THE UNIVERSITY OF CHICAGO

THE FOOD PROBLEM AND THE AGGREGATE PRODUCTIVITY CONSEQUENCES OF CLIMATE CHANGE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
ISHAN BROWNELL NATH

CHICAGO, ILLINOIS
AUGUST 2019

I thank my advisers, Michael Greenstone, Chang-Tai Hsieh, and Pete Klenow for their guidance, mentorship, and patience.

I also wish to acknowledge the love and support of my family members - Julia Brownell Nath, Rekha Nath, Harsh Nath, Kimberly Pesavento, Bill Brownell, Rachel Brownell, Aruna Arora, Pratibha Nath, Ajay Arora, Suresh Nath, Rosie Nath, and Zoe Nath - without whom this work would not have been possible.

All errors and omissions are my own.

Contents

List of Tables vi

List of Figures xi

Chapter 1: The Food Problem and the Aggregate Productivity Consequences of Climate Change 1

Abstract 1

1 Introduction 2

2 Data 8

3 Empirical Strategy 11
- 3.1 Conceptual Framework 12
- 3.2 Causal Effect of Temperature 13
- 3.3 Heterogeneity and Adaptation 15

4 Empirical Results 17
- 4.1 Main Regression Results 17
- 4.2 Robustness 20
- 4.3 U.S. Results 21
- 4.4 Projected Global Sensitivity to Extreme Temperatures 22

5 Model 25
- 5.1 Model Ingredients 25
- 5.2 Comparative Statics 29

6 Model Estimation 32
- 6.1 Parameter Estimates 32
- 6.2 Model Fit 33

7 Model Counterfactuals 37
- 7.1 Estimated Productivity Impacts 37
- 7.2 Comparative Advantage and Trade 40

7.3 Sectoral Reallocation . . . 41
7.4 Aggregate Productivity and Willingness-to-Pay . . . 43
7.5 Low Trade Cost Counterfactual . . . 47
7.6 Future Projections . . . 49

8 Supporting Empirical Evidence **51**

9 Policy Implications **54**

10 Conclusion **56**

Chapter 2: Do Renewable Portfolio Standards Deliver Cost-Effective Carbon Abatement? **58**

Abstract **58**

1 Introduction **59**

2 Renewable Portfolio Standards **64**

3 Conceptual Framework **69**
3.1 Representative Utility Model . . . 70
3.2 Empirical Requirements for Estimating the Full Costs of RPS . . . 73

4 Data Sources and Summary Statistics **75**
4.1 RPS Program Data . . . 76
4.2 Electricity Sector . . . 77
4.3 Manufacturing Employment . . . 78
4.4 Summary Statistics . . . 79

5 Empirical Strategy **79**

6 Results **83**
6.1 Net RPS Requirements and Retail Electricity Prices . . . 83
6.2 Heterogeneity in RPS Price Effects . . . 89
6.3 Economic Activity . . . 92
6.4 Generation . . . 93

