted acreage and yield for each of corn, soy, cotton, and the aggregate of all three for GM. We divide the acreage values by the total size of the county. We then convert the acreage share and yield values into a percentile relative to all counties in the continental US. The GAEZ yield percentiles are calculated as described in *Methods*.






Figure B.1. GM crop suitability and increases in glyphosate for rural counties.

(a) Rural counties using 2003 Rural-Urban Continuum codes from the US Department of Agriculture (USDA) to classify counties as rural. A rural county is any non-metro county, where the USDA defines a metro county as, “broad labor-market areas that include central counties with one or more urban areas with populations of 50,000 or more people. **(b)** Percentile of attainable yield for GM crops equals the difference in attainable yield between high- and low-input scenarios from FAO GAEZ FAO and IIASA (2021) for corn, soy, and cotton. We rescale each crop to be a national percentile, take the maximum over the three crops, and finally scaling again to be a national percentile. Here we filter to only rural counties. **(c)** Change in glyphosate censored at the 1st and 99th percentiles and then filtered to only rural counties.



(a) Time series of GM crop adoption.



(b) Spatial variation in GM crop adoption.

Figure B.2. GM seeds were rapidly adopted after their 1996 release.

(a) Shows the percent of crops with any GM technology by year. The bold line is the entire United States and the grey lines are specific states. **(b)** Shows spatial variation across states in GM adoption rates in 2000, 2005, and 2010. Data from the USDA (USDA, 2024).




Figure B.3. Interpolation of fertilizer data.

This figure shows the fitted spline values and raw data for seven example counties. The dots are raw data and lines are the fitted spline functions. We fit a separate spline for each county and fertilizer type, using all data available—a value every five years between 1950 and 2017 from USGS Falcone (2021). We then generate annual-county-fertilizer type predictions to use with the rest of our data.




Figure B.4. Counties with high suitability for GM crops increased glyphosate intensity and reduced non-glyphosate pesticides with the introduction of glyphosate-resistant seeds.

Each event study comes from separate regressions where the given pesticide is regressed on local GM max attainable yield percentile interacted with year dummies with year and county fixed effects. All coefficients are scaled by the standard deviation of their respective variables. *Herbicide* and *Insecticide* each aggregate all other herbicides and insecticides not individually analyzed. Results from rural US counties. Standard errors are clustered by state and year. A unit of observation is county by year; regressions are weighted by total number of births.




Figure B.5. Perinatal health declined in GM-crop suitable counties after the introduction of glyphosate-resistant seeds

The subfigures extend Figure 13b to additional health outcomes—i.e. the estimated effect of local GM max attainable yield percentile on perinatal health outcomes relative to 1995. All regressions include county and year by month fixed effects and cluster errors by state and year. All regressions also control for family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.

	BW	LBW	VLBW	Gestation	Preterm	C-section	Health Index
Panel A: Policy effect							
GLY/km ²	-843.5 (329.8)	0.167 (0.079)	0.047 (0.017)	-6.13 (1.73)	0.606 (0.223)	0.324 (0.208)	-0.398 (0.119)
<i>Controls</i> (No additional controls)							
Panel B: GLY effect							
GLY/km ²	-1,280.7 (496.5)	0.279 (0.103)	0.087 (0.029)	-9.14 (2.47)	0.920 (0.309)	0.375 (0.374)	-0.549 (0.195)
<i>Controls</i>							
Pesticides	Y	Y	Y	Y	Y	Y	Y
Fertilizers	Y	Y	Y	Y	Y	Y	Y
Employment	Y	Y	Y	Y	Y	Y	Y
Income	Y	Y	Y	Y	Y	Y	Y
Age Shares	Y	Y	Y	Y	Y	Y	Y
Race Shares	Y	Y	Y	Y	Y	Y	Y
Population	Y	Y	Y	Y	Y	Y	Y
Fixed-effects (Both panels)							
Family Demog.	Y	Y	Y	Y	Y	Y	Y
County	Y	Y	Y	Y	Y	Y	Y
Yr × Mo	Y	Y	Y	Y	Y	Y	Y
Summaries (Both panels)							
N (millions)	10.73	10.73	10.73	10.71	10.71	9.510	10.71
2012 mean	3,271.1	0.081	0.014	38.6	0.207	0.278	

Table B.3. 2SLS estimates of the policy and direct GLY effects on perinatal health.

Each coefficient estimate (column-panel combination) provides results from a separate 2SLS regression. The six outcomes are birthweight (BW), the probabilities of low birthweight (LBW; BW < 2500g) and very low birthweight (VLBW; BW < 2500g), gestation length, and the probability of a preterm birth (gestation < 37 weeks). Both panels include family demographic, county, and year by month fixed effects. GLY effect (Panel B) additionally controls for other pesticides and unemployment. Sample restricted to births occurring in rural counties or to mothers residing in rural counties. Instruments are the attainable yield percentile for GM crops in each county interacted with year. Family demographic controls include mother's age, mother's race, mother's origin, mother's education, sex of child, total birth order, mother's residence status, and birth facility. Pesticide controls include alachlor, atrazine, cyanizine, fluazifop, metolachlor, metribuzin, and nicosulfuron. GLY/km² is kg/km². Standard errors in parentheses. We two-way cluster errors by year and state.

	Reduced-form DID results				Two-stage results	
	Raw reduced form		(1-2) scaled by gly.			
Outcome (unit)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Health outcomes						
Birthweight (g)	-29.42	-23.79	-29.2	-48.9	-19.6	-29.8
Gestation (wk)	-.148	-.119	-.146	-.244	-.143	-.213
LBW (%pt)	.507	.469	.502	.964	.399	.650
VLBW (%pt)	.111	.085	.110	.176	.110	.204
Preterm (%pt)	1.64	1.07	1.62	2.19	1.41	2.14
Health index	-.011	-.010	-.011	-.020	-.009	-.013
Panel B: Glyphosate						
Glyphosate (kg/km ²)	.024	.011				
Controls						
Ag. and econ.		Yes		Yes		Yes
Family demog.	Yes	Yes	Yes	Yes	Yes	Yes
County	Yes	Yes	Yes	Yes	Yes	Yes
Yr. × Mo.	Yes	Yes	Yes	Yes	Yes	Yes

Table B.4. Comparing difference-in-differences and 2SLS results

Sample restricts to births from mothers residing in a rural county. Instruments are the maximum attainable yield percentile for GM crops in each county (interacted with year in the 2SLS results in columns 5-6). All regressions include county and month-of-sample fixed effects and control for family and infant demographic controls (mother's age, mother's race, mother's origin, mother's education, sex of child, total birth order, mother's residence status, and birth facility). Glyphosate effect include additional *Ag. and econ.* controls (unemployment rate, employment rate, percent farm employment, and farm employment per capita, non-farm income per capita and farm income per capita, population, age shares, race/ethnicity shares, non-glyphosate pesticides, and fertilizer).




Figure B.6. The estimated effect of glyphosate on birthweight is robust to alternative specifications.

Coefficients are the estimated marginal effect of glyphosate (kg/km^2) on birthweight. Our main specifications are highlighted. All regressions include county of residence, county of occurrence, and family demographic fixed effects, standard errors are clustered by state and year. Pesticide controls include alachlor, atrazine, cyanazine, fluazifop, metolachlor, metribuzin, and nicosulfuron. Employment controls include unemployment rate, employment rate, farm employment per capita, and farm employment share. Income controls include farm and nonfarm income per capita. Age shares controls are share of population in seven decade wide bins from ages 0 to 70, with over 70 as the omitted category. Race share controls are proportion of the population white, Black, and Hispanic. Fertilizer controls are commercial nitrogen, commercial phosphorous, manure nitrogen, and manure phosphorous. Acre controls are corn, soy, and cotton acres, as well as an aggregate of all other crop acreage. Family demographic FEs include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. We vary the construction of GM attainable yield: "GM Max Percentile" is our main specification, "GM Avg Percentile" takes the average standardized attainable yield among corn, soy, and cotton (rather than the average) before re-scaling into a percentile, "GM Average, Split at Median" uses a binary high vs low attainable yield, where a county is high attainable yield if they are above the median attainable yield, "GM Max Top vs Bottom Quartile" is also binary, but only compares the top and bottom quartiles, omitting the middle group, and "1990-1995 GM Max Yield Percentile" is the percentile of observed yield in each county for corn, soy, and cotton between 1990 and 1995 using data from USDA NASS. "Eastern US" measures filter to counties east of the 100th meridian then calculate percentiles. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.7. Glyphosate effects for infants born to non-white mothers are larger for birthweight and for the probabilities of preterm birth, LBW, and VLBW.

Policy and Glyphosate effects for all outcomes at the mean level of glyphosate in 2012, estimated separately by mother's race. All regressions include county and year by month fixed effects, and control for family demographics. Standard errors are clustered by state and year. The Glyphosate Effect adds controls for other pesticides, employment, income, population, age and race shares, and fertilizers. The sample is restricted to births occurring in rural counties or to mothers residing in rural counties.

	(1)	(2)	(3)
Panel A: Birthweight			
Glyphosate/km ²	−843.5 (329.8)	−885.4 (324.8)	−855.0 (323.2)
Panel B: Gestation			
Glyphosate/km ²	−6.13 (1.73)	−5.69 (1.67)	−5.65 (1.66)
Panel C: Health index			
Glyphosate/km ²	−0.398 (0.119)	−0.361 (0.118)	−0.347 (0.116)
<i>Control sets</i>			
Age Shares		Yes	Yes
Race Shares			Yes
<i>Fixed-effects</i>			
Family demog.	Yes	Yes	Yes
County	Yes	Yes	Yes
Yr. × Mo.	Yes	Yes	Yes

Table B.5. Effect of glyphosate on birthweight, gestation, and health index:
Robustness to age and race share controls

Sample restricted to births from mothers residing in a rural county. Instruments are the maximum attainable yield percentile for GM crops in each county interacted with year. Family demographic controls include mother’s age, mother’s race, mother’s origin, mother’s education, sex of child, total birth order, mother’s residence status, and birth facility. Age shares include the share of population in each county in seven 10 year age bins from age 0 to 70. We omit the over 70 category. Race shares include the share of the population in each county that is black, share white, and share hispanic. Glyphosate/ km^2 is kg/km^2 .





Figure B.8. Predicted birthweights closely match actual birthweights across the predicted birthweight distribution.

At each predicted birthweight percentile (x-axis), we take the average actual birthweight and average predicted birthweight, which are both plotted in the y-axis. Sample includes births to mothers with rural residences from 1990 to 2013.


Outcome (<i>units</i>):	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Birthweight (<i>g</i>)		Percentile (<i>pp</i>)		Decile (<i>1–10</i>)		Quntile (<i>1–5</i>)	
Metric:	MAE	MAPE	MAE	MAPE	MAE	MAPE	MAE	MAPE
Panel A: <i>Actual predictions (across three samples)</i>								
<i>Sample</i>								
Full	406.7	15.7	26.1	290.7	2.58	79.0	1.25	59.1
1–99 pctl.	381.0	12.6	26.1	127.9	2.58	78.3	1.26	59.2
5–95 pctl.	329.0	10.3	26.1	91.7	2.59	73.0	1.27	58.8
Panel B: <i>Null models (on the full sample)</i>								
<i>Approach</i>								
Means	443.4	17.2	24.8	488.3	2.48	93.3	1.19	64.2
At random	643.6	23.4	33.3	324.5	3.30	106.1	1.60	77.0

Table B.6. ML prediction performance

Panel A describes the performance of the random-forest model across three samples—the full sample, births in percentiles 1–99, and births in percentiles 5–95. We evaluate the predictions on two metrics: mean absolute error (MAE: $|y - \hat{y}|$) and mean absolute percent error (MAPE: MAE/y). The model trained to predict birthweight (columns 1–2). We also evaluate the performance for predicting birthweight percentile (3–4), decile (5–6), and quintile (7–8). *Panel B* provides two *null* models that (1) predict the sample *means* or (2) predict *at random*. The table uses births to rural-residence mothers to match the main analyses throughout the paper.



(a) *First-stage* effect of GM suitability on glyphosate intensity by predicted birthweight quintile



(b) *Reduced-form* effect of GM suitability on birthweight by predicted birthweight quintile

Figure B.9. First-stage event study coefficients are similar across predicted BW quintiles, reduced form shows larger effects in lower quintiles.

(a) Estimated event-study coefficients for the effect of local GM max attainable yield percentile on glyphosate by year relative to 1995 by predicted birthweight quintile. Pesticide data only go back to 1992—there are no coefficients in 1990–1991. (b) Similar event study but with birthweight as outcome. Estimates from each predicted birthweight quintile come from separate regressions. All regressions include county, year by month, and family demographic fixed effects. Standard errors are clustered by state and year. Family demographics include mother’s age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father’s age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.10. Heterogeneity in Glyphosate Effect is consistent across various predicted birthweight bin sizes, greater disparities among birthweight outcomes.

Estimated Glyphosate Effect at mean of glyphosate/km² on various perinatal health outcomes instrumented with GM attainable yield interacted with year. All regressions include county, year by month, and family demographic fixed effects, and control for other pesticides, employment, income, population, age and race shares, and fertilizers. Standard errors are clustered by state and year. Family demographics include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.11. Limited evidence of heterogeneous marginal effects by sex within predicted BW quintile.

Estimated marginal effect of glyphosate/km² on various perinatal health outcomes instrumented with GM max attainable yield interacted with year. All regressions include county, year by month, and family demographic fixed effects and control for other pesticides, employment, income, population, age and race shares, and fertilizers. Standard errors are clustered by state and year. Family demographics include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.

	BW	LBW	VLBW	Gestation	Preterm	C-section	Health Index
Panel A: Policy effect							
GLY/km ²	24.6	-0.023	-0.003	-0.511	0.030	0.021	-0.047
	(75.9)	(0.022)	(0.006)	(0.373)	(0.050)	(0.068)	(0.029)
<i>Controls</i> (No additional controls)							
Panel B: GLY effect							
GLY/km ²	19.9	0.005	-0.002	-0.368	0.046	-0.011	-0.006
	(60.2)	(0.015)	(0.006)	(0.328)	(0.040)	(0.071)	(0.028)
<i>Controls</i>							
Pesticides	Y	Y	Y	Y	Y	Y	Y
Fertilizers	Y	Y	Y	Y	Y	Y	Y
Employment	Y	Y	Y	Y	Y	Y	Y
Income	Y	Y	Y	Y	Y	Y	Y
Age Shares	Y	Y	Y	Y	Y	Y	Y
Race Shares	Y	Y	Y	Y	Y	Y	Y
Population	Y	Y	Y	Y	Y	Y	Y
Fixed-effects (Both panels)							
Family Demog	Y	Y	Y	Y	Y	Y	Y
County	Y	Y	Y	Y	Y	Y	Y
Yr × Mo	Y	Y	Y	Y	Y	Y	Y
Summaries (Both panels)							
N obs. (millions)	10.73	10.73	10.73	10.71	10.71	9.51	10.73

Table B.7. OLS estimates of the policy and direct GLY effects on perinatal health.

Each coefficient estimate (column-panel combination) provides results from a separate OLS regression. Both panels include family demographic, county, and year by month fixed effects. GLY effect (Panel B) additionally controls for other pesticides, employment, income, population, age and race shares, and fertilizers. Sample restricted to births occurring in rural countries or from mothers residing in rural counties. Family demographic controls include mother's age, mother's race, mother's origin, mother's education, sex of child, total birth order, mother's residence status, and birth facility. Pesticide controls include alachlor, atrazine, cyanazine, fluazifop, metolachlor, metribuzin, and nicosulfuron. GLY/km² is kg/km². Standard errors in parentheses. We two-way cluster errors by year and state.

	BW	LBW	VLBW	Gestation	Preterm	C-section	Health Index
Panel A: Policy effect							
GLY/km ²	-545.3 (206.0)	0.112 (0.052)	0.032 (0.012)	-4.06 (1.00)	0.418 (0.142)	0.174 (0.149)	-0.247 (0.069)
<i>Controls</i> (No additional controls)							
Panel B: GLY effect							
GLY/km ²	-933.3 (329.1)	0.212 (0.082)	0.062 (0.024)	-6.21 (1.61)	0.664 (0.213)	0.100 (0.268)	-0.353 (0.119)
<i>Controls</i>							
Pesticides	Y	Y	Y	Y	Y	Y	Y
Fertilizers	Y	Y	Y	Y	Y	Y	Y
Employment	Y	Y	Y	Y	Y	Y	Y
Income	Y	Y	Y	Y	Y	Y	Y
Age Shares	Y	Y	Y	Y	Y	Y	Y
Race Shares	Y	Y	Y	Y	Y	Y	Y
Population	Y	Y	Y	Y	Y	Y	Y
Fixed-effects (Both panels)							
Family Demog	Y	Y	Y	Y	Y	Y	Y
County	Y	Y	Y	Y	Y	Y	Y
Yr × Mo	Y	Y	Y	Y	Y	Y	Y
Summaries (Both panels)							
N obs. (millions)	10.73	10.73	10.73	10.71	10.71	9.51	10.73

Table B.8. Effect of GLY on perinatal health estimated with 2SLS shift-share instrument.

Each coefficient estimate (column-panel combination) provides results from a separate 2SLS regression. Instruments are the measure of suitability in each county interacted with national glyphosate usage, excluding glyphosate from counties within 100km or upstream. The six outcomes are birthweight (BW), the probabilities of low birthweight (LBW; BW < 2500g) and very low birthweight (VLBW; BW < 2500g), gestation length, and the probability of a preterm birth (gestation < 37 weeks). Both panels include family demographic, county, and year by month fixed effects. GLY effect (Panel B) additionally controls for other pesticides and unemployment. Sample restricted to births occurring in rural counties or to mothers residing in rural counties. Instruments are the attainable yield percentile for GM crops in each county interacted with year. Family demographic controls include mother's age, mother's race, mother's origin, mother's education, sex of child, total birth order, mother's residence status, and birth facility. Pesticide controls include alachlor, atrazine, cyanazine, fluazifop, metolachlor, metribuzin, and nicosulfuron. GLY/km² is kg/km². Standard errors in parentheses. We two-way cluster errors by year and state.






Figure B.12. Robustness of birthweight effect to alternative economic controls.

Estimated effect of local GM max attainable yield percentile on birthweight relative to 1995. Employment controls include the unemployment rate, employment rate, farm employment per capita, and farm employment share. Income controls include farm and nonfarm income per capita. Age share controls are the share of population in seven decade wide age bins, with the over 70-population as the reference group. Race shares are the proportion of the population white, Hispanic, and Black. All regressions include fixed effects for family demographics, county, and year-month and standard errors are clustered by state and year. Family demographics include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




(a) First-stage effect of local GM attainable yield on glyphosate




(b) Reduced-form effect of local GM attainable yield on birthweight

Figure B.13. Robustness of birthweight effect to alternative farm controls.

Estimated effect of local GM max attainable yield percentile on birthweights relative to 1995. Fertilizer controls include commercial nitrogen, commercial phosphorous, manure nitrogen, and manure phosphorous. Acre controls include soy, corn, and cotton acres, as well as an aggregate of total acreage from other crops. “Acres and all others” specification uses all of the economic controls—pesticides, fertilizers, acres, employment, income, age and race shares, and population. All regressions include fixed effects for family demographics, county, and year-month and standard errors are clustered by state and year. Family demographics include mother’s age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father’s age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




(a) First-stage effect of various instruments on glyphosate/ km^2




(b) Reduced-form effect of various instruments on birthweight

Figure B.14. Robustness of birthweight effect to alternative instruments.