7 Interpretation **96**

8 Conclusion 99

Chapter 3: A Global View of Creative Destruction 101

Abstract 101

1 Introduction 102

2 Facts from Canadian and U.S. Manufacturing 105

3 Exogenous Innovation 112

3.1 Static Equilibrium 112

3.2 Innovation 115

3.3 Calibration 121

3.4 Firm Dynamics 122

4 Endogenous Innovation 129

5 Additional Exercises 139

5.1 U.S.-Canada Simulations 139

5.2 Alternative Assumptions about Idea Flows 139

6 Conclusion 143

References 145

Appendix A: Additional Regression Results 152

Appendix B: U.S. Results 163

Appendix C: China Results 165

Appendix D: Adaptation Benefits and Costs 168

Appendix E: Additional Model Fit Figures 176

Appendix F: Country-by-Country Model Counterfactual Results 180

Appendix G: Additional RPS Tables and Figures 239

List of Tables

1 Global Firm-Level Panel Microdata 9
2 Effects of Daily Temperature on Annual Revenue per Worker 18
3 Model Parameters and Target Moments 33
4 Parameter Estimates 34
5 Summary of Model Fit 35
6 Counterfactual Ag GDP Shares - Selected Countries 42
7 Counterfactual GDP Losses (Share of GDP) - Selected Countries 44
8 Equivalent Variation Willingness-to-Pay (Share of GDP) - Selected Countries 45
9 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases 48
10 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits 50
11 Country-Level Panel Data 52
12 Country-Level Panel Regression 53
13 Summary Statistics 80
14 Estimates of RPS Impact on Retail Electricity Prices 86
15 Robustness Checks for RPS Impact 89
16 Heterogeneous Effects of RPS Programs on Retail Electricity Prices 90
17 RPS Effect on Sales and Employment 92
18 Estimates of RPS Impact on Generation and CO_2 Emissions (Trend Break) 94
19 Estimated Cost of Abating CO_2 Emissions from RPS 97
20 Job Flows in the U.S. and Canada 107
21 Job Flows in Canada 108
22 Job Flows in the U.S. 110
23 Export Product Churn 112
24 Markups 113
25 Channels of Innovation 116
26 Probability of Creative Destruction 117
27 Data Moments used for Calibration 121
28 Estimates of Model Parameters 122
29 Firm Dynamics, Data vs. Simulations 124
30 Estimates of Model Parameters, Endogenous Innovation 133

31 Gains From Trade . . . 138
32 Gains From Trade Liberalization . . . 138
33 Job Flows in Canada — Post-CUSFTA versus Pre-CUSFTA . . . 140
34 Counterfactuals with Alternative Assumptions on Idea Flows . . . 142
A-1 U.S. Results . . . 163
A-2 U.S. Energy Results . . . 164
A-3 Counterfactual Ag Net Export Share of GDP - Sub-Saharan Africa . . . 180
A-4 Counterfactual Ag Net Export Share of GDP - Sub-Saharan Africa . . . 181
A-5 Counterfactual Ag Net Export Share of GDP - Middle East and North Africa . . . 182
A-6 Counterfactual Ag Net Export Share of GDP - Asia . . . 183
A-7 Counterfactual Ag Net Export Share of GDP - South America . . . 184
A-8 Counterfactual Ag Net Export Share of GDP - North and Central America . . . 184
A-9 Counterfactual Ag Net Export Share of GDP - Europe . . . 185
A-10 Counterfactual Ag Net Export Share of GDP - Europe . . . 186
A-11 Counterfactual Ag Net Export Share of GDP - Western Pacific and Oceania . . . 186
A-12 Counterfactual Ag Domestic Expenditure Shares - Sub-Saharan Africa . . . 187
A-13 Counterfactual Ag Domestic Expenditure Shares - Sub-Saharan Africa . . . 188
A-14 Counterfactual Ag Domestic Expenditure Shares - Middle East and North Africa . . . 189
A-15 Counterfactual Ag Domestic Expenditure Shares - Asia . . . 190
A-16 Counterfactual Ag Domestic Expenditure Shares - South America . . . 191
A-17 Counterfactual Ag Domestic Expenditure Shares - North and Central America . . . 191
A-18 Counterfactual Ag Domestic Expenditure Shares - Europe . . . 192
A-19 Counterfactual Ag Domestic Expenditure Shares - Europe . . . 193
A-20 Counterfactual Ag Domestic Expenditure Shares - Western Pacific and Oceania . . . 193
A-21 Counterfactual Ag GDP Shares - Sub-Saharan Africa . . . 194
A-22 Counterfactual Ag GDP Shares - Sub-Saharan Africa . . . 195
A-23 Counterfactual Ag GDP Shares - Middle East and North Africa . . . 196
A-24 Counterfactual Ag GDP Shares - Asia . . . 197
A-25 Counterfactual Ag GDP Shares - South America . . . 198
A-26 Counterfactual Ag GDP Shares - North and Central America . . . 198
A-27 Counterfactual Ag GDP Shares - Europe . . . 199
A-28 Counterfactual Ag GDP Shares - Europe . . . 200