We vary the construction of our GM suitability measure: “GM Max Percentile” is our main specification, “GM Avg Percentile” takes the average standardized attainable yield among corn, soy, and cotton (rather than the max) before re-scaling into a percentile, “1990-1995 GM Max Yield Percentile” is constructed using pre-period realized yields for corn, soy, and cotton, “GM Average, Split at Median” uses a binary high vs low attainable yield, where a county is high attainable yield if they are above the median attainable yield, and “GM Max Top vs Bottom Quartile” is another binary treatment definition that compares just the top and bottom quartiles in GM Max Percentile, omitting the middle group. All regressions include family demographics, county, and year by month fixed effects and standard errors are clustered by state and year. The regressions control for fertilizers, other pesticides, employment, income, age and race shares, and population. Family demographic fixed effects include mother’s age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father’s age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




(a) First-stage effect of various instruments on glyphosate kg/km^2




(b) Reduced-form effect of various instruments on birthweight

Figure B.15. Robustness of birthweight effect to alternative fixed effects.

We vary the fixed effects included. All regressions include family demographics, county, and year by month fixed effects and standard errors are clustered by state and year. The regressions control for fertilizers, other pesticides, employment, income, age and race shares, and population. Family demographic fixed effects include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.



(a) First-stage effect of various instruments on glyphosate kg/km^2



(b) Reduced-form effect of various instruments on birthweight

Figure B.16. Robustness of birthweight effect to alternative fixed effects—within state.

We vary the fixed effects included. All regressions include family demographics, county, and year by month fixed effects and standard errors are clustered by state and year. The regressions control for fertilizers, other pesticides, employment, income, age and race shares, and population. Family demographic fixed effects include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.






Figure B.17. Heterogeneity in Birthweight Effect by Geographic Subsets.

Estimated effect of local max GM attainable yield percentile on birthweight relative to 1995. The geographic subsets are primarily defined using census regions (Midwest, Northeast, or South). Fig B.14 shows results with just the eastern US. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for other pesticides, employment, income, population, age and race shares, fertilizers and family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.



(a) First-stage effect on glyphosate kg/km², rural and non-rural counties.



(b) Reduced-form effect on birthweight, rural and non-rural counties.

Figure B.18. Birthweight event studies by rural and non-rural counties.

Estimated effect of local max GM attainable yield percentile on birthweight relative to 1995 for births to mothers residing and occurring in rural and non-rural counties. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for other pesticides, employment, income, population, age and race shares, and fertilizers, and family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race.




Figure B.19. The estimated effect of glyphosate on gestation is robust to alternative specifications.

Coefficients are the estimated marginal effect of glyphosate (kg/km^2) on gestation. Our main specifications are highlighted. All regressions include county of residence, county of occurrence, and family demographic fixed effects, standard errors are clustered by state and year. Pesticide controls include alachlor, atrazine, cyanazine, fluazifop, metolachlor, metribuzin, and nicosulfuron. Employment controls include unemployment rate, employment rate, farm employment per capita, and farm employment share. Income controls include farm and nonfarm income per capita. Age shares controls are share of population in seven decade wide bins from ages 0 to 70, with over 70 as the omitted category. Race share controls are proportion of the population white, Black, and Hispanic. Fertilizer controls are commercial nitrogen, commercial phosphorous, manure nitrogen, and manure phosphorous. Acre controls are corn, soy, and cotton acres, as well as an aggregate of all other crop acreage. Family demographic FEs include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. We vary the construction of GM attainable yield: "GM Max Percentile" is our main specification, "GM Avg Percentile" takes the average standardized attainable yield among corn, soy, and cotton (rather than the average) before re-scaling into a percentile, "GM Average, Split at Median" uses a binary high vs low attainable yield, where a county is high attainable yield if they are above the median attainable yield, "GM Max Top vs Bottom Quartile" is also binary, but only compares the top and bottom quartiles, omitting the middle group, and "1990-1995 GM Max Yield Percentile" is the percentile of observed yield in each county for corn, soy, and cotton between 1990 and 1995 using data from USDA NASS. "Eastern US" measures filter to counties east of the 100th meridian then calculate percentiles. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.20. The estimated effect of glyphosate on the health index is robust to alternative specifications.

Coefficients are the estimated marginal effect of glyphosate (kg/km^2) on the health index. Our main specifications are highlighted. All regressions include county of residence, county of occurrence, and family demographic fixed effects, standard errors are clustered by state and year. Pesticide controls include alachlor, atrazine, cyanazine, fluazifop, metolachlor, metribuzin, and nicosulfuron. Employment controls include unemployment rate, employment rate, farm employment per capita, and farm employment share. Income controls include farm and nonfarm income per capita. Age shares controls are share of population in seven decade wide bins from ages 0 to 70, with over 70 as the omitted category. Race share controls are proportion of the population white, Black, and Hispanic. Fertilizer controls are commercial nitrogen, commercial phosphorous, manure nitrogen, and manure phosphorous. Acre controls are corn, soy, and cotton acres, as well as an aggregate of all other crop acreage. Family demographic FEs include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. We vary the construction of GM attainable yield: "GM Max Percentile" is our main specification, "GM Avg Percentile" takes the average standardized attainable yield among corn, soy, and cotton (rather than the average) before re-scaling into a percentile, "GM Average, Split at Median" uses a binary high vs low attainable yield, where a county is high attainable yield if they are above the median attainable yield, "GM Max Top vs Bottom Quartile" is also binary, but only compares the top and bottom quartiles, omitting the middle group, and "1990-1995 GM Max Yield Percentile" is the percentile of observed yield in each county for corn, soy, and cotton between 1990 and 1995 using data from USDA NASS. "Eastern US" measures filter to counties east of the 100th meridian then calculate percentiles. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.21. Robustness to spatial subsets, all outcomes.

Estimated effect of glyphosate/ km^2 on various perinatal health outcomes instrumented with GM max attainable yield interacted with year. All regressions include county, year by month, and family demographic fixed effects and control for other pesticides, employment, income, population, age and race shares, and fertilizers. Standard errors are clustered by state and year. Family demographics include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race.




Figure B.22. Demographic event studies.

Estimated effect of local GM max attainable yield percentile on various demographics outcomes relative to 1995. All regressions include county and year by month fixed effects and control for family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race (controls exclude the outcome). Standard errors are clustered by state and year. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.23. Reduced form heterogeneity by mother's race.

Estimated effect of local GM max attainable yield percentile on various perinatal health outcomes relative to 1995. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for other pesticides, employment, income, population, age and race shares, fertilizers, and family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.24. Reduced form heterogeneity by mother's education.

Estimated effect of local GM max attainable yield percentile on various perinatal health outcomes relative to 1995. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for other pesticides, employment, income, population, age and race shares, fertilizers, and family demographics, including mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race. Sample restricted to births occurring in rural counties or to mothers residing in rural counties.




Figure B.25. Heterogeneity in effect by different month of birth, all outcomes.

Estimated effect of glyphosate/ km^2 on various perinatal health outcomes instrumented with GM max attainable yield interacted with year. All regressions include county, year, and month fixed effects and standard errors are clustered by state and year. Family demographics include mother's age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father's age and race.






Figure B.26. Effect of local Max GM attainable yield on crop acreage.


Standard errors are clustered by state and year. All regressions include county and year fixed effects with no other controls and are weighted by the number of births. Sample restricted to rural counties.



(a) Soy yield




(b) Corn yield




(c) Cotton yield

Figure B.27. Effect of local GM max attainable yield on crop yield.


Soy and corn yield is measured in bushels/acre, while cotton is measured in lbs/acre. Standard errors are clustered by state and year




(a) Effect of GM suitability on non-farm income.



(b) Effect of GM suitability on farm income



(c) Effect of GM suitability on unemployment rate



(d) Effect of GM suitability on farm employment

Figure B.28. Coefficients from an event study regression of various socioeconomic variables on GM suitability for rural counties.

Standard errors are clustered by state and year. Regressions are weighted by total births. Sample restricted to rural counties.




Figure B.29. Predicted glyphosate disaggregated into watersheds in 2004.

These predictions come from our first stage model regressing glyphosate on local GM attainable yield percentile with county and year fixed effects. We disaggregate from county into watersheds using the portion of the county's area that is covered by each watershed. We generate predictions for each year, but only show 2004 to accompany the exposition.




Figure B.30. Spatial variation in water ML predictors.

Each map depicts a watershed-level average of the given variable. See text for details.






Figure B.31. Capturing upstream and downstream watersheds.


For an example watershed in Illinois (highlighted in red), we show all of the watersheds upstream and all of the watersheds downstream. We calculate distance upstream and downstream using the distance between the centroids of watersheds along the path, then categorize these into 50-kilometer distance bins.




(a) CV Results for AMPA Random Forest Model.



(b) CV Results for glyphosate Random Forest Model.



(c) CV Results for AMPA LASSO Model.



(d) CV Results for glyphosate LASSO Model.

Figure B.32. Cross Validation Results.




Figure B.33. Out-of-sample prediction performance for LASSO and Random Forest models.

Predictions are made on the 757 held-out observations in order to assess the model fit for LASSO and random forest models. Smooth lines are that of a generalized additive model.




Figure B.34. Density of out-of-sample predictions relative to the actual values.




Figure B.35. Aggregating watersheds to counties.

For the same watershed as in Fig B.31 with the red border, on the left, we have population weights for Washington County (black outline). On the right, we have our predicted AMPA in July, 2004 using the LASSO model in each watershed touching Washington County. Thus, to generate county-month level predicted AMPA, we take the weighted average of predictions (left), where the weights come from the population in each watershed (right).




Figure B.36. Predicted county-level AMPA in July of 2004.

This is a map for one month (July), in one year (2004), using one of four predictive models (LASSO predicting AMPA). We generate county level predictions like this for all months and years between 1992 and 2017 with LASSO and random forest models predicting glyphosate and AMPA.

Dep Var	BW			
Model:	(1)	(2)	(3)	(4)
Predicted AMPA (LASSO)	16.1 (11.2)			
Predicted AMPA (RF)		-3.14 (8.83)		
Predicted GLY (LASSO)			-6.01 (16.7)	
Predicted GLY (RF)				-5.94 (6.69)
Local attainable yield	Yes	Yes	Yes	Yes
Local pesticides	Yes	Yes	Yes	Yes
Unemployment	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Family Demog	Yes	Yes	Yes	Yes
Yr x Mo + Cnty	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
N (millions)	7.910	7.910	7.910	7.910
<i>Clustered (Year & State) standard-errors in parentheses</i>				

Table B.9. Effect of predicted GLY or AMPA in water on birthweight.




Figure B.37. Effect of upstream glyphosate by distance bin.

Estimated effect of upstream GM max attainable yield percentile on various perinatal health outcomes relative to 1995. Bin labels represent the lower bound distance between the county and the upstream watershed, thus “50” is an aggregate of watersheds 50 to 100km upstream of a county. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for local attainable yield interacted with year, other pesticides, employment, income, population, age and race shares, fertilizers, and family demographics, including mother’s age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father’s age and race. Sample restricted to births from mothers with rural residence.




Figure B.38. Effect of upstream glyphosate by distance bin by high and low soil erodibility and precipitation.

Estimated effect of upstream GM max attainable yield percentile on various perinatal health outcomes relative to 1995. Bin labels represent the lower bound distance between the county and the upstream watershed, thus “50” is an aggregate of watersheds 50 to 100km upstream of a county. All regressions include county and year by month fixed effects and standard errors are clustered by state and year. All regressions also control for local attainable yield interacted with year, other pesticides, employment, income, population, age and race shares, fertilizers, and family demographics, including mother’s age, race, education, marital status, birth facility, resident status, previous births, sex of infant, and father’s age and race. Sample restricted to births from mothers with rural residence.

APPENDIX C

APPENDIX FOR OPTIMAL SUBSIDIES FOR RESIDENTIAL SOLAR

C.1 Data Appendix

C.1.1 Deep Solar. We obtain data on solar panel installation from the Deepsolar database, which is created by applying a novel semi-supervised deep-learning framework to satellite imagery from Google Static Maps from the year 2016 (Yu et al., 2018). The Deepsolar model predicts the number and size of solar panel installations across the contiguous United States. We use these tract-level data on the number and size of residential solar systems to give us our empirical analogs of M_t and installation size N_i .

Deepsolar also estimates the daily solar radiation in each census tract, measured in kWh per square meter per day, which we show in Figure C.1. For any missing tracts in the Deepsolar data, we impute daily solar radiation by taking the simple mean of any bordering tract with non-missing values.

C.1.2 Google Project Sunroof. For data on solar irradiance, A_i , and number of available spaces for panels $\bar{N}_i$, we utilize tract-level data from Google Project Sunroof (GPS). GPS begins with satellite imagery from Google




Figure C.1. Daily solar radiation ($kWh/m^2/day$) by census tract from Deepsolar.

Maps. It then applies a deep-learning algorithm to create 3D models of rooftops. These 3D models allow GPS to estimate the amount of sunlight a given rooftop receives over the course of the year, taking into account changes in the position of the sun over the course of the day and year. These 3D models are used to calculate the amount of available space for solar panels.

We assume that all households within a given tract have access to the same solar irradiance, which we measure as total solar energy generation potential for the average panel in a given tract. For number of potential panels $\bar{N}_i$, the GPS data provide the number of buildings in each tract with differing amounts of space available for solar panel installations. This effectively gives us the full distribution of $\bar{N}_i$ for households within a given tract.

One potential issue with the GPS data is that it might also capture potential space for solar panels that is not suitable for residential solar (for example not being part of someone’s house). To deal with this, we limit potential solar sites in Google Project Sunroof to those with space available to 42 MW of solar panels, corresponding to the 99.9th percentile of the largest solar panel in the Tracking the Sun data described in Online Appendix C.1.3. The results are not sensitive to this censoring.

C.1.3 Tracking the Sun. Tracking the Sun is an aggregation of solar system installation data created by the Lawrence Berkeley National Lab. The Lawrence Berkeley National Lab collects these data from existing public databases and directly from state agencies, utilities, and other organizations. The result is 2.5 million solar installations from the last two decades, with installation price, system size, and subsidies geographically identified at the zip code level, along with other information about the installed solar system. The installations cover nearly 80

percent of all installed solar systems in the U.S. but include only 25 states. Some of these states do not include price data for any installations. Across all states, about a quarter of observations for residential solar system installations are missing price data.

We use the Tracking the Sun data to estimate prices for solar systems, using total cost and number of panels installed to estimate a fixed cost of installation and variable, per-panel cost. Since many states have no data, we assume pricing functions are common within each census region. We filter the Tracking the Sun data to include residential installations between 2014 and 2018 that are not missing price or the total number of panels, which leaves us with over 720 thousand observations. Additionally, we censor installation costs at the 0.5th and 99.5th percentiles and convert them into 2014 dollars.

C.1.4 System Advisor Model. While we obtain annual electricity generation for solar panels from Google Project Sunroof, those data do not include any information on how that production varies by hour within a year. Thus, we use the System Advisor Model (SAM) from the National Renewable Energy Laboratory to estimate hourly electricity profiles for each county (Blair et al., 2013). SAM is an open-source program that estimates the performance of solar systems and other renewable power systems. We follow the methodology in Sexton et al. (2021), calculating electricity generation for a system with typical parameters where tilt matches latitude and panels point south. The only difference is that we estimate generation for systems at the centroid of each county. We use weather data from Sengupta et al. (2018) for a county’s typical meteorological year.

The model’s output is the hourly production of a solar system over the course of a year in each county. We create hourly profiles by dividing the hourly




Figure C.2. Hourly electricity generation for a standard solar panel for six example counties.

generation by each county's total annual generation. Figure C.2 for hourly production for examples of the results for six counties.

C.1.5 State Electricity Prices. Figure C.3 presents the state-level electricity prices we use in our empirical analysis. California and states in the Northeast have the highest electricity prices at over 15 cents per kWh. Most of the country has prices between 8 and 10 cents per kWh.

C.1.6 Subsidies. We calculate s_j^{kWh} as the sum of per-kWh rebates and the average price of Solar Renewable Energy Certificates. In some states (e.g. Massachusetts and New Jersey), households can only sell Solar Renewable Energy Certificates for a certain number of years after installation. For these states, we only calculate the value of Solar Renewable Energy Certificates for years in which households are permitted to sell the credits. We calculate s_j^{Cost} as the sum of federal investment tax credits, state investment tax credits, sales tax exemptions,




Figure C.3. State electricity prices (\$/kWh)




Figure C.4. Expected subsidy for a 15-panel system by subsidy type in 2014 dollars.

and the net present value of property tax exemptions. We translate per-KW rebates to s_j^{Panel} by assuming a constant 0.25 KW per panel. Maryland has a fixed rebate of \$1000 per system. We translate this into a per-panel subsidy by dividing this amount by the average number of panels in an installation (15). Many states place a cap on the maximum amount of a type of subsidy a household can receive. We enforce these state-level maxima in estimation. Figure C.4 shows the state level expected subsidies for a 15-panel system for each type of subsidy: per-panel, cost, and kWh. Most of the current subsidies take the form of cost-based subsidies, while few states offer kWh and panel-based subsidies.

C.1.7 Power Plants. The following describes how we use the Open Grid Emissions (OGE) data.¹ These data have several advantages over the commonly used raw electricity data from the EIA and EPA, which we describe below.

Plant level The hourly, plant-level data from OGE give net electricity generation and emissions of SO₂, NO_x, CO₂, and CO₂ equivalent. OGE derives these data primarily from the EPA CAMD, which reports hourly gross electricity generation and emissions at the unit level, where units typically correspond to generators connected to a single emissions stack. OGE adjusts gross generation to account for electricity losses before entering the grid and aggregates these units to the facility level, which we refer to as power plants. Additionally, OGE removes the portion of emissions from fuel burned for heat for combined heat and power plants. OGE’s static plant attributes table has a latitude and longitude for each plant, allowing us to match each power plant to a county.

We also collect each plant’s nameplate capacity and stack height. Nameplate capacities are from EIA Form 860 and represent the maximum amount of electricity that a generating unit is rated to produce. We sum the nameplate capacities of generators in a plant to calculate plant-level nameplate capacities. Thirty-seven plants are missing from the EIA 860 data, for which we use nameplate capacities from the EPA’s eGRID files from 2019 and 2020. We obtain stack heights from the EPA CAMD and set a plant’s stack height as the median stack height of units the within that plant. Thirty plants are missing from these data, for which we set the stack height equal to the median stack height of all plants of the same primary fuel

¹OGE is a relatively new dataset and under active development. We use v0.3.0 in this analysis.

category. We use a power plant’s stack height and location to assign the marginal damage coefficient in the AP3 model. We opt for CO₂e over CO₂ when calculating damages as it includes emissions of the more potent greenhouse gasses methane and nitrous oxide in addition to CO₂.²

The EPA CAMD hourly unit-level data only include fossil-fuel plants with greater than 25 MW of generating capacity, leaving a non-negligible portion of generation and emissions unreported. One of the main goals of OGE is to ensure complete coverage of the electricity generation sector. In essence, they combine the reported hourly plant-level data from the EPA CAMD with hourly balancing authority-fuel category level data from the EIA to calculate a ‘residual’ profile, the unreported production from small or non-fossil-fuel power plants.

There are 9,167 plants with hourly production and emissions in the 2019 OGE data. About a third of the plants do not have observations for every hour in 2019. We fill in any generation and emissions values between the first and the last hour a plant appears in the data with zeros. After removing 4,312 non-dispatchable plants (wind and solar), 193 plants with zero or negative reported net electricity generation, and 64 plants with no variation in net electricity generation, we have 4,598 power plants—giving us nearly 40 million plant-hour observations.