A-29 Counterfactual Ag GDP Shares - Western Pacific and Oceania 200
A-30 Counterfactual GDP Losses (Share of GDP) - Sub-Saharan Africa 201
A-31 Counterfactual GDP Losses (Share of GDP) - Sub-Saharan Africa 202
A-32 Counterfactual GDP Losses (Share of GDP) - Middle East and North Africa 203
A-33 Counterfactual GDP Losses (Share of GDP) - Asia 204
A-34 Counterfactual GDP Losses (Share of GDP) - South America 205
A-35 Counterfactual GDP Losses (Share of GDP) - North and Central America 205
A-36 Counterfactual GDP Losses (Share of GDP) - Europe 206
A-37 Counterfactual GDP Losses (Share of GDP) - Europe 207
A-38 Counterfactual GDP Losses (Share of GDP) - Western Pacific and Oceania 207
A-39 Equivalent Variation Willingness-to-Pay (Share of GDP) - Sub-Saharan Africa 208
A-40 Equivalent Variation Willingness-to-Pay (Share of GDP) - Sub-Saharan Africa 209
A-41 Equivalent Variation Willingness-to-Pay (Share of GDP) - Middle East and North Africa 210
A-42 Equivalent Variation Willingness-to-Pay (Share of GDP) - Asia 211
A-43 Equivalent Variation Willingness-to-Pay (Share of GDP) - South America 212
A-44 Equivalent Variation Willingness-to-Pay (Share of GDP) - North and Central America 212
A-45 Equivalent Variation Willingness-to-Pay (Share of GDP) - Europe 213
A-46 Equivalent Variation Willingness-to-Pay (Share of GDP) - Europe 214
A-47 Equivalent Variation Willingness-to-Pay (Share of GDP) - Western Pacific and Oceania 214
A-48 Counterfactual Change in Food Prices - Sub-Saharan Africa 215
A-49 Counterfactual Change in Food Prices - Sub-Saharan Africa 216
A-50 Counterfactual Change in Food Prices - Middle East and North Africa 217
A-51 Counterfactual Change in Food Prices - Asia 218
A-52 Counterfactual Change in Food Prices - South America 219
A-53 Counterfactual Change in Food Prices - North and Central America 219
A-54 Counterfactual Change in Food Prices - Europe 220
A-55 Counterfactual Change in Food Prices - Europe 221
A-56 Counterfactual Change in Food Prices - Western Pacific and Oceania 221
A-57 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Sub-Saharan Africa 222

A-58 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Sub-Saharan Africa 223
A-59 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Middle East and North Africa 224
A-60 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Asia 225
A-61 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - South America 226
A-62 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - North and Central America 226
A-63 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Europe 227
A-64 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Europe 228
A-65 Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Western Pacific and Oceania 228
A-66 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Sub-Saharan Africa 229
A-67 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Sub-Saharan Africa 230
A-68 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Middle East and North Africa 231
A-69 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Asia 232
A-70 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - South America 233
A-71 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - North and Central America 234
A-72 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Europe 235
A-73 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Europe 236

A-74 Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Western Pacific and Oceania 237
A-75 Country-Level Panel Regression 238

List of Figures

1 Cline (2007) Projected Impact of Climate Change on Agricultural Productivity, 2080-2099 . . . 2
2 Comparative Advantage and Specialization in Agriculture . . . 3
3 Predicted Heterogeneous Response of Annual Manufacturing Revenue per Worker to Daily Maximum Temperature . . . 19
4 Estimated Response of U.S. Annual Manufacturing Revenue per Worker to Daily Maximum Temperature . . . 21
5 Predicted Effect of a 40°C Day on Annual Manufacturing Revenue per Worker . . . 23
6 Predicted Effect of a -5°C Day on Annual Manufacturing Revenue per Worker . . . 23
7 Agriculture Share of GDP - Data vs. Simulation . . . 35
8 Relative Price of Food - Data vs. Simulation . . . 36
9 Projected Impact of Climate Change on Manufacturing Productivity . . . 38
10 Projected Impact of Climate Change on Agricultural Relative Productivity . . . 38
11 Projected Impact of Climate Change on Agricultural Net Exports . . . 40
12 Projected Impact of Climate Change on Agricultural GDP Share . . . 43
13 Willingness-to-Pay to Avoid Climate Change . . . 46
14 Projected Percentage Change in Food Prices . . . 46
15 Direct Costs to Import a 20-Foot Long Container (USD) . . . 55
16 Days to Import a Container . . . 56
17 RPS Passage by State . . . 65
18 RPS Passage by State . . . 68
19 RPS Passage by State . . . 69
20 Estimated Effects of RPS Programs on Net Renewable Requirements . . . 83
21 Estimated Effects of RPS Programs on Retail Electricity Prices . . . 84
22 Δ Job Destruction in Canada vs. Δ Canadian Tariffs . . . 109
23 Δ Job Creation in Canada from Exportsvs. Δ U.S. Tariffs . . . 109
24 Distribution of Employment and Labor Productivity . . . 111
25 Adjustment to Increase in U.S. Innovation Rate . . . 119
26 Simulated TFP Growth Rate and Relative Wage vs. Trade Costs . . . 125
27 Simulated Job Creation and Destruction vs. Trade Costs . . . 126
28 Simulated Job Destruction from *Large* Firms vs. Trade Costs . . . 127