Emissions rates OGE’s hourly data does not include PM_{2.5} emissions, as the EPA CAMD and the EIA do not report PM_{2.5} emissions from power plants. As a part of the eGRID project, the EPA has collected annual PM_{2.5} emissions from the National Emissions Inventory (NEI) and matched those emissions to electricity-

²As detailed in the OGE documentation, they calculate CO₂e using the global warming potential of each GHG according to the IPCC’s 5th Assessment Report. They calculate methane and nitrous oxide emissions using a constant, fuel-specific emissions factor.

generating units to calculate an average PM2.5 emissions rate. We match these estimated annual rates to our power plants, taking the production-weighted average over units within a power plant. We use the fuel category median value for the power plants missing PM2.5 emissions rates. This imputation only applies to less than 10% of total electricity production. We censor PM 2.5 emissions rates at the 95th percentile for each fuel category. A few power plants have emissions rates orders of magnitude larger than is reasonable—these are small plants that do not directly report their hourly emissions or generation in the EPA CAMD data, and thus must be imputed by OGE.³ For these plants, we set emissions rates equal to the median for their fuel type.

Regions We follow Holland et al. (2016) in our definitions of regions for the electricity generation model. OGE assigns plants to the balancing authority in charge of dispatching the plant. We then assign balancing authorities to regions. There are six NERC regions in the contiguous US. Four of these (MRO, RFC, NPCC, and SERC) fall within the Eastern Interconnection, while the other two (WECC and TRE) are in the Western and Texas Interconnections, respectively.

Most BAs fall entirely within one NERC region, but some BAs have generating units in multiple NERC regions. For all BAs except MISO and PJM, we assign the BA to the NERC region with the most overlapping generating units between the BA and NERC region using the static plant attributes data from OGE. We assign the MISO BA to the MRO NERC region and the PJM BA to the RFC NERC region. Finally, we give California its own NERC region, consisting of five BAs: BANC, CISO, IID, LDWP, and TIDC. Figure C.5 shows a map of these

³These plants' EIA ID's are 2221, 2528, 2503, 50626, and 50931.




Figure C.5. Map of region service areas

Table C.1. Summary statistics on power plants by region

Region	Number of plants	Net Generation		Emissions (lb/MWh)		
		Total (TWH)	% Fossil Fuel	NOx	SO2	CO2e
CAL	642	1,623	58.9	0.66	0.07	663
MRO	1,168	8,116	79.7	0.96	1.03	1,343
NPCC	725	2,291	46.1	0.41	0.06	540
RFC	645	7,831	61.8	0.50	0.51	861
SERC	572	9,396	67.4	0.49	0.35	918
TRE	179	3,348	87.2	0.64	0.70	1,069
WECC	667	4,776	61.5	0.75	0.36	1,011
Total	4,598	37,382	68.2	0.64	0.53	990

regions. We used the eGRID power profiler to assign approximate service areas for each region.

Table C.1 shows summary statistics describing generation and average emissions in each region, highlighting the heterogeneity in average emissions between regions. This is largely driven by differences in the fuel mix between regions. Table C.2 shows a summary of generation and emissions by fuel category.

Excess Load We calculate the excess load (total demand minus production from nondispatchable generating units) within each region using OGE's power

Table C.2. Summary statistics on power plants by fuel category

Fuel	Number of plants	Net generation (TWH)	Emissions (lb/MWh)		
			NO _x	SO ₂	CO ₂ e
Biomass	577	542	5.54	1.12	1,569
Coal	258	9,457	1.44	1.95	2,221
Natural Gas	1,685	16,027	0.41	0.03	906
Nuclear	60	8,136	0.00	0.00	5
Petroleum	517	25	3.98	4.98	1,907
Geothermal	60	155	0.00	0.35	137
Hydro	1,332	2,854	0.00	0.00	0
Other	43	49	0.08	0.00	105
Waste	66	135	5.83	1.32	3,722
Total	4,598	37,382	0.64	0.53	990

sector-level data. These data give hourly net generation by fuel category for each balancing authority. We perform minor data cleaning to ensure that misreporting in the underlying data does not impact our estimates. We consider values above 1.5 times the 99th percentile for that balancing authority and fuel category to be outliers. We replace outliers with the value from the previous hour as long as the previous hour's value is not also an outlier. If the previous hour is an outlier, then we use the value from the same hour in the previous day as long as that is not an outlier. If the previous hour and the previous day are outlier values, we censor to 1.5 times the 99th percentile.⁴ We then calculate the total load within a region as the sum of net generation across all balancing authorities and fuel categories within a region. Excess load is the total load in a region minus net generation from solar and wind, the two nondispatchable energy sources.

⁴This process mimics that used by the EIA when aggregating net generation by balancing authority to the region level, see the "Net Generation" section on this page.

C.2 Theory and Quantitative Appendix: For Online Publication

C.2.1 States Without Net Metering.

In the general model above, we assumed that households could sell back electricity produced by their solar panels at price p_j . This is the case if the state offers net-metering, which is offered in all but 9 in our sample.⁵ In the states where net metering is not offered, households can sell back electricity to the grid at price $p_j^{sale} \leq p_j$. Let $A_i^{home} \left(N_i, \{e_{it}\}_{t=0}^T, \{A_{it}\}_{t=0}^T \right)$ give the discounted sum of energy that is used at home, written as a function of panels installed, electricity consumption, and the stream of solar irradiance. Let $A_i^{grid} \left(N_i, \{e_{it}\}_{t=0}^T, \{A_{it}\}_{t=0}^T \right)$ be the discounted sum of energy that is sold back to grid, such that $A_i^{home}(\cdot) + A_i^{grid}(\cdot) = A_i$. We can write the budget constraint for households in states without net metering as

$$c + \underbrace{p_j (e - m_i N_i A_i^{home}(\cdot))}_{\text{Cost of electricity}} + \underbrace{m_i [(1 - s_j^{\text{Cost}}) p_j^{\text{Ins}}(N_i)]}_{\text{Net cost of installation}} = y_i + m_i \left(\underbrace{N_i A_i s_j^{\text{kWh}}}_{\text{kWh Subsidy}} + \underbrace{N_i s_j^{\text{Panel}}}_{\text{Per-Panel Subsidy}} + \underbrace{p_j^{sale} N_i A_i^{grid}(\cdot)}_{\text{Electricity sold to grid}} \right). \quad (\text{C.1})$$

In estimating and simulating the model, we assume that the household's optimal electricity consumption, $\{e_{it}^*\}_{t=0}^T$, is independent of the household's installation decision. Again letting N_i^* represent the optimal choice of panels, we can then summarize the decision for households in states without net metering as

$$V_i = \max_{N_i, m_i \in \{0,1\}} + m_i [\hat{\mu}_{ij}(N_i) + \gamma_i(N_i^*)],$$

where

$$\hat{\mu}_{ij}(N_i) = N_i A_i (\hat{p}_j + s_j^{\text{kWh}}) - (1 - s_j^{\text{Cost}}) p_j^{\text{Ins}}(N_i) + N_i s_j^{\text{Panel}}$$

⁵Idaho, Tennessee, Texas and Alabama do not have statewide mandatory net-metering policies. Idaho does not have state net-metering policy but each of the state's three investor-owned utilities have a net-metering policy. Five other states in our sample have distributed generation rules other than net-metering.

and where $\hat{p}_j = p_j^{sale} \frac{A_i^{grid}(N_i, \{e_{it}^*\}_{t=0}^T, \{A_{it}\}_{t=0}^T)}{A_i} + p_j \frac{A_i^{home}(N_i, \{e_{it}^*\}_{t=0}^T, \{A_{it}\}_{t=0}^T)}{A_i}$ is the average of the purchasing and sales price of electricity, weighted by the fractions of electricity the household uses at home and sells back to the grid at the optimum.

For data on p_j^{sale} , we use the marginal cost of electricity as measured by Borenstein and Bushnell (2022). One challenge empirically is that we do not have disaggregated data on $A_i^{grid}(\cdot)$ or $A_i^{home}(\cdot)$. Therefore, we assume that the amount of electricity that is sold back to the grid is given by the reduced-form expression $A_i^{grid} = \left(N_i, \{e_{it}^*\}_{t=0}^T, \{A_{it}\}_{t=0}^T \right) = A_i C N_i$, where C is a constant.

The household's optimal number of panels is then given by

$$N_i^* = \min \left[\bar{N}_i, - \left(\frac{\frac{\partial \mu_{ij}}{\partial N_i} + \gamma_{1N}}{2(\gamma_{2N} - C A_i (p_j - p_j^{sale}))} \right) \right].$$

where, as before,

$$\mu_{ij}(N_i) = N_i A_i (p_j + s_j^{\text{kWh}}) - (1 - s_j^{\text{Cost}}) p_j^{\text{Ins}}(N_i) + N_i s_j^{\text{Panel}}.$$

We calibrate C such that a household with the average number of panels sells 30% of their electricity back to the grid.⁶

C.2.2 Maximum Likelihood Estimation of Power Plant Policy

Functions. Let y_{kt}^o denote observed production from power plant k in time t , and let $\hat{f}_k(\text{ELoad}_t | \psi_k) = f_k(\text{ELoad}_t, \varepsilon_{kt}) - \varepsilon_{kt}$ denote the deterministic portion of the latent variable for power plant k in time t , written as a function the ψ_k^0 , ψ_{Rk}^1 and ψ_{Rk}^2 parameters, which we collective denote by ψ_k . The log-likelihood contribution of a given hour of power plant k 's production is

⁶<https://www.seia.org/initiatives/net-metering>

$$\begin{aligned}
\log \mathcal{L}_{kt}(\text{ELoad}_t | \psi_k, \sigma_k^2) = & \mathbb{1}(y_{kt}^o = 0) \times \log \left(\Phi \left(\frac{\hat{f}_k(\text{ELoad}_t | \psi_k)}{\sigma_k} \right) \right) + \\
& \mathbb{1}(y_{kt}^o \in (0, \bar{y}_k)) \times \log \left(\frac{1}{\sigma_k} \phi \left(\frac{y_{kt}^o - \hat{f}_k(\text{ELoad}_t | \psi_k)}{\sigma_k} \right) \right) + \\
& \mathbb{1}(y_{kt}^o \geq \bar{y}_k) \times \log \left(1 - \Phi \left(\frac{\bar{y}_k - \hat{f}_k(\text{ELoad}_t | \psi_k)}{\sigma_k} \right) \right),
\end{aligned} \tag{C.2}$$

where $\mathbb{1}(\cdot)$ represents an indicator functions which turns on if y_{kt}^o is equal to a given value or falls within a certain range, Φ is the standard normal CDF, and ϕ is the standard normal PDF. We choose the structural parameters for each power plant k by maximizing the sum of log likelihood contributions over all hours for that power plant. We restrict the parameter estimates such that output is weakly increasing in excess load for each region over the range of excess load observed in the data.

C.2.3 Details: Cost-Neutral Reforms. We can express the government's constrained maximization problem as the Lagrangian

$$W = \int_i V_i di - D(\text{ELoad}) - \lambda \left(\sum_j \int_{i \in I_j} s_{ij} m_i^* di - G \right), \tag{C.3}$$

where $D(\text{ELoad})$ is total environmental damages, $s_{ij} = s_j^{\text{Panel}} N_i^* + s_j^{\text{kWh}} A_i N_i^* + s_j^{\text{Cost}} p_j^{\text{Ins}}(N_i^*)$ is the total subsidy paid to household i conditional on installation, and G is the maximum amount the government can spend on subsidies.

The nationally-optimal set of subsidies must satisfy the first-order conditions of the government's problem. Taking the derivative of W with respect to a given subsidy type $\theta \in \{\text{kWh}, \text{Panel}, \text{Cost}\}$ in state j yields

$$\begin{aligned} \frac{\partial W}{\partial s_j^\theta} = & \int_i \frac{\partial V_i}{\partial s_j^\theta} di + \int_i \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\text{ELoad}_t^{SB})}{\partial E_{Rt}^{\text{Solar}}} \right| \left(\vec{m}_i^\theta N_i^* + m_i^* \frac{\partial N_i}{\partial s_j^\theta} \right) di \\ & - \lambda \left(\int_i \vec{m}_i^\theta s_{ij} di - \int_i m_i^* \frac{\partial N_i}{\partial s_j^\theta} \frac{\partial s_{ij}}{\partial N_i} di - \int_i m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta} di \right), \quad (\text{C.4}) \end{aligned}$$

where ELoad_t^{SB} denotes the excess load in time t evaluated at the nationally-optimal (welfare-maximizing) system of subsidies.

By the envelope theorem we have $\frac{\partial V_i}{\partial s_j^\theta} = m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta}$, which tells us that the utility gain for households is simply equal to the value of the increase in subsidy for non-additional households, holding the number and size of installations constant.

Plugging this into (C.4) and setting the derivative equal to 0 yields

$$\begin{aligned} \int_i \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\text{ELoad}_t^{SB})}{\partial E_{Rt}^{\text{Solar}}} \right| \left(\vec{m}_i^\theta N_i^* + m_i^* \frac{\partial N_i}{\partial s_j^\theta} \right) di \\ - \lambda \left(\int_i \vec{m}_i^\theta s_{ij} di - \int_i m_i^* \frac{\partial N_i}{\partial s_j^\theta} \frac{\partial s_{ij}}{\partial N_i} di \right) - (\lambda - 1) \int_i m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta} di = 0, \quad (\text{C.5}) \end{aligned}$$

This can be rewritten as

$$\begin{aligned} \int_i \vec{m}_i^\theta di \times \left(\frac{\int_i \vec{m}_i^\theta \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\text{ELoad}_t^{SB})}{\partial E_{Rt}^{\text{Solar}}} \right| N_i^*}{\int_i \vec{m}_i^\theta di} - \lambda \frac{\int_i \vec{m}_i^\theta s_{ij} di}{\int_i \vec{m}_i^\theta di} \right) + \\ \int_i \frac{\partial N_i}{\partial s_j^\theta} \times \left(\frac{\int_i m_i^* \frac{\partial N_i}{\partial s_j^\theta} \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\text{ELoad}_t^{SB})}{\partial E_{Rt}^{\text{Solar}}} \right|}{\int_i \frac{\partial N_i}{\partial s_j^\theta}} - \lambda \frac{\int_i \frac{\partial N_i}{\partial s_j^\theta} \frac{\partial s_{ij}}{\partial N_i}}{\int_i \frac{\partial N_i}{\partial s_j^\theta}} \right) + \\ (1 - \lambda) M_j \frac{\int_i m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta} di}{\int_i m_i^* di}. \quad (\text{C.6}) \end{aligned}$$

Finally, plugging in the definitions of $\frac{\partial M_j}{\partial s_j^\theta}$, $\overrightarrow{\Delta D}_j^{\theta, \text{ext}}$, $\vec{s}_j^{\theta, \text{ext}}$, $\frac{\partial N_j}{\partial s_j^\theta} \Big|_{M_j^{\text{st}}}$, $\overrightarrow{\Delta D}_j^{\theta, \text{int}}$, $\frac{\partial s}{\partial N_j} \theta, \text{int}$, and $\overrightarrow{\frac{\partial s_{ij}}{\partial s_j^\theta}}$ yields (4.8).

C.2.4 Numerical Algorithm for Calculating Optimal Subsidies.

In this appendix, we outline the numerical algorithm we use to solve for the welfare-maximizing subsidies.

1. Make a guess of the marginal cost of public funds, λ . Call this guess $\hat{\lambda}$.
2. Make a guess of the set of subsidies. Let this matrix of all types of subsidies in all states be denoted by $\hat{s}$.
3. Given the current guess of subsidies, $\hat{s}$, and the guess of the marginal cost of public funds, $\hat{\lambda}$, calculate the first-order conditions of the government's problem for each subsidy type and each state given by (4.8). We use analytical derivatives to evaluate $\frac{\partial M_j}{\partial s_j^\theta}$, $\overrightarrow{\Delta D}_j^{\theta, \text{ext}}$, $\overrightarrow{s}_j^{\theta, \text{ext}}$, $\frac{\partial N_j}{\partial s_j^\theta} \bigg|_{M_j^{\text{st}}}$, $\overrightarrow{\Delta D}_j^{\theta, \text{int}}$, and $\frac{\overrightarrow{\partial s}}{\partial N_j} \theta, \text{int}$.
4. If all of the first-order conditions are sufficiently close to 0, move on to the next step. If not, update the guess of the subsidies and return to Step 3.
5. Given the current guess of subsidies, calculate the total government cost.
6. If the government cost is sufficiently close to G , then the current guesses, $\hat{\lambda}$ and $\hat{s}$, solve the constrained maximization problem. If not, take a new guess for $\hat{\lambda}$ and return to Step 2.

C.2.5 Details: Damage-Minimizing Subsidies.

The government's problem is to choose subsidies to minimize national damages, $D(\text{ELoad})$, subject to the budget constraint that the total spending on subsidies cannot exceed some value G :

$$\sum_j \int_{i \in I_j} s_{ij} m_i^* di \leq G,$$

where $s_{ij} = s_j^{\text{Panel}} N_i^* + s_j^{\text{kWh}} A_i N_i^* + s_j^{\text{Cost}} p_j^{\text{Ins}} (N_i^*)$ is the total subsidy paid to household i conditional on installation, and G is the maximum amount the government can spend on subsidies.

We can express this constrained optimization problem as the Lagrangian

$$W = -D(\text{ELoad}) - \lambda \left(\sum_j \int_{i \in I_j} s_{ij} m_i^* di - G \right). \quad (\text{C.7})$$

Taking the derivative of W with respect to s_j^θ yields

$$\begin{aligned} \frac{\partial W}{\partial s_j^\theta} = \int_i \sum_{t=0}^T \frac{A_{it}}{(1+r)^t} \left| \frac{\partial D_t(\text{ELoad}_t^{MD})}{\partial E_{Rt}^{\text{Solar}}} \right| \left(\vec{m}_i^\theta N_i^* + m_i^* \frac{\partial N_i}{\partial s_j^\theta} \right) di \\ - \lambda \left(\int_i \vec{m}_i^\theta s_{ij} di - \int_i m_i^* \frac{\partial N_i}{\partial s_j^\theta} \frac{\partial s_{ij}}{\partial N_i} di - \int_i m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta} di \right), \end{aligned} \quad (\text{C.8})$$

where ELoad_t^{MD} denotes the excess load in time t evaluated at the nationally-optimal (damage-minimizing) system of subsidies.