29 Simulated Job Flows from Trade vs. Trade Costs 127
30 Simulated Distribution of Employment and Labor Productivity 128
31 Simulated arrival rates after trade liberalization 134
32 Research labor shares after trade liberalization 134
33 Simulated real consumption after trade liberalization 135
34 Simulated steady-state real consumption vs. various trade costs 136
35 Simulated job flows from trade liberalization 137
A-1 Predicted Heterogeneous Response of Annual Manufacturing Revenue to Daily Maximum Temperature 152
A-2 Predicted Heterogeneous Response of Annual Manufacturing Employment to Daily Maximum Temperature 153
A-3 Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature 154
A-4 Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature 155
A-5 Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE 156
A-6 Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE 157
A-7 Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE 158
A-8 Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - Controls for Capital 159
A-9 Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature 160
A-10 Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature 161
A-11 Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature 162
A-12 Estimated Response of U.S. Annual Manufacturing TFPR to Daily Maximum Temperature 163

A-13 Estimated Response of U.S. Annual Manufacturing Plant-Level Energy Expenditures to Daily Maximum Temperature . . . 164
A-14 China Replication - Overlapping Years . . . 166
A-15 China Manufacturing Temperature Sensitivity - Estimated and Predicted . . . 167
A-16 China Replication - Different Years . . . 167
A-17 Predicted Effect of a 40°C Day on Annual Manufacturing Revenue per Worker At 2080 Average Temperatures . . . 169
A-18 Firm-Level Adaptation Costs (Share of Manufacturing Output) . . . 170
A-19 Firm-Level Adaptation Net Benefits (Share of Manufacturing Output) . . . 170
A-20 Predicted Effect of a 40°C Day on Annual Services Revenue per Worker . . . 171
A-21 Predicted Effect of a -5°C Day on Annual Services Revenue per Worker . . . 171
A-22 Projected Change in Exposure to Extreme Heat . . . 172
A-23 Projected Change in Exposure to Extreme Cold . . . 172
A-24 Projected Impact of Climate Change on Services Productivity . . . 173
A-25 Projected Impact of Climate Change on Manufacturing Productivity Accounting for Economic Growth and Adaptation . . . 174
A-26 Projected Impact of Climate Change on Services Productivity Accounting for Economic Growth and Adaptation . . . 175
A-27 Manufacturing Share of GDP - Data vs. Simulation . . . 176
A-28 Services Share of GDP - Data vs. Simulation . . . 176
A-29 Agriculture Share of GDP - Data vs. Simulation . . . 177
A-30 Domestic Production Share of Agriculture Expenditures - Data vs. Simulation . . . 177
A-31 Manufacturing Domestic Production Share of Expenditures - Data vs. Simulation . 178
A-32 Log GDP Per Capita - Data vs. Simulation . . . 178
A-33 Domestic Production Share of Expenditures in Agriculture - Model Simulation . . . 179
A-34 Agriculture Share of GDP - Data vs. Simulation Stone-Geary Specification . . . 179
A-35 REC Tracking Markets . . . 239
A-36 Implementation of Energy Programs by State and Year . . . 240
A-37 Electricity Prices Before and After RPS Passage, by Sector . . . 241
A-38 Estimated Effects of RPS Programs on Net Renewable Requirements (Extended Post Period) . . . 242

A-39 Estimated Effects of RPS Programs on Gross Renewable Requirements (Extended Post Period) 243
A-40 Estimated Effects of RPS Programs on Retail Electricity Prices (Extended Post Period) 244
A-41 CO_2 Emissions Intensity Before and After RPS Passage 245

Chapter 1: The Food Problem and the Aggregate Productivity Consequences of Climate Change

Abstract

Climate change is projected to sharply reduce agricultural productivity in hot developing countries and raise it in temperate regions. Reallocation of labor across sectors could temper the aggregate impacts of these changes if hotter regions shift toward importing food and specializing in manufacturing or exacerbate them if subsistence food requirements push labor toward agriculture where its productivity suffers most. I quantify these effects in two steps. First, I project changes in global comparative advantage by using firm-level micro-data from 17 countries covering over half the world's population to estimate the heterogeneous effect of temperature on output per worker in manufacturing and services. I find large effects of extremely hot and cold temperatures on non-agricultural output per worker, but treatment effects diminish with income and expectations of temperature such that the projected impact of climate change is larger in agriculture than non-agriculture. Second, I embed my estimates in an open-economy model of structural transformation that matches moments on output-per-worker, sectoral specialization, and trade for 158 countries. Simulations suggest that subsistence food requirements dominate labor reallocation in response to climate change on average and the global decline in GDP is 12.0% larger, and 52.1% larger for the poorest quartile of the world, when accounting for sectoral reallocation than in the counterfactual with fixed sectoral shares. The aggregate willingness-to-pay to avoid climate change is 1.5-2.7% of annual GDP and 6.2-10.0% for the poorest quartile. Trade reduces the welfare costs of climate change relative to autarky by only 7.4% under existing policy, but by 30.7% overall and by 68.2% for the poorest quartile in an alternative scenario with reduced trade costs.[1]