As in Online Appendix C.2.3, we can again use $\frac{\partial V_i}{\partial s_j^\theta} = m_i^* \frac{\partial s_{ij}}{\partial s_j^\theta}$ by the envelope theorem. Plugging this in and using the definitions of $\frac{\partial M_j}{\partial s_j^\theta}$, $\overrightarrow{\Delta D}_j^{\theta, \text{ext}}$, $\vec{s}_j^{\theta, \text{ext}}$, $\frac{\partial N_j}{\partial s_j^\theta} \Big|_{M_j^{\text{st}}}$, $\overrightarrow{\Delta D}_j^{\theta, \text{int}}$, $\frac{\partial s_{ij}}{\partial N_j}$, and $\frac{\partial s_{ij}}{\partial s_j^\theta}$ yields (C.9), which gives the first-order condition for each subsidy type in each state j :

$$\underbrace{\frac{\partial M_j}{\partial s_j^\theta} \times \left(\overrightarrow{\Delta D}_j^{\theta, \text{ext}} - \lambda \vec{s}_j^{\theta, \text{ext}} \right)}_{\text{Extensive Margin}} + \underbrace{\frac{\partial N_j}{\partial s_j^\theta} \Big|_{M_j^{\text{st}}} \times \left(\overrightarrow{\Delta D}_j^{\theta, \text{int}} - \lambda \frac{\partial s_{ij}}{\partial N_j} \right)}_{\text{Intensive Margin}} - \underbrace{\lambda M_j \frac{\partial s_{ij}}{\partial s_j^\theta}}_{\text{Mechanical Effect}} = 0. \quad (\text{C.9})$$

These optimality conditions for a damage-minimizing planner share a similar structure to those of the welfare-maximizing planner given by (4.8). The exception is how the two planners value increases in subsidies given to non-additional households, which are represented the third term in each of the first-order conditions (“Mechanical Effect”). For the damage-minimizing planner, increased subsidies for these non-additional households entail a fiscal cost with

no additional decrease in damages. Therefore, the number of non-additional households (M_j) enters negatively into the first order condition. The welfare-maximizing planner, on the other hand, values the increase in utility associated with increases in subsidies for non-additional households. Therefore, each additional dollar of subsidies for a non-additional household is valued at $(1 - \lambda)$, reflecting both this increase in utility and the fiscal cost.

C.3 Results Appendix: For Online Publication

C.3.1 Installation Prices. Table C.3 shows the results for estimating solar system installation prices using the Tracking the Sun data using the following regression,

$$p_{ijt}^{\text{Ins}} = p_{R(j)}^{0,\text{Ins}} + p_{R(j)}^{1,\text{Ins}} N_{ijt} + p_{R(j)}^{2,\text{Ins}} (t - 2017) + \varepsilon_{ijt}, \quad (\text{C.10})$$

where p_{ijt}^{Ins} is the installation cost paid by household i in state j and year t , N_{ijt} is the number panels in the installation, $p_{R(j)}^{0,\text{Ins}}$ is a fixed cost, $p_{R(j)}^{1,\text{Ins}}$ is a per-panel cost, $p_{R(j)}^{2,\text{Ins}}$ is a linear time trend for year t , and $R(j)$ is the Census Region containing state j . The table shows results for the full sample and each region, where the intercept gives the fixed installation cost in 2017, and the coefficient on the number of panels is the per-panel cost in 2017. The linear model is a good fit for the data, as seen in Figure C.6, which shows our fitted line against a flexible smoothing function for each region.

C.3.2 Relationship Between Installations and Monetary

Incentives. Table C.4 regresses tract-level log installations on the monetary benefits of installation, where again we calculate the monetary benefits of installation as $\mu_{ij}(N_i^*)$ evaluated at $N_i^* = 15$, the average number of panels in a solar system in the data. Specifications with “Demographic Controls” include

Table C.3. Solar system installation prices

Dependent Variable:	Total Cost				
Census Region	Full sample	Midwest	Northeast	South	West
Model:	(1)	(2)	(3)	(4)	(5)
<i>Variables</i>					
Constant	5,899.5*** (32.9)	4,784.8*** (634.1)	1,825.2*** (62.0)	8,626.4*** (375.3)	7,048.1*** (39.6)
Number of Panels	918.1*** (1.6)	1,060.0*** (26.3)	1,062.1*** (2.6)	703.0*** (14.9)	863.8*** (2.1)
Year (Relative to 2017)	-577.5*** (9.2)	-2,806.6*** (192.6)	-1,165.5*** (15.3)	630.1*** (77.9)	-338.3*** (11.2)
<i>Fit statistics</i>					
Observations	720,665	1,097	171,306	12,354	535,908
R ²	0.53	0.61	0.71	0.46	0.44
Adjusted R ²	0.53	0.61	0.71	0.46	0.44

Heteroskedasticity-robust standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Sample limited to between 2014 and 2018, prices measured in 2014 dollars.




Figure C.6. Estimation results for solar system price regression

The dashed black line is our estimated model in 2017 and the solid blue line shows the fit of a generalized additive model for installations in 2017. Installation prices are measured in 2014 dollars.

controls for tract-level college completion percentage and percent of voters who voted Democrat in the 2016 presidential election. Columns (3) and (4) add Census region fixed effects while columns (5) and (6) include Census division fixed effects. Across all specifications, we find that a \$1000 increase in monetary benefits for a 15-panel installation is associated with a 8.2% to 10.6% increase in installations.

C.3.3 Installation Size Regressions. Table C.5 regresses tract-level data on average number of panels per installation on monetary benefits of installation, where again we calculate the monetary benefits of installation as $\mu_{ij}(N_i^*)$ evaluated at $N_i^* = 15$, the average number of panels in a solar system in the data. Across all specifications, we find that a \$1000 increase in monetary benefits is associated with a 0.07 to 0.11 increase in average panels per installation.

	(1)	(2)	(3)	(4)	(5)	(6)
Monetary Benefits	0.106*** (0.0353)	0.105*** (0.0350)	0.0818*** (0.0166)	0.0850*** (0.0196)	0.0908*** (0.0162)	0.0938*** (0.0182)
Observations	41,776	41,776	41,776	41,776	41,776	41,776
R-squared	0.187	0.201	0.388	0.411	0.421	0.445
Demographic Controls	NO	YES	NO	YES	NO	YES
Region FE	NO	NO	YES	YES	NO	NO
Division FE	NO	NO	NO	NO	YES	YES

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

Table C.4. Regression of log installations on the net present value of total monetary benefits associated with solar panel installations.

Monetary benefits measured in thousands of 2014 dollars. Standard errors clustered by state.

	(1)	(2)	(3)	(4)	(5)	(6)
Monetary Benefits	0.0737*** (0.0222)	0.0856*** (0.0217)	0.0787*** (0.0228)	0.0878*** (0.0221)	0.115*** (0.0316)	0.117*** (0.0321)
Observations	41,776	41,776	41,776	41,776	41,776	41,776
R-squared	0.013	0.018	0.027	0.030	0.042	0.043
Demographic Controls	NO	YES	NO	YES	NO	YES
Region FE	NO	NO	YES	YES	NO	NO
Division FE	NO	NO	NO	NO	YES	YES

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

Table C.5. Regression of average panel size on the total monetary benefits associated with solar panel installations.

Monetary benefits measured in thousands of 2014 dollars. Standard errors clustered by state.




Figure C.7. Border Discontinuities in Household Characteristics.

For each graph, we regress the variable in question on state-border fixed effects and dummy variables for these locations bins and plot these estimated location fixed effects.

C.3.4 Border Discontinuities in Household Characteristics.

Here, we look for evidence of sorting based on preferences for solar panels on either side of state borders. For each graph, we regress the variable in question on state-border fixed effects and dummy variables for these locations bins and plot these estimated location fixed effects. Figures C.7a, C.7b, and C.7c plot these fixed effects for percent with college degree, percent of voters who voted Democrat in the 2016 presidential election, and average household income, respectively. There is no clear pattern in these characteristics on either side of the border.

C.3.5 Border Regressions with Controls for Non-Subsidy

Incentives. In order to measure the prevalence of non-subsidy programs aimed at increasing solar adoption, we collect data from DSIRE on the number of state-implemented policies for solar photovoltaics in the residential sector for each state. We categorized non-subsidy programs into three categories: 1) “Financing Programs,” policies providing loans and financing options to pay for solar installations, 2) “Access Rules,” policies that protect households’ ability to

install solar panels, and 3) “Building Incentives,” programs which create incentives for new building projects to include rooftop solar.⁷ We calculated the number of distinct policies in each of these three categories and added these policy counts as controls in our border discontinuity regressions.

Table C.6 shows the results. The first column shows the border discontinuity regression without these additional policy controls, column (2) adds a control variable for the number of Financing Programs in the states, column (3) adds the number of Access Rules in the state, column (4) add the number of Building Incentives. In all regressions, our main coefficient estimating the elasticity of installations with respect to subsidies is very similar to the baseline results, and all of the non-subsidy programs have the expected positive sign.

In Table C.7, we repeat this analysis but replace the count of policies in each program type with a dummy equal to one if the state has any programs of that type. The results are very similar to those in C.6.

C.3.6 Power Plant Model. Here we present additional information about the power plant model estimation results. Figures C.8 and C.9 show the model fit graphs for the main paper, broken out into regions and evaluated on out-of-sample data from 2020. These show that model performance is consistent within each interconnection.

C.10 shows how fuel mix varies for each region, using in-sample 2019 data.

C.3.7 Type of Subsidy in Nationally-Optimal System. We now analyze how the government should optimally allocate subsidies across the three subsidy types: cost-based subsidies, per-panel subsidies, and production-based

⁷“Financing Programs” are programs classified as “Loan Program” and “PACE Financing” in DSIRE, “Access Rules” are programs classified as “Solar/Wind Access Policy” and “Solar/Wind Permitting Standards” which target solar photovoltaics, and “Building Incentives” are programs classified as “Building Energy Codes” and “Green Building Incentives”.

	Dependent Variable:			
	Log Installations per Capita			
	(1)	(2)	(3)	(4)
NPV Subsidies (\$1000s)	0.0826*** (0.00761)	0.105*** (0.00882)	0.0854*** (0.0142)	0.0810*** (0.0132)
Financing Programs		0.328*** (0.0641)	0.302*** (0.0605)	0.307*** (0.0611)
Access Rules			0.289* (0.163)	0.271 (0.166)
Building Incentives				0.208* (0.110) (0.104)
Demographic Controls	YES	YES	YES	YES
Border FE	YES	YES	YES	YES
Distance Bandwidth	10	10	10	10
Tax Controls	Yes	Yes	Yes	Yes
Observations	6,052	6,052	6,052	6,052

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

Table C.6. Regression of log installations per capita on the net present value of subsidies

For a 15-panel installation within 10 miles of state borders with controls for non-subsidy programs. Subsidies are measured in thousands of 2014 dollars. State-clustered standard errors in parentheses. All regressions contain border fixed effects, tract-level demographic controls, and controls for state tax rates. Non-subsidy policy variables are measured as the number of policies of a given type in the state.

	Dependent Variable:			
	Log Installations per Capita			
	(1)	(2)	(3)	(4)
NPV Subsidies (\$1000s)	0.0826*** (0.00761)	0.107*** (0.00984)	0.103*** (0.0131)	0.0933*** (0.0137)
Has Financing Programs		0.587*** (0.148)	0.563*** (0.125)	0.575*** (0.129)
Has Access Rules			0.0968 (0.212)	0.0827 (0.216)
Has Building Incentives				0.308* (0.171)
Demographic Controls	YES	YES	YES	YES
Border FE	YES	YES	YES	YES
Distance Bandwidth	10	10	10	10
Tax Controls	Yes	Yes	Yes	Yes
Observations	6,052	6,052	6,052	6,052

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

Table C.7. Regression of log installations per capita on the net present value of subsidies

For a 15-panel installation within 10 miles of state borders with controls for non-subsidy programs. Subsidies are measured in thousands of 2014 dollars. State-clustered standard errors in parentheses. All regressions contain border fixed effects, tract-level demographic controls, and controls for state tax rates. Non-subsidy policy variables are dummy variables equal to one if the state has at least one policy of a given type.

	(1)	(2)	(3)
	Baseline	State-Specific Subsidies	
		Welfare Max	Damage Min
Unit Subsidies	7.1	0.8	1.0
Cost Subsidies	80.9	1.5	0.8
kWh Subsidies	8.4	97.7	98.2
Total	100.0	100.0	100.0

Table C.8. Percent of total subsidy value from each type of subsidy.

For a 15-panel installation averaged across all households in the model. Each column shows the subsidy values for a different simulation.




Figure C.8. Model fit at the region level in 2020.

excluding Texas as there is only one region in the Texas interconnection. Dots represent an hour of production for each region, smoothed lines show the fit of a generalized additive model.

(per-kWh) subsidies. To facilitate comparison, we calculate the present discounted value an “average installation” would receive. Specifically, we calculate the subsidy value every household in the model would receive if they purchased a 15-panel installation.⁸ We then average this hypothetical subsidy value over all households. Table C.8 shows the percent of the total subsidy value coming from each subsidy type in each simulation. Under the current system, over 80% of the value of subsidies comes from cost-based subsidies, which include the federal investment tax credit, state investment tax credits, sales tax exemptions, and property tax exemptions. On the other hand, the welfare-maximizing subsidies almost exclusively consist of production-based subsidies. Intuitively, production-based subsidies incentivize installations for households where sunlight, and therefore

⁸We define households as rooftops suitable for solar panel installations as defined by GPS data.




Figure C.9. Model fit at the region level by hour and season in 2020.

The dashed green line gives electricity production in the data while the solid orange line gives predicted production.




Figure C.10. Fuel mix of production by region.

The X-axis gives excess load at the interconnection level and the Y-axis gives the percent of electricity production that is produced by each of the fuel types using data from 2020. The dashed lines show the fuel mix in the data while the solid lines show the simulated fuel mix.

	(1)	(2)	(3)	(4)	(5)
		State-Specific Subsidies		Tract-Specific Subsidies	
	Baseline	Welfare Max	Damage Min	Welfare Max	Damage Min
I. Average Number of Panels per installation					
Midwest	14.7	14.9	15.0	14.9	15.0
Northeast	15.1	14.9	15.0	14.9	15.0
South	14.8	15.1	15.1	15.1	15.1
West	15.0	15.2	15.1	15.2	15.1

Table C.9. Average number of solar panels in a solar installation by Census region.

environmental benefits, are high.⁹ However, as we show in Section 4.6.5, the gains to reallocating across subsidy types within states are small relative to the gains from reallocating subsidies *across* states.

C.3.8 Average Panel Size Across Counterfactuals. Table C.9 shows the average number of panels per installation across Census regions in each simulation. We can see that panel size does not significantly change across regions or across simulations. These results suggest that extensive-margin adjustments shown in the body of the paper play a much more important role quantitatively than the intensive-margin adjustments shown here.

C.3.9 State-Level Results. The first columns of Table C.10 gives the baseline and welfare-maximizing subsidy in each state. The following columns show the simulated number of solar panel installations per 1000 households under the current subsidies and under the welfare-maximizing subsidies.

The first two columns of Table C.11 shows the state-level subsidies given the current system and the damage-minimizing subsidies. The following columns of Table C.11 show the simulated number of solar panel installations per 1000 households given the current subsidies and under the damage-minimizing subsidies.

C.3.10 Unconstrained Nationally-Optimal Subsidies with $\lambda = 1$.

⁹One important caveat is that we assume households' discount rate is given by the inverse of the real interest rate.

	Expected Subsidy		Installs/1000HH			Expected Subsidy		Installs/1000HH	
	Baseline	Optimal	Baseline	Optimal		Baseline	Optimal	Baseline	Optimal
Alabama	5.8	15.1	3.3	7.2	Nebraska	7.6	13.8	4.1	6.9
Arizona	11.2	14.0	8.7	11.1	Nevada	6.6	13.9	4.8	9.0
Arkansas	5.8	13.9	3.3	6.6	New Hampshire	18.1	9.4	16.7	8.0
California	11.7	13.0	15.7	17.7	New Jersey	27.6	17.7	38.7	16.8
Colorado	9.4	12.6	6.2	8.1	New Mexico	10.6	14.4	8.3	11.5
Connecticut	15.7	9.5	18.3	10.9	New York	17.1	9.5	16.2	8.6
Delaware	10.0	18.2	7.0	14.2	North Carolina	8.9	15.1	5.7	9.7
Florida	10.3	16.0	7.2	11.8	North Dakota	10.0	12.6	3.8	4.8
Georgia	5.8	15.1	4.8	10.6	Ohio	6.6	17.2	3.6	8.7
Idaho	7.4	11.7	3.3	4.7	Oklahoma	5.8	14.7	3.3	7.1
Illinois	12.5	17.1	7.0	10.2	Oregon	11.1	9.9	6.2	5.5
Indiana	9.6	17.7	4.4	8.7	Pennsylvania	5.7	17.5	4.2	11.1
Iowa	16.9	13.2	8.0	5.8	Rhode Island	15.9	9.6	16.4	9.6
Kansas	11.7	14.2	6.6	8.1	South Carolina	9.3	15.3	4.7	8.7
Kentucky	6.3	14.3	3.6	7.2	South Dakota	11.3	12.9	4.0	4.5
Louisiana	15.2	14.2	8.2	7.6	Tennessee	7.6	14.7	3.9	7.1
Maine	5.3	9.5	4.7	6.7	Texas	16.0	12.1	8.9	6.4
Maryland	12.5	18.3	10.3	16.8	Utah	8.0	12.3	3.6	5.4
Massachusetts	24.7	9.5	36.6	10.4	Vermont	12.8	9.2	14.8	11.0
Michigan	6.2	12.5	3.7	6.3	Virginia	5.8	15.3	4.1	9.2
Minnesota	12.1	12.8	6.3	6.6	Washington	18.8	9.4	10.3	4.6
Mississippi	5.8	14.9	3.5	7.6	West Virginia	5.8	17.6	2.6	7.1
Missouri	9.4	13.8	5.2	7.4	Wisconsin	16.9	16.0	9.2	8.5
Montana	9.0	10.8	4.1	4.7	Wyoming	6.0	12.1	2.4	4.1

Table C.10. Cost neutral state level subsidy results.

The first two columns shows the state-level subsidies given the current system and the welfare-maximizing subsidies. Subsidies are measured as the average present discounted value of subsidies for a 15-panel installation, measured in thousands of 2014 dollars. The following columns show the simulated number of solar panel installations per 1000 households given the current subsidies and under the welfare-maximizing subsidies.

	Expected Subsidy			Installs/1000HH				Expected Subsidy			Installs/1000HH		
	Baseline	Optimal		Baseline	Optimal			Baseline	Optimal		Baseline	Optimal	
Alabama	5.8	15.6		3.3	7.5		Nebraska	7.6	12.9		4.1	6.3	
Arizona	11.2	13.4		8.7	10.5		Nevada	6.6	13.1		4.8	8.4	
Arkansas	5.8	13.2		3.3	6.2		New Hampshire	18.1	3.9		16.7	5.0	
California	11.7	11.3		15.7	15.3		New Jersey	27.6	20.6		38.7	21.4	
Colorado	9.4	10.5		6.2	6.8		New Mexico	10.6	14.1		8.3	11.2	
Connecticut	15.7	4.1		18.3	6.9		New York	17.1	4.1		16.2	5.4	
Delaware	10.0	21.8		7.0	19.2		North Carolina	8.9	15.5		5.7	10.0	
Florida	10.3	17.3		7.2	13.3		North Dakota	10.0	10.5		3.8	4.0	
Georgia	5.8	15.5		4.8	11.0		Ohio	6.6	19.7		3.6	10.8	
Idaho	7.4	8.7		3.3	3.6		Oklahoma	5.8	14.7		3.3	7.1	
Illinois	12.5	19.3		7.0	12.4		Oregon	11.1	5.0		6.2	3.7	
Indiana	9.6	20.8		4.4	11.4		Pennsylvania	5.7	20.2		4.2	14.0	
Iowa	16.9	11.5		8.0	5.1		Rhode Island	15.9	4.1		16.4	6.1	
Kansas	11.7	13.7		6.6	7.7		South Carolina	9.3	16.0		4.7	9.3	
Kentucky	6.3	13.9		3.6	6.9		South Dakota	11.3	10.8		4.0	3.8	
Louisiana	15.2	13.8		8.2	7.3		Tennessee	7.6	14.6		3.9	7.1	
Maine	5.3	4.1		4.7	4.2		Texas	16.0	9.6		8.9	5.1	
Maryland	12.5	21.9		10.3	22.7		Utah	8.0	9.9		3.6	4.4	
Massachusetts	24.7	3.9		36.6	6.6		Vermont	12.8	3.4		14.8	6.7	
Michigan	6.2	10.0		3.7	5.1		Virginia	5.8	15.9		4.1	9.7	
Minnesota	12.1	10.6		6.3	5.5		Washington	18.8	4.2		10.3	3.0	
Mississippi	5.8	15.1		3.5	7.7		West Virginia	5.8	20.6		2.6	9.2	
Missouri	9.4	12.7		5.2	6.8		Wisconsin	16.9	17.1		9.2	9.4	
Montana	9.0	6.7		4.1	3.3		Wyoming	6.0	9.5		2.4	3.2	

Table C.11. Damage minimizing state level subsidy results.