[1]I appreciate helpful comments from Tamma Carleton, Michael Dinerstein, Alessandra Voena, Manasi Deshpande, Lauren Bergquist, Will Rafey, Jeremy Pearce, Liangjie Wu, Eric English, Michele Carter, and Karthik Nagarajan. I thank Henry Zhang, Nick Tsivanidis, and Cian Ruane for help processing and cleaning my datasets, Steve Mohr and Theodor Kulczycki for guidance on coding and computing, and the Chicago Booth Initiative on Global Markets (IGM), and Jennifer Williams and Peggy Eppink in particular, for providing access to the Amadeus dataset from Bureau van Dijk. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

Existing evidence suggests that climate change will cause large and heterogeneous changes in agricultural productivity across the world during the 21st century. Figure 1 shows estimates of the country-level impact of climate change on agricultural productivity from Cline (2007), which synthesizes evidence from economics, agronomy, and climate science.[2]

Figure 1: Cline (2007) Projected Impact of Climate Change on Agricultural Productivity, 2080-2099

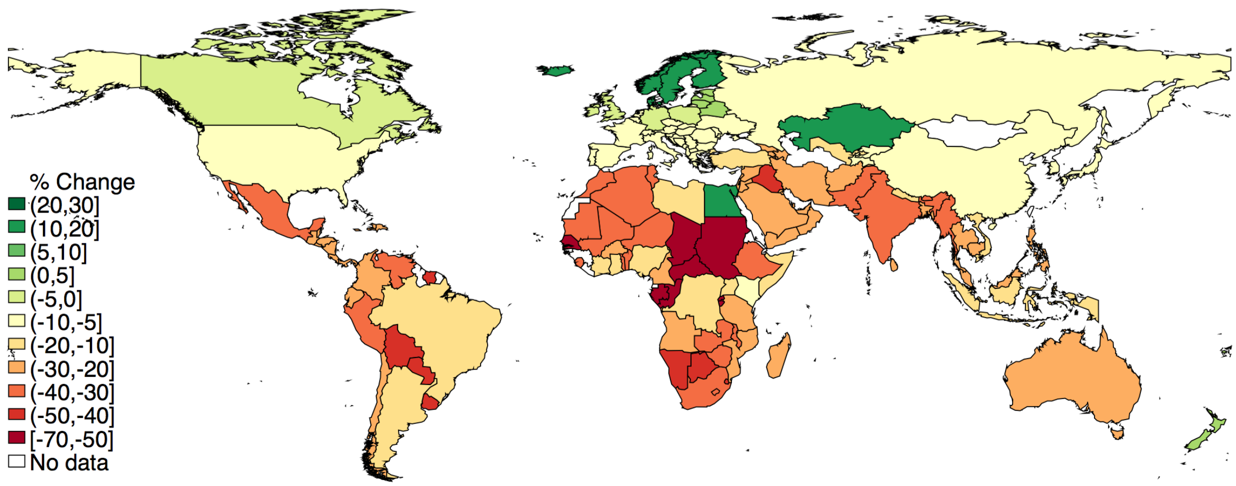

Notes: Figure shows the projected change in revenue per acre from producing grains, vegetables, fruits, and livestock according to analysis by Cline (2007).

The projections in Figure 1 show large declines in agricultural productivity of 30-60% in hot regions such as Sub-Saharan Africa and South Asia, with neutral or positive effects in cold regions such as Canada and northern Europe. This pattern suggests large potential gains from reallocating the geography of agricultural production. If productivity suffers greatly in some places and improves in others, intuition suggests that market forces will push agriculture toward temperate climates and substantially reduce the welfare consequences of these changes. This

[2] I explain the methods used in Cline (2007) more in Section 7.1. The findings are broadly consistent with a large body of economics research on the impacts of climate change on agriculture, which includes Mendelsohn, Nordhaus and Shaw (1994), Deschenes and Greenstone (2007), Schlenker and Roberts (2009), and Schlenker and Lobell (2010), among many others. I use Cline (2007) in this paper because it is the best available source for country-level impact estimates that use globally representative data and account for adaptation.