The first two columns shows the state-level subsidies given the current system and the damage-minimizing subsidies. Subsidies are measured as the average present discounted value of subsidies for a 15-panel installation, measured in thousands of 2014 dollars. The following columns show the simulated number of solar panel installations per 1000 households given the current subsidies and under the damage-minimizing subsidies.

Theory We define welfare as

$$W = \underbrace{\int_i V_i di}_{\text{Utility}} - \underbrace{D(\text{ELoad})}_{\text{Damages}} - \underbrace{\sum_j \int_{i \in I_j} s_{ij} m_i^* di}_{\text{Government Cost}}. \quad (\text{C.11})$$

The nationally-optimal system of subsidies must satisfy $\frac{\partial W}{\partial s_j^\theta} = 0$ for each type of subsidy in each state, which implies

$$\underbrace{\frac{\partial M_j}{\partial s_j^\theta} \times \left(\overrightarrow{\Delta D_j^{\theta, \text{ext}}} - \overrightarrow{s_j^{\theta, \text{ext}}} \right)}_{\text{Extensive Margin}} + \underbrace{\frac{\partial N_j}{\partial s_j^\theta} \bigg|_{M_j^{\text{st}}} \times \left(\overrightarrow{\Delta D_j^{\theta, \text{int}}} - \frac{\overrightarrow{\partial s}}{\partial N_j} \overrightarrow{\theta, \text{int}} \right)}_{\text{Intensive Margin}} = 0, \quad (\text{C.12})$$

where all objects are as defined in Section 4.2.3.

The optimal policy balances two forces: 1) the decrease in damages and 2) the increase in cost due to an increase in the number of panel installed, through both extensive and intensive margin adjustments. Importantly, note that household utility does not show up in this formula. This is because there is no first-order welfare effect on households for marginal households (i.e. households who choose to install solar panels in response to the increase in subsidies) because of the envelope theorem. Further, the utility increase for non-additional households (i.e. households who already chose to install solar panels before the increase in subsidies) associated with receiving a larger subsidy for existing panels is exactly offset by the cost of increasing subsidies for these households.¹⁰

Results Table C.12 presents the baseline and nationally-optimal subsidy in each state. Nationally-optimal subsidies are lowest in Vermont, at \$6,600.

¹⁰This is a direct consequence of 1) quasilinear utility and a utilitarian welfare function with equal Pareto weights, which together imply that marginal social welfare weights (Saez & Stantcheva, 2016) are equalized and there are no effects of total welfare of wealth redistribution, and 2) the assumption of that the marginal cost of public funds is equal to one: the social planner values an increase in consumption for a given household the same as an increase in government revenue.

Nationally-optimal subsidies are over twice as high in most of the Mid-Atlantic, with the highest subsidies at over \$14,400 in Delaware and Maryland.

C.4 Extensions and Robustness Appendix: For Online Publication

C.4.1 Alternative Specifications of Household Utility. Table C.13 recalculates our main results under alternative specifications of household utility. Each entry shows the change in average subsidies, installations, and environmental benefits associated with moving from the current system of subsidies to the welfare-maximizing system of subsidies, given the specification in question. Column (1) considers a specification in which the nonpecuniary component does not depend on tract-level demographics. In (2), we add the tract-level fraction of individuals with a college education and in (3) we also add the percent of voters who voted Democrat in the 2016 presidential election. To these variables, (4) adds home-ownership rate and (5) instead adds average income. Note that Column (3) is the same as our baseline specification. The results are qualitatively very similar across specifications.

C.4.2 Alternative Discounting of Future Subsidy Payments. The optimal subsidy results where households discount future subsidy payments at an implicit interest rate of 15% are summarized in Table C.14. In all three specifications, the distribution of subsidies and installations across space is similar to those in the baseline model.

Table C.15 shows how the government should optimally allocate subsidies across the three subsidy types: cost-based subsidies, per-panel subsidies, and production-based (per-kWh) subsidies. For this, we calculate the subsidy value every household in the model would receive if they purchased a 15-panel installation and then average this hypothetical subsidy value over all households.

	Expected Subsidy			Expected Subsidy	
	Baseline	Optimal		Baseline	Optimal
Alabama	5.8	11.7	Nebraska	7.6	10.6
Arizona	11.2	10.7	Nevada	6.6	10.6
Arkansas	5.8	10.7	New Hampshire	18.1	6.7
California	11.7	9.8	New Jersey	27.6	13.9
Colorado	9.4	9.5	New Mexico	10.6	11.1
Connecticut	15.7	6.8	New York	17.1	6.8
Delaware	10.0	14.4	North Carolina	8.9	11.7
Florida	10.3	12.4	North Dakota	10.0	9.6
Georgia	5.8	11.7	Ohio	6.6	13.5
Idaho	7.4	8.7	Oklahoma	5.8	11.3
Illinois	12.5	13.3	Oregon	11.1	7.1
Indiana	9.6	14.0	Pennsylvania	5.7	13.8
Iowa	16.9	10.0	Rhode Island	15.9	6.8
Kansas	11.7	10.9	South Carolina	9.3	11.9
Kentucky	6.3	11.0	South Dakota	11.3	9.7
Louisiana	15.2	10.9	Tennessee	7.6	11.3
Maine	5.3	6.8	Texas	16.0	9.1
Maryland	12.5	14.5	Utah	8.0	9.3
Massachusetts	24.7	6.8	Vermont	12.8	6.6
Michigan	6.2	9.4	Virginia	5.8	11.9
Minnesota	12.1	9.6	Washington	18.8	6.8
Mississippi	5.8	11.5	West Virginia	5.8	13.9
Missouri	9.4	10.5	Wisconsin	16.9	12.5
Montana	9.0	7.9	Wyoming	6.0	9.1

Table C.12. Unconstrained state level subsidy results.

The first two columns shows the state-level subsidies given the current system and the unconstrained nationally-optimal subsidies. Subsidies are measured as the average present discounted value of subsidies for a 15-panel installation, measured in thousands of 2014 dollars. The following columns show the simulated number of solar panel installations per 1000 households given the current subsidies and under the welfare-maximizing subsidies.

	(1)	(2)	(3)	(4)	(5)
I. Δ Average Subsidy (\$Thousands)					
Midwest	2.6	2.6	2.7	2.7	2.7
Northeast	-8.3	-8.6	-9.1	-9.1	-9.3
South	2.4	2.4	2.4	2.5	2.4
West	1.0	1.1	1.2	1.2	1.4
II. Δ Installations per 1000HHs					
Midwest	5.1	5.2	5.4	5.4	5.5
Northeast	-4.3	-4.2	-4.0	-4.1	-3.9
South	3.7	3.8	3.9	3.9	4.1
West	0.7	0.8	0.9	0.9	1.0
III. Δ Annual Damages Offset (\$Millions)					
Total	29.4	30.0	30.3	30.3	30.1
Nonpecuniary Component Depends On:					
College Share	No	Yes	Yes	Yes	Yes
Percent Democrat	No	No	Yes	Yes	Yes
Homeowner Share	No	No	No	Yes	No
Average Income	No	No	No	No	Yes

Table C.13. Counterfactual results under alternative model specifications.

Each entry shows the change of moving from the current system of subsidies to the nationally-optimal cost-neutral system of subsidies given the specification in question. The first panel shows the change in the average present discounted value of subsidies for a 15-panel installation for each census region. The second panel gives the change in the simulated number of solar installations per 1000 households in the model for each Census region. The final panel gives the change in total damages offset by rooftop solar. All monetary values are measured in 2014 dollars. See text for details on each model specification. “Percent Democrat” refers to the percent of voters who voted Democrat in the 2016 presidential election.

	(1)	(2)	(3)	(4)
	State-Specific Subsidies			
	Baseline	Welfare Max	Damage Min	Unconstrained
I. Average Subsidy (\$Thousands)				
Midwest	10.0	15.4	15.7	11.7
Northeast	17.1	13.1	11.7	9.9
South	10.6	14.6	14.5	11.2
West	11.7	12.6	10.7	9.6
II. Installations per 1000HHs				
Midwest	5.3	8.2	8.8	5.9
Northeast	20.7	11.4	11.3	8.6
South	6.5	9.1	9.3	6.7
West	11.9	13.2	11.5	10.1
National	9.9	10.4	10.1	7.8
III. Annual Damages Offset (\$Millions)				
Total	263.6	293.6	296.5	217.0
IV. Annuitized Total Subsidies Paid (\$Millions)				
	383.3	383.3	383.3	217.0

Table C.14. Nationally-optimal subsidies with alternative household discounting.

The first panel shows the average present discounted value of subsidies received for a 15-panel installation for each census region. The second panel gives the simulated number of solar installations per 1000 households in the model for each Census region. The third panel gives the total damages offset by rooftop solar. The final panel gives total government cost under each subsidy scheme converted to an annuity value. “Unconstrained” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$). All monetary values are measured in 2014 dollars.

	(1)	(2)	(3)	
	Current	State-Specific Optimal Subsidies		
	Subsidies	Welfare Max	Damage Min	Unconstrained
Unit Subsidies	7.1	16.9	16.4	10.4
Cost Subsidies	80.9	83.1	83.6	89.6
kWh Subsidies	8.4	0.0	0.0	0.0
Total	100.0	100.0	100.0	100.0

Table C.15. Percent of total subsidy value from each type of subsidy

for a 15-panel installation averaged across all households in the model with alternative household discounting. Each column shows the subsidy values for a different simulation. “Unconstrained” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$).

We then calculate the percent of the total subsidy value coming from each subsidy type in each simulation. In all of the optimal subsidy systems, the government relies entirely on upfront subsidies. This makes sense, as production subsidies are not as highly valued, per dollar, as these other two subsidies types.

Table C.16 shows the damages offset per additional dollar of government cost associated with marginal subsidy increases around the current system of subsidies. As before, we calculate this by first simulating the model 1) under the current system of subsidies and 2) with marginally more generous subsidies of a given type in a given state. We then calculate the damages offset per dollar of this particular subsidy as the difference in damages between the two simulations divided by the difference in the fiscal cost. We repeat this process for each subsidy type in each state.

We can see that production subsidies are much less cost-effective than investment and per-panel subsidies. This makes sense, as households in this model heavily discount subsidies received in the future.

	Subsidy Type				Subsidy Type		
	Panel	Cost	kWh		Panel	Cost	kWh
Alabama	0.68	0.68	0.32	Nebraska	0.56	0.55	0.28
Arizona	0.41	0.41	0.24	Nevada	0.60	0.59	0.29
Arkansas	0.62	0.62	0.30	New Hampshire	0.23	0.23	0.14
California	0.43	0.43	0.24	New Jersey	0.36	0.36	0.24
Colorado	0.46	0.46	0.24	New Mexico	0.51	0.50	0.27
Connecticut	0.25	0.25	0.15	New York	0.24	0.24	0.14
Delaware	0.68	0.67	0.36	North Carolina	0.58	0.57	0.30
Florida	0.58	0.58	0.31	North Dakota	0.45	0.45	0.24
Georgia	0.69	0.68	0.32	Ohio	0.75	0.75	0.36
Idaho	0.47	0.47	0.23	Oklahoma	0.67	0.66	0.32
Illinois	0.56	0.56	0.31	Oregon	0.27	0.27	0.16
Indiana	0.67	0.66	0.35	Pennsylvania	0.80	0.79	0.38
Iowa	0.36	0.35	0.21	Rhode Island	0.25	0.25	0.15
Kansas	0.47	0.47	0.26	South Carolina	0.55	0.54	0.29
Kentucky	0.39	0.39	0.24	South Dakota	0.43	0.43	0.23
Louisiana	0.42	0.41	0.24	Tennessee	0.60	0.59	0.30
Maine	0.40	0.40	0.19	Texas	0.34	0.33	0.20
Maryland	0.61	0.60	0.34	Utah	0.42	0.41	0.23
Massachusetts	0.18	0.18	0.12	Vermont	0.27	0.27	0.15
Michigan	0.53	0.53	0.25	Virginia	0.69	0.68	0.33
Minnesota	0.41	0.41	0.23	Washington	0.22	0.22	0.14
Mississippi	0.67	0.66	0.32	West Virginia	0.81	0.80	0.38
Missouri	0.50	0.50	0.26	Wisconsin	0.44	0.44	0.26
Montana	0.39	0.38	0.20	Wyoming	0.52	0.52	0.25

Table C.16. Damages offset per additional dollar of government funds

associated with marginal subsidy increases around the current system of subsidies with alternative household discounting. The first column (“Panel”) gives the damages offset per dollar associated with marginal increases in panel-based subsidies, s_j^{Panel} , the second column (“Cost”) gives the damages offset per dollar associated with marginal increases in cost-based subsidies, s_j^{Cost} , and the third column (“kWh”) gives the damages offset per dollar associated with marginal increases in production-based subsidies, s_j^{kWh} .

C.4.3 Line Losses. We use the methodology from Borenstein and Bushnell (2022) to account for line losses between the power plant and households. Formally, losses for each region come from a constant plus a factor proportional to the square of flow on the line: $L_{Rt} = \alpha_{1R} + \alpha_{2R} (\text{Load}_{Rt} - E_{Rt}^{\text{Solar}})^2$. Note that the parameters α_{1R} and α_{2R} are both allowed to vary by region to reflect differences in grid characteristics across regions. We then adjust excess load by those losses, $\text{ELoad}_{Rt}^{LL} = \text{Load}_{Rt} - E_{Rt}^{\text{NonD}} - E_{Rt}^{\text{Solar}} + L_{Rt}$. Losses enter positively since power plants must produce not only the amount of electricity demanded by households but also must make up for the losses incurred in transporting electricity to the household. Adding line losses changes the marginal damages offset by residential solar to

$$\left| \frac{\partial D(\text{ELoad})}{\partial N_i} \right| = \sum_{t=0}^T \frac{1}{(1+r)^t} A_{it} (1 + 2\alpha_{2R} (\text{Load}_{Rt} + E_{Rt}^{\text{Solar}})) \left| \frac{\partial D_t(\cdot)}{\partial \text{ELoad}_{Rt}} \right|.$$

The installation of solar panels now has two benefits. As we have in our primary model, solar panels reduce the electricity demand fulfilled by power plants, generating benefits equal to the electricity produced by a panel, A_{it} , times the change in damages, $\left| \frac{\partial D_t(\cdot)}{\partial \text{ELoad}_{Rt}} \right|$. Now, there is an additional benefit from offsetting line losses, captured by the term $2\alpha_{2R} (\text{Load}_{Rt} - E_{Rt}^{\text{Solar}})$, which is the marginal change in losses. Including line losses increase the average damages offset by marginal installers, $\overrightarrow{\Delta D}_j^{\theta, \text{ext}}$, and by intensive margin installers, $\overrightarrow{\Delta D}_j^{\theta, \text{int}}$, when calculating nationally-optimal subsidies.¹¹

Borenstein and Bushnell (2022) estimate line losses as a proportion of total production for over 1,600 utilities in the United States. We take the weighted average of these estimates to create values for each region, weighting by the total

¹¹One caveat is that we do not adjust electricity production from solar panels to account for line losses when residential solar panels transmit electricity back into the grid.

electricity production of each utility. Let γ_R be line losses as a proportion of total production in region R . We then follow their assumption that 25% of line losses are independent of flow on the line, which allows us to back out $\alpha_1 = 0.25\gamma_R \sum_t (\text{Load}_{Rt} - E_{Rt}^{\text{Solar}})$ and $\alpha_2 = 0.75\gamma_R \frac{\sum_t (\text{Load}_{Rt} - E_{Rt}^{\text{Solar}})}{\sum_t (\text{Load}_{Rt} - E_{Rt}^{\text{Solar}})^2}$.

Results The results are summarized in Table C.17. The first column gives cost-neutral subsidies, installations, and damages offset given the current system of subsidies when we account for line losses. The annual damages offset are slightly larger than the baseline model in which we do not account for line losses.

The following summarize the results under (2) state-specific welfare-maximizing subsidies, (3) state-specific damage-minimizing subsidies, (4) state-specific unconstrained nationally-optimal subsidies. In all three counterfactuals, the subsidies and installations are similar to those in the baseline model when we do not account for line losses. However, the environmental gains are larger than in the baseline model.

C.4.4 Transmission Constraints. The results in the model with transmission constraints are summarized in Table C.17. The first column gives cost-neutral subsidies, installations, and damages offset given the current system of subsidies when we account for transmission constraints. The annual damages offset are slightly larger than the baseline model in which we do not account for line losses.

The following columns summarize the results under (2) state-specific welfare-maximizing subsidies, (3) state-specific damage-minimizing subsidies, (4) state-specific unconstrained nationally-optimal subsidies. We find that current spending on subsidies exceeds the optimal amount by roughly 45% in this case.

	(1)	(2)	(3)	(4)
	State-Specific Subsidies			
	Baseline	Welfare Max	Damage Min	Unconstrained
I. Average Subsidy (\$Thousands)				
Midwest	10.0	15.4	16.1	13.1
Northeast	17.1	13.1	11.6	10.9
South	10.6	14.5	14.2	12.2
West	11.7	12.5	10.8	10.4
II. Installations per 1000HHs				
Midwest	5.4	8.2	9.0	6.7
Northeast	20.3	11.3	11.2	9.4
South	6.6	8.9	9.1	7.4
West	11.9	13.1	11.6	11.0
National	9.9	10.3	10.1	8.5
III. Annual Damages Offset (\$Millions)				
Total	288.5	321.6	324.0	264.2
IV. Annuitized Total Subsidies Paid (\$Millions)				
	380.8	380.8	380.8	263.9

Table C.17. Nationally-optimal subsidies when accounting for line losses.

The first panel shows the average present discounted value of subsidies received for a 15-panel installation for each census region. The second panel gives the simulated number of solar installations per 1000 households in the model for each Census region. Households are defined as rooftops which are suitable for solar panel installations as defined by GPS data. The third panel gives the total damages offset by rooftop solar. The final panel gives total government cost under each subsidy scheme converted to an annuity value. “Unconstrained” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$). All monetary values are measured in 2014 dollars.

	(1)	(2)	(3)	(4)
	State-Specific Subsidies			
	Baseline	Welfare Max	Damage Min	Unconstrained
I. Average Subsidy (\$Thousands)				
Midwest	10.0	14.4	14.4	12.0
Northeast	17.1	15.5	16.6	13.1
South	10.6	12.8	11.4	10.6
West	11.7	13.6	13.1	11.4
II. Installations per 1000HHs				
Midwest	5.4	7.4	7.5	6.1
Northeast	20.3	13.7	15.1	11.1
South	6.6	7.6	6.9	6.3
West	11.9	14.4	14.2	11.8
National	9.9	10.4	10.3	8.5
III. Annual Damages Offset (\$Millions)				
Total	299.1	311.7	313.5	256.3
IV. Annuitized Total Subsidies Paid (\$Millions)				
	380.8	380.8	380.8	259.7

Table C.18. Nationally-optimal subsidies when accounting for transmission constraints.