paper investigates the conditions under which this intuition does and does not hold true, and quantifies the aggregate productivity consequences of climate change in the presence of the general equilibrium changes in sectoral specialization likely to occur in practice.[3]

Two key elements of sectoral allocation complicate the idea that the changes in Figure 1 will push agriculture away from the equator. First, these estimates show the change in the absolute advantage of agricultural production, whereas comparative advantage across sectors drives international trade. Ricardian models of trade will only predict that Canada will export more food and India will import more food if the *relative* productivity of agriculture rises in Canada and falls in India.[4] Given existing evidence that temperature also affects non-agricultural productivity, the change in comparative advantage is not immediately clear.[5]

Second, comparative advantage does not exclusively, or even primarily, determine sectoral specialization. Figure 2 shows that poor countries have much higher agricultural labor shares

Figure 2: Comparative Advantage and Specialization in Agriculture

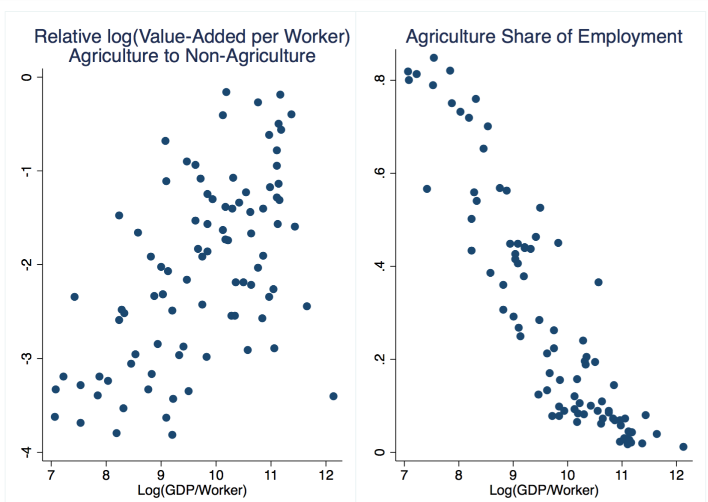

Notes: Figure shows data from Tombe (2015) that adjusts for prices for the global cross-section in 2005. Poor countries specialize heavily in agriculture despite low productivity relative to other sectors.

[3] I use the phrase climate change to refer to shifting distributions of temperature in this paper. Other consequences of climate change, such as sea-level rise or intensified hurricanes, are beyond the scope of the analysis.

[4] I use the word food interchangeably with agricultural production in this paper because subsistence requirements for food drive the key features of consumer preferences in my analysis.

[5] This evidence includes work by Zhang, Deschenes, Meng and Zhang (2018) and Somanathan, Somanathan, Sudarshan, Tewari et al. (2015).

despite lower relative value-added per worker in agriculture compared to non-agriculture. Lagakos and Waugh (2013) calculate that, adjusting for prices, the gap in aggregate output per worker between the 90th to 10th percentile of the world's income distribution is 45 to 1 in agriculture, but just 4 to 1 in non-agriculture. Yet agriculture's share of employment averages 65% in 10th percentile countries and only 3% in 90th percentile countries. Trade in agriculture plays only a small role in developing countries. The average person in the poorest quartile of the world consumes 91.3% domestically produced food, compared with 45.1% in the richest quartile. In these relatively closed economies, high agricultural production and labor shares follow from the high consumption shares necessary for people with low incomes to meet subsistence requirements for food. Projecting the effects of climate change on sectoral reallocation requires accounting for the forces driving this existing global equilibrium in which poor countries specialize in agriculture despite low absolute and relative productivity, a fact which the literature on the general equilibrium effects of climate change has not yet confronted.

I address both these challenges in my analysis. First, to project changes in agricultural comparative advantage, I estimate the country-level impact of climate change on productivity in manufacturing and services using a global dataset of nationally representative firm-level panel data from 17 countries covering over half the world's population, and representing nearly the full range of current temperatures and income levels. Using methods developed by Carleton et al. (2018), I use my data to estimate plausibly causal treatment effects of extreme temperatures on output-per-worker, and account for firm-level adaptation by allowing these treatment effects to vary with income and expectations of temperature.