The first panel shows the average present discounted value of subsidies received for a 15-panel installation for each census region. The second panel gives the simulated number of solar installations per 1000 households in the model for each Census region. Households are defined as rooftops which are suitable for solar panel installations as defined by GPS data. The third panel gives the total damages offset by rooftop solar. The final panel gives total government cost under each subsidy scheme converted to an annuity value. “Unconstrained” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$). All monetary values are measured in 2014 dollars.

C.4.5 Improved Storage of Nondispatchable Electricity.

Because of intermittent nature of many renewable energy sources, times when renewable energy generation is high may not correspond with times when electricity demand is high. Improvements in energy storage technology would allow electricity generated by nondispatchable energy sources to be stored for times when it is most needed. What would be the environmental benefits of these improvements in energy storage technology? And how would the introduction of improved electricity storage technology change the nationally-optimal system of solar subsidies?

As a simple way to try to answer these questions, we consider a setting in which electricity produced by nondispatchable sources (including household solar) can be imperfectly reallocated over time. Specifically, given the total amount of electricity produced by nondispatchable sources in a year, we assume a proportion ω of this electricity is reallocated over time such that the profile of *usage* of this reallocated electricity is proportional to electricity demand.¹² Formally, we write excess demand as

$$\text{ELoad}_{Rt}^{\text{storage}} = \left(1 - \underbrace{\omega (A_R^{\text{NonD}} + A_R^{\text{Solar}})}_{\text{Reallocated Electricity}} \right) \text{Load}_{Rt} - \underbrace{(1 - \omega) (E_{Rt}^{\text{NonD}} + E_{Rt}^{\text{Solar}})}_{\text{Non-Reallocated Electricity}}$$

where $A_R^{\text{NonD}} = \frac{\sum_t E_{Rt}^{\text{NonD}}}{\sum_t \text{Load}_{Rt}}$ and $A_R^{\text{Solar}} = \frac{\sum_t E_{Rt}^{\text{Solar}}}{\sum_t \text{Load}_{Rt}}$ are region-specific constants

which ensure that total amount of nondispatchable energy utilized is equal to total nondispatchable energy generated.¹³

¹²This is highly stylized model of electricity storage. More generally, optimal storage and withdrawal of electricity will depend on the distribution of the cost of electricity production by other sources over time and space. See Holland, Mansur, and Yates (2022) for a richer model of electricity storage. It would be straightforward to only be reallocated within the same day it is generated.

¹³Similar to Holland, Mansur, and Yates (2022), we assume that there are no electricity losses associated with electricity storage, e.g. from charging batteries or decay of electricity over time.

Results We calculate the environmental benefits of this improved storage technology and the nationally-optimal subsidies given the new storage technology for three values of ω in Table C.19. Column (2), for example, shows the effects of this alternative storage technology with $\omega = .25$, holding the system of solar subsidies at their current levels. As subsidies do not change, the distribution of installations is the same as in the case without storage technology. Panel III shows that the improved storage technology leads to a decrease in environmental damages valued at over \$125 million annually. Column (3) recalculates the nationally-optimal cost-neutral subsidies given that the new storage technology is in place. We find that the nationally-optimal subsidies are very similar to the baseline case and that implementing the nationally-optimal subsidies leads to similar reductions in environmental damages as we find without the improved storage technology. Column (4) shows the unconstrained nationally-optimal subsidies given the new storage technology. The nationally-optimal unconstrained subsidies are very similar to the baseline case.

The remaining columns repeat this exercise for $\omega = .5$ and $\omega = .75$. In both scenarios, we find large environmental benefits to the new technology. However, the nationally-optimal subsidies and the environmental benefits associated with implementing those subsidies are similar to those in the baseline case.

C.4.6 Cleaner Electricity Production. We present our results when we allow changes in electricity production in Table C.20. The first column gives the results under the current technology, as in our baseline results. The first panel gives the change in average subsidy, measured in thousands of dollars when moving from the current subsidies to the cost-neutral welfare-maximizing subsidies. The second panel gives the change in installations per 1000 households. The final

(1)	(2)	(3) $\omega = .25$		(4)		(5)		(6) $\omega = .5$		(7)		(8)		(9) $\omega = .75$		(10)
	Baseline	+Optimal		+Unconstrained		Improved		+Optimal		+Unconstrained		Improved		+Optimal		+Unconstrained
I. Average Subsidy (\$Thousands)	Storage	Subsidies		Optimal		Storage		Subsidies		Optimal		Storage		Subsidies		Optimal
Midwest	10.0	15.2	12.9	11.9	9.9	10.0	17.3	15.1	12.9	11.9	9.9	10.0	17.3	15.1	12.8	11.8
Northeast	17.1	12.9	11.0	9.9	11.2	17.3	11.0	12.9	14.4	9.9	11.2	17.3	11.0	12.8	14.3	9.8
South	10.6	14.4	12.8	11.2	9.8	11.0	12.4	12.9	12.9	11.2	9.9	11.0	12.4	13.0	13.0	11.2
West	11.7	12.4	12.8	9.8		12.4						12.4				10.0
II. Installations per 1000HHs																
Midwest	5.4	8.0		6.0		5.4		7.9		6.0		5.4		7.9		6.0
Northeast	20.3	11.1		8.5		20.3		11.0		8.5		20.3		11.0		8.5
South	6.6	8.9		6.7		6.6		8.8		6.7		6.6		8.8		6.7
West	11.9	13.4		10.4		11.9		13.6		10.5		11.9		13.7		10.6
III. Annual Damages Offset (\$Millions)																
Total	262.3	382.1	412.6	340.0		472.8		503.4		431.6		533.3		563.9		492.1
IV. Annualized Total Subsidies Paid (\$Millions)																
National	380.8	380.8	380.8	222.4		380.8		380.8		224.5		380.8		380.8		224.6

Table C.19. Nationally-optimal cost-neutral and unconstrained subsidies with improved storage technology.

The first panel shows the average present discounted value of subsidies received for a 15-panel installation for each census region. The second panel gives the simulated number of solar installations per 1000 households in the model for each Census region. The third panel gives the total damages offset by rooftop solar and by the increased storage technology of renewable energy. The final panel gives total government cost under each subsidy scheme converted to an annuity value. “Unconstrained Optimal” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$). All monetary values are measured in 2014 dollars.

panel gives the percentage change in environmental benefits when moving from the current to welfare-maximizing subsidies. As before, we can see that moving to the welfare-maximizing cost-neutral subsidies given the baseline technology leads to an increase in environmental benefits of 6.2%.

The next three columns show the results when we recalculate welfare-maximizing subsidies given that the scale of utility-scale solar and wind expand based on three scenarios of projected renewable expansion by 2030 from the EIA (Nalley & LaRose, 2022). Specifically, we expand wind and solar based on their “reference case” projection, low-cost projection, and high cost-projection. The high-cost scenario is associated with the smallest increase in utility-scale solar and wind production, while the low-cost scenario is associated with the largest increases.¹⁴ We refer to their reference case projection as the mid-cost projection. Across the three scenarios, we find that moving to the welfare-maximizing cost-neutral subsidies leads to a 17-20% increase in aggregate environmental benefits.

In the fifth column, we recalculate results considering each coal plant to have “cleaned up” by adjusting marginal damages from coal plants so that the distribution of marginal damages from coal plants matches that of natural gas plants. Moving to the welfare-maximizing cost-neutral subsidies leads to a 11.6% increase in aggregate environmental benefits in this case.

Columns 6 through 10 repeat this exercise with unconstrained nationally-optimal subsidies. In all cases, current subsidies are overfunded relative to the optimum. Moving to unconstrained nationally-optimal subsidies involves cutting funding for subsidies by 55% to 65% across specifications.

¹⁴Specifically, utility-scale solar increases by roughly 200%, 350%, and 500% in the three scenarios, while wind increases by 45%, 50%, and 55%.

(1)	(2)	(3) Cost Neutral			(4)	(5)	(6)	(7)	(8) Unconstrained	(9)	(10)
	Current Tech	High Cost	Increased Renewables	Mid Cost	Low Cost	Clean Coal	Current Tech	High Cost	Increased Renewables	Mid Cost	Clean Coal
I. Δ Average Subsidy (\$Thousands)											
Midwest	5.4	6.7		6.9	7.1	3.5	1.9	1.6	1.5	1.3	-2.2
Northeast	-4.0	-3.0		-2.8	-2.7	-4.3	-7.2	-7.6	-7.8	-8.0	-9.9
South	3.9	4.9		5.0	5.0	3.9	0.6	0.2	-0.1	-0.5	-2.1
West	0.9	-1.1		-1.5	-1.7	1.6	-2.2	-5.0	-5.6	-6.0	-4.0
II. Δ Installations per 1000HHs											
Midwest	2.7	3.7		3.8	4.0	1.6	0.7	0.5	0.4	0.3	-1.1
Northeast	-9.1	-8.0		-7.9	-7.7	-9.5	-11.8	-12.0	-12.2	-12.3	-13.5
South	2.4	3.2		3.3	3.3	2.4	0.2	-0.1	-0.3	-0.4	-1.2
West	1.2	-1.4		-1.8	-1.9	2.5	-1.9	-4.4	-4.8	-5.0	-3.1
III. % Δ Environmental Benefits											
Total	11.6	17.4		19.9	20.5	11.6	-16.6	-22.8	-24.0	-25.9	-33.1
IV. % Δ Total Subsidies Paid											
Total	0	0		0	0	0	-42.6	-55.0	-58.0	-60.5	-64.8

Table C.20. Nationally-optimal cost-neutral and unconstrained subsidies under alternative assumptions about central generation energy production.

Each entry of Columns (1) through (4) shows the change of moving from the current system of subsidies to the welfare-maximizing cost-neutral system of subsidies given the specification in question. Each entry of Columns (5) through (8) shows the change of moving from the current system of subsidies to the unconstrained nationally-optimal system of subsidies given the specification in question. “Unconstrained” refers to the optimal unconstrained subsidies with the marginal value of public funds equal to one ($\lambda = 1$). All monetary values are measured in 2014 dollars. See text for details on each model specification.

REFERENCES CITED

- Albouy, D., Graf, W., Kellogg, R., & Wolff, H. (2016, March). Climate Amenities, Climate Change, and American Quality of Life. *Journal of the Association of Environmental and Resource Economists*, 3(1), 205–246. doi: 10.1086/684573
- Alexander, D., & Schwandt, H. (2022, February). The impact of car pollution on infant and child health: Evidence from emissions cheating. *The Review of Economic Studies*, 89(6), 2872–2910. Retrieved from <http://dx.doi.org/10.1093/restud/rdac007> doi: 10.1093/restud/rdac007
- Almond, D., Hoynes, H. W., & Schanzenbach, D. W. (2011, May). Inside the war on poverty: The impact of food stamps on birth outcomes. *Review of Economics and Statistics*, 93(2), 387–403. Retrieved from http://dx.doi.org/10.1162/REST_a_00089 doi: 10.1162/rest_a_00089
- Anderson, M. L. (2008, December). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American Statistical Association*, 103(484), 1481–1495. Retrieved from <http://dx.doi.org/10.1198/016214508000000841> doi: 10.1198/016214508000000841
- Anderson, S. T., Kellogg, R., & Saltee, J. M. (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management*, 66(3), 383–403.
- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Antier, C., Kudsk, P., Reboud, X., Ulber, L., Baret, P. V., & Messéan, A. (2020, January). Glyphosate Use in the European Agricultural Sector and a Framework for Its Further Monitoring. *Sustainability*, 12(14), 5682. doi: 10.3390/su12145682
- Arbuckle, T. E., Lin, Z., & Mery, L. S. (2001, August). An exploratory analysis of the effect of pesticide exposure on the risk of spontaneous abortion in an Ontario farm population. *Environmental Health Perspectives*, 109(8), 851–857. doi: 10.1289/ehp.01109851
- Arkolakis, C., & Walsh, C. (2022). Clean growth. *Working Paper*.

- Athey, S., & Wager, S. (2019). *Estimating treatment effects with causal forests: An application*. arXiv. Retrieved from <https://arxiv.org/abs/1902.07409>
doi: 10.48550/ARXIV.1902.07409
- Auffhammer, M. (2022, May). Climate Adaptive Response Estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption. *Journal of Environmental Economics and Management*, 102669. doi: 10.1016/j.jeem.2022.102669
- Auffhammer, M., & Mansur, E. T. (2014, November). Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics*, 46, 522–530. doi: 10.1016/j.eneco.2014.04.017
- Baer, KM., & Marcel, BJ. (2014). *Glyphosate* (3rd ed., Vol. 1). Academic Press.
- Baker, N. T., & Stone, W. W. (2015). *Estimated annual agricultural pesticide use for counties of the conterminous United States, 2008-12* (Tech. Rep. No. 907). U.S. Geological Survey. doi: 10.3133/ds907
- Ballard, C. L., Shoven, J. B., & Whalley, J. (1985). General equilibrium computations of the marginal welfare costs of taxes in the united states. *The American Economic Review*, 75(1), 128–138.
- Banzhaf, S., Ma, L., & Timmins, C. (2019a, February). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives*, 33(1), 185–208. doi: 10.1257/jep.33.1.185
- Banzhaf, S., Ma, L., & Timmins, C. (2019b, February). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives*, 33(1), 185–208. doi: 10.1257/jep.33.1.185
- Barbose, G., & Darghouth, N. (2019). *Tracking the Sun: Pricing and Design Trends for Distributed Photovoltaic Systems in the United States - 2019 Edition*.
- Barreca, A., Clay, K., Deschenes, O., Greenstone, M., & Shapiro, J. S. (2016, February). Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century. *Journal of Political Economy*, 124(1), 105–159. doi: 10.1086/684582
- Barrows, G., Sexton, S., & Zilberman, D. (2014). Agricultural Biotechnology: The Promise and Prospects of Genetically Modified Crops. *Journal of Economic Perspectives*, 28(1), 99–120. doi: 10.1257/jep.28.1.99
- Batini, N., Di Serio, M., Fragetta, M., Melina, G., & Waldron, A. (2022). Building back better: How big are green spending multipliers? *Ecological Economics*, 193, 107305.

- Battaglin, W., Meyer, M., Kuivila, K., & Dietze, J. (2014). Glyphosate and Its Degradation Product AMPA Occur Frequently and Widely in U.S. Soils, Surface Water, Groundwater, and Precipitation. *JAWRA Journal of the American Water Resources Association*, 50(2), 275–290. doi: 10.1111/jawr.12159
- Baum-Snow, N., & Pavan, R. (2012, January). Understanding the City Size Wage Gap. *The Review of Economic Studies*, 79(1), 88–127. doi: 10.1093/restud/rdr022
- Bayer, P., Ferreira, F., & McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4), 588–638.
- Bayer, P., Keohane, N., & Timmins, C. (2009, July). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, 58(1), 1–14. doi: 10.1016/j.jeem.2008.08.004
- Beam, A. L., Fried, I., Palmer, N., Agniel, D., Brat, G., Fox, K., ... Armstrong, J. (2020, February). Estimates of healthcare spending for preterm and low-birthweight infants in a commercially insured population: 2008–2016. *Journal of Perinatology*, 40(7), 1091–1099. Retrieved from <http://dx.doi.org/10.1038/s41372-020-0635-z> doi: 10.1038/s41372-020-0635-z
- Benachour, N., Sipahutar, H., Moslemi, S., Gasnier, C., Travert, C., & Séralini, G. E. (2007, July). Time- and Dose-Dependent Effects of Roundup on Human Embryonic and Placental Cells. *Archives of Environmental Contamination and Toxicology*, 53(1), 126–133. doi: 10.1007/s00244-006-0154-8
- Benbrook, C. (2020, April). Shining a light on glyphosate-based herbicide hazard, exposures and risk: Role of non-hodgkin lymphoma litigation in the usa. *European Journal of Risk Regulation*, 11(3), 498–519. Retrieved from <http://dx.doi.org/10.1017/err.2020.16> doi: 10.1017/err.2020.16
- Benbrook, C. M. (2018, July). Why regulators lost track and control of pesticide risks: Lessons from the case of glyphosate-based herbicides and genetically engineered-crop technology. *Current Environmental Health Reports*, 5(3), 387–395. Retrieved from <http://dx.doi.org/10.1007/s40572-018-0207-y> doi: 10.1007/s40572-018-0207-y

- Benbrook, C. M. (2019, January). How did the US EPA and IARC reach diametrically opposed conclusions on the genotoxicity of glyphosate-based herbicides? *Environmental Sciences Europe*, 31(1). Retrieved from <http://dx.doi.org/10.1186/s12302-018-0184-7> doi: 10.1186/s12302-018-0184-7
- Bentham, A. v., Gillingham, K., & Sweeney, J. (2008). Learning-by-doing and the optimal solar policy in California. *The Energy Journal*, 29(3), 131–152.
- Berry, S., Levinsohn, J., & Pakes, A. (2004, February). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy*, 112(1), 68–105. doi: 10.1086/379939
- Bharadwaj, P., Lundborg, P., & Rooth, D.-O. (2017, January). Birth weight in the long run. *Journal of Human Resources*, 53(1), 189–231. Retrieved from <http://dx.doi.org/10.3368/jhr.53.1.0715-7235R> doi: 10.3368/jhr.53.1.0715-7235r
- Black, S. E. (1999). Do better schools matter? parental valuation of elementary education. *The Quarterly Journal of Economics*, 114(2), 577–599.
- Blair, N. J., Dobos, A. P., & Gilman, P. (2013, August). *Comparison of Photovoltaic Models in the System Advisor Model: Preprint* (Tech. Rep. No. NREL/CP-6A20-58057). National Renewable Energy Lab. (NREL), Golden, CO (United States). Retrieved 2022-11-17, from <https://www.osti.gov/biblio/1096104>
- Bollinger, B., & Gillingham, K. (2012). Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6), 900–912.
- Bollinger, B., & Gillingham, K. (2019). Learning-by-doing in solar photovoltaic installations. *Available at SSRN 2342406*.
- Bollinger, B., Gillingham, K., Kirkpatrick, A. J., & Sexton, S. (2022). Visibility and peer influence in durable good adoption. *Marketing Science*.
- Borenstein, S. (2012). The private and public economics of renewable electricity generation. *Journal of Economic Perspectives*, 26(1), 67–92.
- Borenstein, S. (2017). Private net benefits of residential solar pv: The role of electricity tariffs, tax incentives, and rebates. *Journal of the Association of Environmental and Resource Economists*, 4(S1), S85–S122.