I find that extreme heat and extreme cold can both have important effects on non-agricultural productivity, but with strong evidence of adaptation in rich countries and to temperatures with which agents are accustomed. In poor countries with moderate climates, an extreme day with daily maximum temperature of 40°C or -5°C reduces annual output-per-worker by up to 0.4%, approximately the equivalent of one full working day.[6] Effects are about half as large in middle-

[6]I find similar effects for manufacturing and services firms, though I lack data coverage for services firms in poor countries where the effects of temperature are most detectable.

income countries, and smaller still in those places that experience given extremes more frequently. The effects of extreme days in rich countries are negligible, with some evidence of mild effects from unexpected extremes caused by hot days in cold places and cold days in hot places. I combine these estimates of predicted temperature sensitivity with global climate model predictions of future temperatures to project the country-level effects of climate change on manufacturing and services productivity. The effects of climate change on non-agricultural productivity are non-trivial in some poor countries, but generally small relative to productivity losses in agriculture. Thus, the change in the global relative productivity in agriculture is qualitatively similar to the change in absolute productivity.

Second, I construct a global open economy model of structural transformation that explains the existing distribution of sectoral specialization as a function of sector-level productivities. The model incorporates two key features of consumer preferences - nonhomothetic preferences and low substitutability across sectors - that explain the high agricultural share of consumption in poor countries with high relative prices for food. Gollin, Parente and Rogerson (2007) refer to the macro-development effects of these subsistence requirements as "the food problem," which drives developing countries to specialize in a relatively low-productivity sector because people need food to survive. My model also includes Ricardian comparative advantage within and across sectors, which tends to force countries with low relative productivity in agriculture toward specializing production in other sectors, but only to the extent that they are open to trade.[7]

Thus, my model shows that two competing effects govern the response of sectoral specialization to climate change, and that their net effect could either temper or exacerbate the aggregate

[7]While rural-urban migration within countries plays a key implicit role in the sectoral reallocation captured by my model, I hold the global distribution of population fixed across countries rather than allowing for international migration. To justify this assumption, I note that some combination of home-bias and barriers to migration are sufficient to maintain welfare differences of two orders of magnitude between the poorest and richest countries in the existing global equilibrium. While climate change is likely to exacerbate global income differences, it seems plausible that the strength of these forces will continue to keep most people confined to the places where they already live. To the extent that climate change does cause substantial international migration, my analysis captures the welfare consequences for those people left behind in the countries suffering major impacts.

consequences of the sector-level changes. If the trade effect dominates, then countries can dampen the effect of falling agricultural productivity by shifting production to other sectors; exporting more manufactured goods and importing more food. To the extent that climate change exacerbates 'the food problem' by reducing agricultural productivity, however, the general equilibrium response could drive labor toward the sector suffering large declines in productivity and worsen the aggregate impact.

To quantify the relative strengths of these mechanisms, I estimate my model to match data on income levels, trade flows, and sectoral specialization for 158 countries covering over 99.9% of global GDP. I embed the empirically estimated projected impacts of climate change on productivity in agriculture, manufacturing, and services into the estimated model, and conduct counterfactual simulations that calculate the effects of climate change on sectoral specialization, trade, prices, GDP, and welfare.[8] I disentangle the effects of 'the food problem' and trade by running separate counterfactuals with no reallocation, in autarky, with estimated trade costs, and in an alternative policy scenario with reduced barriers to trade.

I find that the net effect of sectoral reallocation exacerbates the effects of climate change on aggregate productivity. Climate change raises the agriculture share of GDP by 2.8 percentage points in the poorest quartile of the world, which suffers large falls in relative agricultural productivity, as 'the food problem' outweighs the trade response on average. Comparative advantage predominantly shifts away from the equator and net exports in agriculture increase in colder countries, such as those in northern Europe, and in a few hot countries that suffer declines in agricultural productivity that are small relative to those of their close trading partners or to their decline in manufacturing productivity. Net imports of food rise in most hot countries in the developing world, but only some countries are sufficiently open to trade for this effect to substantially alter sectoral specialization. Overall, climate change reduces global GDP by 12.0% more, and by 52.1% more for the poorest quartile of the world, when accounting for the full effects of sectoral reallocation than in the naive counterfactual with fixed sectoral shares.

[8]My simulations use Cline (2007) for the effects of climate change on agriculture, and my own estimates for manufacturing and services.