- Borenstein, S., & Bushnell, J. B. (2022, November). Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency. *American Economic Journal: Economic Policy*, 14(4), 80–110. Retrieved 2022-12-19, from <https://www.aeaweb.org/articles?id=10.1257/pol.20190758> doi: 10.1257/pol.20190758
- Borenstein, S., & Davis, L. W. (2016). The distributional effects of us clean energy tax credits. *Tax Policy and the Economy*, 30(1), 191–234.
- Bound, J., & Holzer, H. J. (2000, January). Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s. *Journal of Labor Economics*, 18(1), 20–54. doi: 10.1086/209949
- Brown, P. R., & O’Sullivan, F. M. (2020). Spatial and temporal variation in the value of solar power across united states electricity markets. *Renewable and Sustainable Energy Reviews*, 121, 109594.
- Burke, P. J., & Abayasekara, A. (2018). The Price Elasticity of Electricity Demand in the United States: A Three-Dimensional Analysis. *The Energy Journal*, 39(2), 123–146.
- Burlig, F., Knittel, C., Rapson, D., Reguant, M., & Wolfram, C. (2020, November). Machine learning from schools about energy efficiency. *Journal of the Association of Environmental and Resource Economists*, 7(6), 1181–1217. Retrieved from <http://dx.doi.org/10.1086/710606> doi: 10.1086/710606
- Burr, C. (2014). Subsidies, tariffs and investments in the solar power market. *University of Colorado-Boulder Working Paper*.
- Busso, M., Gregory, J., & Kline, P. (2013, April). Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review*, 103(2), 897–947. doi: 10.1257/aer.103.2.897
- Butters, R. A., Dorsey, J., & Gowrisankaran, G. (2021, August). *Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium* [Working Paper]. National Bureau of Economic Research. Retrieved 2022-12-17, from <https://www.nber.org/papers/w29133> doi: 10.3386/w29133
- Cain, L., Hernandez-Cortes, D., Timmins, C., & Weber, P. (2023). *Recent Findings and Methodologies in Economics Research in Environmental Justice* (SSRN Scholarly Paper No. 4368212). Rochester, NY. doi: 10.2139/ssrn.4368212

- Callaway, D. S., Fowlie, M., & McCormick, G. (2018). Location, location, location: The variable value of renewable energy and demand-side efficiency resources. *Journal of the Association of Environmental and Resource Economists*, 5(1), 39–75.
- Calzada, J., Gisbert, M., & Moscoso, B. (2023, November). The hidden cost of bananas: The effects of pesticides on newborns' health. *Journal of the Association of Environmental and Resource Economists*, 10(6), 1623–1663. Retrieved from <http://dx.doi.org/10.1086/725349> doi: 10.1086/725349
- Camacho, A., & Mejia, D. (2017, July). The health consequences of aerial spraying illicit crops: The case of Colombia. *Journal of Health Economics*, 54, 147–160. doi: 10.1016/j.jhealeco.2017.04.005
- Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.
- Card, D. (2009, May). Immigration and Inequality. *American Economic Review*, 99(2), 1–21. doi: 10.1257/aer.99.2.1
- Chowdhury, M. S., Rahman, K. S., Chowdhury, T., Nuthammachot, N., Techato, K., Akhtaruzzaman, M., ... Amin, N. (2020). An overview of solar photovoltaic panels' end-of-life material recycling. *Energy Strategy Reviews*, 27, 100431.
- Christensen, P., Sarmiento-Barbieri, I., & Timmins, C. (2021, November). *Racial Discrimination and Housing Outcomes in the United States Rental Market* (Working Paper No. 29516). National Bureau of Economic Research. doi: 10.3386/w29516
- Christensen, P., & Timmins, C. (2021, July). *The Damages and Distortions from Discrimination in the Rental Housing Market* (Working Paper No. 29049). National Bureau of Economic Research. doi: 10.3386/w29049
- Cicala, S., Holland, S. P., Mansur, E. T., Muller, N. Z., & Yates, A. J. (2021). Expected health effects of reduced air pollution from covid-19 social distancing. *Atmosphere*, 12(8), 951.
- CIESIN. (2017). *U.S. Census Grids (Summary File 1), 2010*. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC). doi: 10.7927/H40Z716C

- Clausing, P., Robinson, C., & Burtscher-Schaden, H. (2018, March). Pesticides and public health: an analysis of the regulatory approach to assessing the carcinogenicity of glyphosate in the european union. *Journal of Epidemiology and Community Health*, 72(8), 668–672. Retrieved from <http://dx.doi.org/10.1136/jech-2017-209776> doi: 10.1136/jech-2017-209776
- Colas, M., Findeisen, S., & Sachs, D. (2021). Optimal need-based financial aid. *Journal of Political Economy*, 129(2), 492–533.
- Colas, M., & Hutchinson, K. (2021, May). Heterogeneous Workers and Federal Income Taxes in a Spatial Equilibrium. *American Economic Journal: Economic Policy*, 13(2), 100–134. doi: 10.1257/pol.20180529
- Colas, M., & Morehouse, J. M. (2022, January). The environmental cost of land-use restrictions. *Quantitative Economics*, 13(1), 179–223. doi: 10.3982/TE1612
- Colas, M., & Reynier, E. (2024). Means-tested solar subsidies. *Working Paper*.
- Colas, M., & Saulnier, E. (2023, July). Vertical migration externalities. *Regional Science and Urban Economics*, 101, 103900. doi: 10.1016/j.regsciurbeco.2023.103900
- Crago, C. L., & Chernyakhovskiy, I. (2017). Are policy incentives for solar power effective? evidence from residential installations in the northeast. *Journal of Environmental Economics and Management*, 81, 132–151.
- Cruz, J.-L., & Rossi-Hansberg, E. (2023, April). The Economic Geography of Global Warming. *The Review of Economic Studies*, rdad042. doi: 10.1093/restud/rdad042
- Currie, J., Graff Zivin, J., Meckel, K., Neidell, M., & Schlenker, W. (2013, August). Something in the water: Contaminated drinking water and infant health. *Canadian Journal of Economics*, 46(3), 791–810. Retrieved from <http://dx.doi.org/10.1111/caje.12039> doi: 10.1111/caje.12039
- Currie, J., Greenstone, M., & Meckel, K. (2017, December). Hydraulic fracturing and infant health: New evidence from Pennsylvania. *Science Advances*, 3(12), e1603021. doi: 10.1126/sciadv.1603021
- Currie, J., Neidell, M., & Schmieder, J. F. (2009, May). Air pollution and infant health: Lessons from new jersey. *Journal of Health Economics*, 28(3), 688–703. Retrieved from <http://dx.doi.org/10.1016/j.jhealeco.2009.02.001> doi: 10.1016/j.jhealeco.2009.02.001

- Currie, J., & Walker, R. (2011, January). Traffic congestion and infant health: Evidence from e-zpass. *American Economic Journal: Applied Economics*, 3(1), 65–90. Retrieved from <http://dx.doi.org/10.1257/app.3.1.65> doi: 10.1257/app.3.1.65
- Dallegrave, E., Mantese, F. D., Coelho, R. S., Pereira, J. D., Dalsenter, P. R., & Langeloh, A. (2003, April). The teratogenic potential of the herbicide glyphosate-Roundup® in Wistar rats. *Toxicology Letters*, 142(1), 45–52. doi: 10.1016/S0378-4274(02)00483-6
- Data Sets Not Requiring IRB Review*. (2024). <https://www.nber.org/programs-projects/projects-and-centers/human-subjects-protection-and-institutional-review-board-irb/guidance-data-sets-not-requiring-irb-review>. (NBER)
- Dauwalter, T. E., & Harris, R. I. (2023). Distributional benefits of rooftop solar capacity. *Journal of the Association of Environmental and Resource Economists*, 10(2), 487–523.
- Davis, L. W., & Gertler, P. J. (2015, May). Contribution of air conditioning adoption to future energy use under global warming. *Proceedings of the National Academy of Sciences*, 112(19), 5962–5967. doi: 10.1073/pnas.1423558112
- Davoren, M. J., & Schiestl, R. H. (2018, July). Glyphosate-based herbicides and cancer risk: a post-iarc decision review of potential mechanisms, policy and avenues of research. *Carcinogenesis*, 39(10), 1207–1215. Retrieved from <http://dx.doi.org/10.1093/carcin/bgy105> doi: 10.1093/carcin/bgy105
- De Groote, O., & Verboven, F. (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review*, 109(6), 2137–72.
- Depro, B., Timmins, C., & O’Neil, M. (2015, September). White Flight and Coming to the Nuisance: Can Residential Mobility Explain Environmental Injustice? *Journal of the Association of Environmental and Resource Economists*, 2(3), 439–468. doi: 10.1086/682716
- Desmet, K., Kopp, R. E., Kulp, S. A., Nagy, D. K., Oppenheimer, M., Rossi-Hansberg, E., & Strauss, B. H. (2021, April). Evaluating the Economic Cost of Coastal Flooding. *American Economic Journal: Macroeconomics*, 13(2), 444–486. doi: 10.1257/mac.20180366

- Desmet, K., & Rossi-Hansberg, E. (2015, July). On the spatial economic impact of global warming. *Journal of Urban Economics*, 88, 16–37. doi: 10.1016/j.jue.2015.04.004
- Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *The American Economic Review*, 106(3), 479–524.
- Dias, M., Rocha, R., & Soares, R. R. (2023, February). Down the River: Glyphosate Use in Agriculture and Birth Outcomes of Surrounding Populations. *The Review of Economic Studies*, 90, 2943–2981. doi: 10.1093/restud/rdad011
- Documentation. (2023). <https://www.ers.usda.gov/data-products/rural-urban-continuum-codes/documentation>. (USDA ERS)
- Doremus, J. M., Jacqz, I., & Johnston, S. (2022, March). Sweating the energy bill: Extreme weather, poor households, and the energy spending gap. *Journal of Environmental Economics and Management*, 112, 102609. doi: 10.1016/j.jeem.2022.102609
- Dorsey, J. (2024). Solar market frictions: The role of platforms and policies. *Working Paper*.
- Duke, S. O., & Powles, S. B. (2008). Glyphosate: A once-in-a-century herbicide. *Pest Management Science*, 64(4), 319–325. doi: 10.1002/ps.1518
- Eichner, T., & Runkel, M. (2014). Subsidizing renewable energy under capital mobility. *Journal of Public Economics*, 117, 50–59.
- Elwert, F., & Winship, C. (2014, July). Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40(1), 31-53. Retrieved from <http://dx.doi.org/10.1146/annurev-soc-071913-043455> doi: 10.1146/annurev-soc-071913-043455
- Engist, D., Guzman, L. M., Larsen, A., Church, T., & Noack, F. (2024, July). The impact of genetically modified crops on bird diversity. *Nature Sustainability*, 7(9), 1149–1159. Retrieved from <http://dx.doi.org/10.1038/s41893-024-01390-y> doi: 10.1038/s41893-024-01390-y
- Executive Office of the President. (2021, January). *Executive Order 14008: Tackling the Climate Crisis at Home and Abroad* (Regulation No. 14008).

- Falcone, J. A. (2021). *Estimates of county-level nitrogen and phosphorus from fertilizer and manure from 1950 through 2017 in the conterminous United States* (Tech. Rep. No. 2020-1153). U.S. Geological Survey. doi: 10.3133/ofr20201153
- FAO, & IIASA. (2021). *Global Agro Ecological Zones version 4 (GAEZ v4)*. <http://www.fao.org/gaez/en>.
- Feger, F., Pavanini, N., & Radulescu, D. (2022). Welfare and redistribution in residential electricity markets with solar power. *The Review of Economic Studies*, 89(6), 3267–3302.
- Finkelstein, A., & Hendren, N. (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives*, 34(4), 146–167.
- Finkelstein, A., & McKnight, R. (2008). What did medicare do? the initial impact of medicare on mortality and out of pocket medical spending. *Journal of public economics*, 92(7), 1644–1668.
- First Street Foundation. (2022a, April). First Street Aggregated Flood Risk Summary Statistics Version 2.0.
doi: 10.5281/zenodo.6498206
- First Street Foundation. (2022b, May). First Street Foundation Aggregated Wildfire Risk Summary Statistics Version 1.0.
doi: 10.5281/zenodo.6564731
- Fletcher, J., & Noghanibehambari, H. (2024, January). The siren song of cicadas: Early-life pesticide exposure and later-life male mortality. *Journal of Environmental Economics and Management*, 123, 102903. doi: 10.1016/j.jeem.2023.102903
- Foucault, A., Vallet, N., Ravalet, N., Picou, F., Bene, M. C., Gyan, E., & Herault, O. (2021). Occupational pesticide exposure increases risk of acute myeloid leukemia: a meta-analysis of case-control studies including 3,955 cases and 9,948 controls. *Scientific Reports*, 11(1), 2007.
- Frank, E., & Taylor, C. A. (2022). *The “Golden Age” of Pesticides? Trade-offs of DDT and Health in the US*. (Unpublished)
- Frank, E. G. (2024, September). The economic impacts of ecosystem disruptions: Costs from substituting biological pest control. *Science*, 385(6713). Retrieved from <http://dx.doi.org/10.1126/science.adg0344> doi: 10.1126/science.adg0344

- Gerarden, T. D. (2023). Demanding innovation: The impact of consumer subsidies on solar panel production costs. *Management Science*, 69(12), 7799–7820.
- Gergel, D. R., Malevich, S. B., McCusker, K. E., Tenezakis, E., Delgado, M. T., Fish, M. A., & Kopp, R. E. (2023, January). Global downscaled projections for climate impacts research (GDPCIR): Preserving extremes for modeling future climate impacts. *EGUsphere*, 1–35. doi: 10.5194/egusphere-2022-1513
- Gerona, R. R., Reiter, J. L., Zakharevich, I., Proctor, C., Ying, J., Mesnage, R., . . . Winchester, P. D. (2022, October). Glyphosate exposure in early pregnancy and reduced fetal growth: A prospective observational study of high-risk pregnancies. *Environmental Health*, 21(1), 95. doi: 10.1186/s12940-022-00906-3
- Gillezeau, C., van Gerwen, M., Shaffer, R. M., Rana, I., Zhang, L., Sheppard, L., & Taioli, E. (2019, January). The evidence of human exposure to glyphosate: a review. *Environmental Health*, 18(1). Retrieved from <http://dx.doi.org/10.1186/s12940-018-0435-5> doi: 10.1186/s12940-018-0435-5
- Gillingham, K., Deng, H., Wiser, R., Darghouth, N., Nemet, G., Barbose, G., . . . Dong, C. (2016). Deconstructing solar photovoltaic pricing: the role of market structure, technology, and policy. *The Energy Journal*, 37(3), 231–250.
- Gillingham, K., & Tsvetanov, T. (2019). Hurdles and steps: Estimating demand for solar photovoltaics. *Quantitative Economics*, 10(1), 275–310.
- Glaeser, E. L., & Kahn, M. E. (2010, May). The greenness of cities: Carbon dioxide emissions and urban development. *Journal of Urban Economics*, 67(3), 404–418. doi: 10.1016/j.jue.2009.11.006
- Goulder, L. H., & Mathai, K. (2010). Optimal co2 abatement in the presence of induced technological change. In *Technological change and the environment* (pp. 210–250). Routledge.
- Graff Zivin, J. S., Kotchen, M. J., & Mansur, E. T. (2014). Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *Journal of Economic Behavior & Organization*, 107, 248–268.
- Grube, A., Donaldson, D., Kiely, T., & Wu, L. (2011, February). *Pesticides Industry Sales and Usage: 2006 and 2007 Market Estimates* (Tech. Rep.). U.S. EPA.

- Hamilton, T. L., & Phaneuf, D. J. (2015, November). An integrated model of regional and local residential sorting with application to air quality. *Journal of Environmental Economics and Management*, 74, 71–93. doi: 10.1016/j.jeem.2015.08.001
- Hansen, C. W., & Wingender, A. M. (2023, June). National and Global Impacts of Genetically Modified Crops. *American Economic Review: Insights*, 5(2), 224–240. doi: 10.1257/aeri.20220144
- Hartmann, A., Jasechko, S., Gleeson, T., Wada, Y., Andreo, B., Barberá, J. A., . . . Wagener, T. (2021, May). Risk of groundwater contamination widely underestimated because of fast flow into aquifers. *Proceedings of the National Academy of Sciences*, 118(20), e2024492118. doi: 10.1073/pnas.2024492118
- Hasna, Z. (2021). The grass is actually greener on the other side: Evidence on green multipliers from the united states. *Job Market Paper*.
- Hausman, C., & Stolper, S. (2021, October). Inequality, information failures, and air pollution. *Journal of Environmental Economics and Management*, 110, 102552. doi: 10.1016/j.jeem.2021.102552
- Heath, G. A., & Mann, M. K. (2012). Background and reflections on the life cycle assessment harmonization project. *Journal of Industrial Ecology*, 16(NREL/JA-6A20-54177).
- Heckman, J. J., Moon, S. H., Pinto, R., Savelyev, P. A., & Yavitz, A. (2010). The rate of return to the highscope perry preschool program. *Journal of public Economics*, 94(1-2), 114–128.
- Heijmans, B. T., Tobi, E. W., Stein, A. D., Putter, H., Blauw, G. J., Susser, E. S., . . . Lumey, L. (2008). Persistent epigenetic differences associated with prenatal exposure to famine in humans. *Proceedings of the National Academy of Sciences*, 105(44), 17046–17049.
- Heilmann, K., Kahn, M. E., & Tang, C. K. (2021, May). The urban crime and heat gradient in high and low poverty areas. *Journal of Public Economics*, 197, 104408. doi: 10.1016/j.jpubeco.2021.104408
- Holland, S. P., Kotchen, M. J., Mansur, E. T., & Yates, A. J. (2022, February). Why marginal CO2 emissions are not decreasing for US electricity: Estimates and implications for climate policy. *Proceedings of the National Academy of Sciences*, 119(8), e2116632119. doi: 10.1073/pnas.2116632119
- Holland, S. P., & Mansur, E. T. (2008). Is real-time pricing green? the environmental impacts of electricity demand variance. *The Review of Economics and Statistics*, 90(3), 550–561.

- Holland, S. P., Mansur, E. T., Muller, N. Z., & Yates, A. J. (2016). Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review*, 106(12), 3700–3729.
- Holland, S. P., Mansur, E. T., Muller, N. Z., & Yates, A. J. (2020). Decompositions and policy consequences of an extraordinary decline in air pollution from electricity generation. *American Economic Journal: Economic Policy*, 12(4), 244–74.
- Holland, S. P., Mansur, E. T., Muller, N. Z., & Yates, A. J. (2021). The environmental benefits of transportation electrification: Urban buses. *Energy policy*, 148, 111921.
- Holland, S. P., Mansur, E. T., & Yates, A. J. (2022). *Decarbonization and electrification in the long run* (Tech. Rep.). National Bureau of Economic Research.
- Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., ... Houser, T. (2017, June). Estimating economic damage from climate change in the United States. *Science*, 356(6345), 1362–1369. doi: 10.1126/science.aal4369
- Hu, Y., Wu, S., Lyu, W., Ning, J., & She, D. (2023). Risk assessment of human exposure to airborne pesticides in rural greenhouses. *Scientific Reports*, 13(1), 5138.
- Hughes, J. E., & Podolefsky, M. (2015). Getting green with solar subsidies: evidence from the california solar initiative. *Journal of the Association of Environmental and Resource Economists*, 2(2), 235–275.
- Ito, K. (2014). Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2), 537–563.
- Jacobs, B. (2018). The marginal cost of public funds is one at the optimal tax system. *International Tax and Public Finance*, 25, 883–912.
- Jones, B. A. (2020, March). Invasive Species Control, Agricultural Pesticide Use, and Infant Health Outcomes. *Land Economics*, 96(2), 149–170. doi: 10.3368/le.96.2.149
- Jordan, D. C., & Kurtz, S. R. (2013). Photovoltaic degradation rates—an analytical review. *Progress in photovoltaics: Research and Applications*, 21(1), 12–29.
- Katz, L. F., & Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1), 35–78. doi: 10.2307/2118323

- Kline, P., & Moretti, E. (2014). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, 6(1), 629–662. doi: 10.1146/annurev-economics-080213-041024
- Kniss, A. R. (2017, April). Long-term trends in the intensity and relative toxicity of herbicide use. *Nature Communications*, 8(1), 14865. doi: 10.1038/ncomms14865
- Lamp, S., & Samano, M. (2023). (mis)allocation of renewable energy sources. *Journal of the Association of Environmental and Resource Economists*, 10(1), 195–229.
- Landrigan, P. J. (2018, January). Pesticides and Human Reproduction. *JAMA internal medicine*, 178(1), 26–27. doi: 10.1001/jamainternmed.2017.5092
- Langer, A., & Lemoine, D. (2022). Designing dynamic subsidies to spur adoption of new technologies. *Journal of the Association of Environmental and Resource Economists*, 9(6), 1197–1234.
- Lark, T. J., Hendricks, N. P., Smith, A., Pates, N., Spawn-Lee, S. A., Bougie, M., ... Gibbs, H. K. (2022, March). Environmental outcomes of the US Renewable Fuel Standard. *Proceedings of the National Academy of Sciences*, 119(9), e2101084119. doi: 10.1073/pnas.2101084119
- Larsen, A. E., Gaines, S. D., & Deschênes, O. (2017, August). Agricultural pesticide use and adverse birth outcomes in the san joaquin valley of california. *Nature Communications*, 8(1). Retrieved from <http://dx.doi.org/10.1038/s41467-017-00349-2> doi: 10.1038/s41467-017-00349-2
- Lazard. (2023, April). *2023 Levelized Cost Of Energy+* (Tech. Rep.). Lazard. Retrieved 2023-08-16, from <https://www.lazard.com/research-insights/2023-levelized-cost-of-energyplus/>
- Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2), 281–355.
- Lee, S., Moschini, G., & Perry, E. D. (2023, September). Genetically engineered varieties and applied pesticide toxicity in U.S. maize and soybeans: Heterogeneous and evolving impacts. *Ecological Economics*, 211, 107873. doi: 10.1016/j.ecolecon.2023.107873
- Lehner, B., & Grill, G. (2013). Global river hydrography and network routing: Baseline data and new approaches to study the world’s large river systems. *Hydrological Processes*, 27(15), 2171–2186. doi: 10.1002/hyp.9740

- Leon, M. E., Schinasi, L. H., Lebailly, P., Beane Freeman, L. E., Nordby, K.-C., Ferro, G., ... Schüz, J. (2019, October). Pesticide use and risk of non-Hodgkin lymphoid malignancies in agricultural cohorts from France, Norway and the USA: A pooled analysis from the AGRICOH consortium. *International Journal of Epidemiology*, 48(5), 1519–1535. doi: 10.1093/ije/dyz017
- Lusk, J. L., Tack, J., & Hendricks, N. P. (2017, June). *Heterogeneous Yield Impacts from Adoption of Genetically Engineered Corn and the Importance of Controlling for Weather* (Working Paper No. 23519). National Bureau of Economic Research. doi: 10.3386/w23519
- Lyubich, E. (2020, June). The Race Gap in Residential Energy Expenditures. *Berkeley Energy Institute Working Paper Series*, 306, 15.
- Marcus, M. (2021, January). Going beneath the surface: Petroleum pollution, regulation, and health. *American Economic Journal: Applied Economics*, 13(1), 72–104. Retrieved from <http://dx.doi.org/10.1257/app.20190130> doi: 10.1257/app.20190130
- Medalie, L., Baker, N. T., Shoda, M. E., Stone, W. W., Meyer, M. T., Stets, E. G., & Wilson, M. (2020, March). Influence of land use and region on glyphosate and aminomethylphosphonic acid in streams in the USA. *Science of The Total Environment*, 707, 136008. doi: 10.1016/j.scitotenv.2019.136008
- Menne, M. J., Durre, I., Korzeniewski, B., McNeill, S., Thomas, K., Yin, X., ... Houston, T. G. (2012). *Global Historical Climatology Network - Daily (GHCN-Daily), Version 3*. NOAA National Climatic Data Center. doi: 10.7289/V5D21VHZ
- Mie, A., Rudén, C., & Grandjean, P. (2018, November). Safety of safety evaluation of pesticides: developmental neurotoxicity of chlorpyrifos and chlorpyrifos-methyl. *Environmental Health*, 17(1). Retrieved from <http://dx.doi.org/10.1186/s12940-018-0421-y> doi: 10.1186/s12940-018-0421-y
- Miller, G., Pease, G., Knowles, M., & Shi, W. (2022, October). *Open Grid Emissions Initiative (v0.3.0)*. Zenodo. Retrieved 2022-11-17, from <https://zenodo.org/record/7255769> doi: 10.5281/zenodo.7255769
- Millstein, D., Wiser, R., Bolinger, M., & Barbose, G. (2017). The climate and air-quality benefits of wind and solar power in the united states. *Nature Energy*, 2(9), 1–10.
- Morehouse, J. M. (2022). Carbon Taxes in Spatial Equilibrium. , 52.

- Muller, B. (2020). Glyphosate-A love story. Ordinary thoughtlessness and responseability in industrial farming. *Journal of Agrarian Change*, 21(1), 160–179. doi: 10.1111/joac.12374
- Muller, N. Z., Mendelsohn, R., & Nordhaus, W. (2011). Environmental accounting for pollution in the united states economy. *American Economic Review*, 101(5), 1649–75.
- Muñoz, J. P., Bleak, T. C., & Calaf, G. M. (2021, May). Glyphosate and the key characteristics of an endocrine disruptor: A review. *Chemosphere*, 270, 128619. Retrieved from <http://dx.doi.org/10.1016/j.chemosphere.2020.128619> doi: 10.1016/j.chemosphere.2020.128619
- Nalley, S., & LaRose, A. (2022). Annual energy outlook 2022 (aeo2022). *Energy Information Agency*, 23.
- National Center for Health Statistics. (1990–2013). *Restricted-use vital statistics data*. (As compiled from data provided by the 57 statistics jurisdictions through the VSCP)
- Newhouse, J. P. (1992). Medical care costs: how much welfare loss? *Journal of Economic perspectives*, 6(3), 3–21.
- Olken, B. A. (2007). Monitoring corruption: evidence from a field experiment in indonesia. *Journal of political Economy*, 115(2), 200–249.
- Ospina, M., Schütze, A., Morales-Agudelo, P., Vidal, M., Wong, L.-Y., & Calafat, A. M. (2022, December). Exposure to glyphosate in the United States: Data from the 2013–2014 National Health and Nutrition Examination Survey. *Environment International*, 170, 107620. doi: 10.1016/j.envint.2022.107620
- Paganelli, A., Gnazzo, V., Acosta, H., López, S. L., & Carrasco, A. E. (2010, October). Glyphosate-Based Herbicides Produce Teratogenic Effects on Vertebrates by Impairing Retinoic Acid Signaling. *Chemical Research in Toxicology*, 23(10), 1586–1595. doi: 10.1021/tx1001749
- Pappas, C., & Foos, B. (2023). Pesticide data program: 30 years of food residue data and trends. *Journal of Exposure Science & Environmental Epidemiology*, 33(5), 805–812.
- Parvez, S., Gerona, R. R., Proctor, C., Friesen, M., Ashby, J. L., Reiter, J. L., . . . Winchester, P. D. (2018, March). Glyphosate exposure in pregnancy and shortened gestational length: A prospective Indiana birth cohort study. *Environmental Health*, 17(1), 23. doi: 10.1186/s12940-018-0367-0

- Pham, T. H., Derian, L., Kervarrec, C., Kernanec, P.-Y., Jégou, B., Smagulova, F., & Gely-Pernot, A. (2019, May). Perinatal Exposure to Glyphosate and a Glyphosate-Based Herbicide Affect Spermatogenesis in Mice. *Toxicological Sciences*, 169(1), 260–271. doi: 10.1093/toxsci/kfz039
- Pigou, A. C. (1920). The economics of welfare. *London, United Kingdom*.
- Piyapromdee, S. (2021, February). The Impact of Immigration on Wages, Internal Migration, and Welfare. *The Review of Economic Studies*, 88(1), 406–453. doi: 10.1093/restud/rdaa029
- Pless, J., Langheim, R., Machak, C., Hellow, H., & Sigrin, B. (2017). *The price-concentration relationship in early residential solar third-party markets* (Tech. Rep.). National Renewable Energy Lab.(NREL), Golden, CO (United States).
- Pless, J., & Van Benthem, A. A. (2019). Pass-through as a test for market power: An application to solar subsidies. *American Economic Journal: Applied Economics*, 11(4), 367–401.
- Popp, D., Vona, F., Marin, G., & Chen, Z. (2020). *The employment impact of green fiscal push: evidence from the american recovery act* (Tech. Rep.). National Bureau of Economic Research.
- Portier, C. J., Armstrong, B. K., Baguley, B. C., Baur, X., Belyaev, I., Bellé, R., . . . Zhou, S.-F. (2016, March). Differences in the carcinogenic evaluation of glyphosate between the international agency for research on cancer (iarc) and the european food safety authority (efsa). *Journal of Epidemiology and Community Health*, 70(8), 741–745. Retrieved from <http://dx.doi.org/10.1136/jech-2015-207005> doi: 10.1136/jech-2015-207005
- Poterba, J. M. (1996). Government intervention in the markets for education and health care: how and why? In *Individual and social responsibility: Child care, education, medical care, and long-term care in america* (pp. 277–308). University of Chicago Press.
- PRISM Climate Group. (2022). *PRISM Gridded Climate Data*. Oregon State University.
- Quigley, J. M., & Rubinfeld, D. L. (1989). Unobservables in Consumer Choice: Residential Energy and the Demand for Comfort. *The Review of Economics and Statistics*, 71(3), 416–425. doi: 10.2307/1926898

- Ramirez Haberkon, N. B., Aparicio, V. C., & Mendez, M. J. (2021, June). First evidence of glyphosate and aminomethylphosphonic acid (AMPA) in the respirable dust (PM10) emitted from unpaved rural roads of Argentina. *Science of The Total Environment*, 773, 145055. doi: 10.1016/j.scitotenv.2021.145055
- Reames, T. G. (2016, October). Targeting energy justice: Exploring spatial, racial/ethnic and socioeconomic disparities in urban residential heating energy efficiency. *Energy Policy*, 97, 549–558. doi: 10.1016/j.enpol.2016.07.048
- Reguant, M. (2019). The efficiency and sectoral distributional impacts of large-scale renewable energy policies. *Journal of the Association of Environmental and Resource Economists*, 6(S1), S129–S168.
- Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., . . . others (2022). Comprehensive evidence implies a higher social cost of co2. *Nature*, 610(7933), 687–692.
- Rhode, P. (2024, June). How Suitable are FAO-GAEZ Crop Suitability Indices for Historical Analysis? In *SI 2024 Development of the American Economy*. Boston, MA: National Bureau of Economic Research.
- Roback, J. (1982, December). Wages, Rents, and the Quality of Life. *Journal of Political Economy*, 90(6), 1257–1278. doi: 10.1086/261120
- Roca, J. D. L., & Puga, D. (2017, January). Learning by Working in Big Cities. *The Review of Economic Studies*, 84(1), 106–142. doi: 10.1093/restud/rdw031
- Rode, A., Carleton, T., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., . . . Yuan, J. (2021, October). Estimating a social cost of carbon for global energy consumption. *Nature*, 598(7880), 308–314. doi: 10.1038/s41586-021-03883-8
- Rosen, S. (1979). Wage-based indexes of urban quality of life. *Current issues in urban economics*, 74–104.
- Rubin, E., & Auffhammer, M. (2023). Quantifying heterogeneity in the price elasticity of residential natural gas.
- Rudik, I., Lyn, G., Tan, W., & Ortiz-Bobea, A. (2021, June). *The Economic Effects of Climate Change in Dynamic Spatial Equilibrium*. doi: 10.31235/osf.io/usghb
- Ruggles, S., Flood, S., Goeken, R., Schouweiler, M., & Sobek, M. (2022). *IPUMS USA: Version 12.0*. Minneapolis, MN: IPUMS. doi: 10.18128/D010.V12.0

- Saez, E., & Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1), 24–45.
- Saiz, A. (2010, August). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics*, 125(3), 1253–1296. doi: 10.1162/qjec.2010.125.3.1253
- Schlenker, W., & Roberts, M. J. (2009, September). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences of the United States of America*, 106(37), 15594–15598. doi: 10.1073/pnas.0906865106
- Schulz, R., Bub, S., Petschick, L. L., Stehle, S., & Wolfram, J. (2021, April). Applied pesticide toxicity shifts toward plants and invertebrates, even in GM crops. *Science*, 372(6537), 81–84. doi: 10.1126/science.abe1148
- Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., & Shelby, J. (2018, June). The National Solar Radiation Data Base (NSRDB). *Renewable and Sustainable Energy Reviews*, 89, 51–60. Retrieved 2022-11-17, from <https://www.sciencedirect.com/science/article/pii/S136403211830087X> doi: 10.1016/j.rser.2018.03.003
- Sexton, S., Kirkpatrick, A. J., Harris, R. I., & Muller, N. Z. (2021). Heterogeneous solar capacity benefits, appropriability, and the costs of suboptimal siting. *Journal of the Association of Environmental and Resource Economists*, 8(6), 1209–1244.
- Shapiro, J. S., & Walker, R. (2020). *Is air pollution regulation too stringent?* (Tech. Rep.). National Bureau of Economic Research.
- Sheppard, L., McGrew, S., & Fenske, R. A. (2020, October). Flawed analysis of an intentional human dosing study and its impact on chlorpyrifos risk assessments. *Environment International*, 143, 105905. Retrieved from <http://dx.doi.org/10.1016/j.envint.2020.105905> doi: 10.1016/j.envint.2020.105905
- Siler-Evans, K., Azevedo, I. L., Morgan, M. G., & Apt, J. (2013). Regional variations in the health, environmental, and climate benefits of wind and solar generation. *Proceedings of the National Academy of Sciences*, 110(29), 11768–11773.
- Skidmore, M. E., Sims, K. M., & Gibbs, H. K. (2023, November). Agricultural intensification and childhood cancer in Brazil. *Proceedings of the National Academy of Sciences*, 120(45), e2306003120. doi: 10.1073/pnas.2306003120

- Sodhi, M., Banaszek, L., Magee, C., & Rivero-Hudec, M. (2022). Economic lifetimes of solar panels. *Procedia CIRP*, 105, 782–787.
- Soil Survey Staff. (2021, October). *Gridded Soil Survey Geographic (gSSURGO) Database for the Conterminous United States*. United States Department of Agriculture, Natural Resources Conservation Service.
- Surveillance, Epidemiology, and End Results (SEER) Program. (2024). *U.S. County Population Data - 1969-2022*.
- Tago, D., Andersson, H., & Treich, N. (2014, January). Pesticides and Health: A Review of Evidence on Health Effects, Valuation of Risks, and Benefit-Cost Analysis. In *Preference Measurement in Health* (Vol. 24, pp. 203–295). Emerald Group Publishing Limited. doi: 10.1108/S0731-219920140000024006
- Taylor, C. A. (2022). *Cicadian Rhythm: Insecticides, Infant Health and Long-term Outcomes*. (Unpublished)
- Thelin, G., & Stone, W. (2013). *Estimation of Annual Agricultural Pesticide Use for Counties of the Conterminous United States, 1992-2009* (Scientific Investigations Report No. 2013-5009). U.S. Geological Survey.
- Tibebu, T. B., Hittinger, E., Miao, Q., & Williams, E. (2021). What is the optimal subsidy for residential solar? *Energy Policy*, 155, 112326.
- U. S. Global Change Research Program. (2023). *Fifth National Climate Assessment* (Tech. Rep.). U.S. Global Change Research Program, Washington, DC.
- United Nations. (2024). *UN-Water SDG 6 Data Portal*. <https://sdg6data.org/>. (Online; accessed 11 June 2024)
- U.S. Environmental Protection Agency. (2020). *Glyphosate interim registration review decision*. Retrieved from <https://www.regulations.gov/document/EPA-HQ-OPP-2009-0361-14442> (Docket ID: EPA-HQ-OPP-2009-0361)
- U.S. EPA. (2019, April). *Glyphosate: Response to Comments, Usage, and Benefits* (Tech. Rep.).
- US EPA. (2023, November). *Biden-Harris Administration Announces \$2 Billion to Fund Environmental and Climate Justice Community Change Grants as Part of Investing in America Agenda* [News Release].

- US EPA. (2024, July). *Biden-Harris Administration Announces More Than \$325 Million in Environmental and Climate Justice Community Change Grants* [News Release].
- USDA. (2024). *Adoption of Genetically Engineered Crops in the U.S.*
- Vogel, S. A. (2008). From 'The Dose Makes the Poison' to 'The Timing Makes the Poison': Conceptualizing Risk in the Synthetic Age. *Environmental History*, 13(4), 667–673.
- vom Saal, F. S., Akingbemi, B. T., Belcher, S. M., Crain, D. A., Crews, D., Guidice, L. C., ... Zoeller, R. T. (2010, February). Flawed experimental design reveals the need for guidelines requiring appropriate positive controls in endocrine disruption research. *Toxicological Sciences*, 115(2), 612–613. Retrieved from <http://dx.doi.org/10.1093/toxsci/kfq048> doi: 10.1093/toxsci/kfq048
- Waitzman, N. J., Jalali, A., & Grosse, S. D. (2021, April). Preterm birth lifetime costs in the United States in 2016: An update. *Seminars in Perinatology*, 45(3), 151390. doi: 10.1016/j.semperi.2021.151390
- Walczak, J., & Drenkard, S. (2017). *State and local sales tax rates in 2017*. Washington: Tax Foundation.
- Walker, K., & Herman, M. (2024). *Tidycensus: Load US Census Boundary and Attribute Data as 'tidyverse' and 'sf'-Ready Data Frames*.
- Wing, O. E. J., Lehman, W., Bates, P. D., Sampson, C. C., Quinn, N., Smith, A. M., ... Kousky, C. (2022, January). Inequitable patterns of US flood risk in the Anthropocene. *Nature Climate Change*, 1–7. doi: 10.1038/s41558-021-01265-6
- World Bank Open Data. (2024). *GDP per capita*. Retrieved 2024-01-25, from <https://data.worldbank.org/>
- Wrenn, D. H. (2023, November). The Effect of Natural Disasters and Extreme Weather on Household Location Choice and Economic Welfare. *Journal of the Association of Environmental and Resource Economists*. doi: 10.1086/728887
- Xu, Y., Li, J., Tan, Q., Peters, A. L., & Yang, C. (2018). Global status of recycling waste solar panels: A review. *Waste Management*, 75, 450–458.
- Xu, Z., Hennessy, D. A., Sardana, K., & Moschini, G. (2013). The Realized Yield Effect of Genetically Engineered Crops: U.S. Maize and Soybean. *Crop Science*, 53(3), 735–745. doi: 10.2135/cropsci2012.06.0399

- Yu, J., Wang, Z., Majumdar, A., & Rajagopal, R. (2018). Deepsolar: A machine learning framework to efficiently construct a solar deployment database in the united states. *Joule*, 2(12), 2605–2617.
- Zhang, L., Rana, I., Shaffer, R. M., Taioli, E., & Sheppard, L. (2019, July). Exposure to glyphosate-based herbicides and risk for non-Hodgkin lymphoma: A meta-analysis and supporting evidence. *Mutation Research/Reviews in Mutation Research*, 781, 186–206. doi: 10.1016/j.mrrev.2019.02.001