The equivalent variation willingness-to-pay (WTP) to avoid each year of climate change is between 1.5% and 2.7% of contemporaneous global GDP, depending on assumptions about economic growth. The worst effects are concentrated in poor countries that comprise a small share of global GDP, but a substantial portion of the population. The average person in the poorest quartile of the global income distribution suffers losses of 6.2%-10.0% of their income. Trade reduces the aggregate global willingness-to-pay to avoid climate change by 7.4% relative to autarky under existing policy, and by 30.7% under the alternative low trade barrier counterfactual. Reducing trade barriers has heterogeneous effects, increasing the costs of climate change in some regions as greater interdependence makes countries less vulnerable to local shocks but more vulnerable to global shocks.[9] Reducing trade barriers is particularly valuable for climate change adaptation in poor countries. Trade reduces WTP for the poorest quartile of the global population by only 4.5% relative to autarky under existing policy, largely because many poor countries are mostly closed to trade, but by 68.2% in the low trade cost counterfactual.

This paper relates to several literatures on the economics of climate change and macroeconomic development. The two most similar papers are Costinot, Donaldson and Smith (2016), who examine reallocation across crops but do not consider income effects or cross-sector reallocation, and Desmet and Rossi-Hansberg (2015), who primarily focus on the important role for international migration in climate change adaptation. The latter paper includes changes in the global distribution of sectoral specialization in the model, but does not attempt to incorporate realistic trade costs or the importance of 'the food problem' in the analysis. My paper is the first to consider the effects of climate change on structural transformation.

My empirical work on temperature and productivity builds on country-level estimates produced by Somanathan, Somanathan, Sudarshan, Tewari et al. (2015) and Zhang, Deschenes, Meng and Zhang (2018) in India and China, and uses methods that closely follow Carleton et al. (2018). The model builds on several papers that consider structural transformation in an open-economy setting, including Tombe (2015), Uy, Yi and Zhang (2013), and Teignier (2018). I also

[9] Note that this nets out gains from trade that are unrelated to climate change adaptation. Thus, these results do not imply that these countries are worse off overall from reducing trade barriers.

use a nonhomothetic CES specification for consumer preferences from Comin, Lashkari and Mestieri (2015) to best capture the observed pattern of structural transformation. Finally, some of my counterfactual predictions about the role of trade and the spatial correlation of shocks relate to the work of Dingel, Meng and Hsiang (2019).

The paper is structured as follows. Sections 2, 3, and 4 describe the data, empirical strategy, and results for the estimation of the relationship between temperature and non-agricultural productivity. Section 5 lays out the model. Section 6 explains the model estimation and describes the model's success in fitting the data. Section 7 contains the counterfactual model simulations. Section 8 provides additional country-level panel regression evidence on the impact of agriculture-biased productivity shocks on sectoral reallocation. Section 9 discusses implications for policy and Section 10 concludes.

2 Data

Firm Data

I assemble a globally representative panel of firm-level microdata to estimate the relationship between temperature and productivity in manufacturing and services. Table 1 lists the countries and years included in the dataset as well as the data source for each country. The data combines surveys administered by national governments with data acquired from the Amadeus database maintained by Bureau van Dijk (BVD). BVD is a private company owned by Moody's Analytics that collects and distributes firm-level financial information from around the world. They collect data both by acquiring administrative data directly from national business registers and by conducting their own surveys.

I restrict my analysis to those countries with nationally representative panels. This includes government-level surveys from India, Colombia, Indonesia, China, and the United States, and Amadeus data from twelve European countries with mandatory filing requirements according to BVD documentation.[10] Bloom, Draca and Van Reenen (2016) report that the data in most

[10]Importantly, the online version of the Amadeus database does not maintain accurate historical records. Thus, I download the data directly from the 2005, 2010, and 2015 vintages (CDs). Each Amadeus vintage contains 10 years of historical data for each firm. I match firms across years using BVD's unique firm identification number, and drop a small subset of observations with inconsistent data across

Table 1: Global Firm-Level Panel Microdata

Country	Data Source	Dataset	Years
Austria	Bureau Van Dijk	Amadeus	1995-2014
Belgium	Bureau Van Dijk	Amadeus	1995-2014
China	National Bureau of Statistics	Chinese Industrial Survey	2003-2012
Colombia	National Administrative Department of Statistics (DANE)	Annual Manufacturing Survey	1977-1991
Finland	Bureau Van Dijk	Amadeus	1995-2014
France	Bureau Van Dijk	Amadeus	1995-2014
Germany	Bureau Van Dijk	Amadeus	1995-2014
Greece	Bureau Van Dijk	Amadeus	1995-2014
India	Central Statistical Office	Annual Survey of Industries	1985-2007
Indonesia	Badan Pusat Statistik	Annual Manufacturing Survey	1975-1995
Italy	Bureau Van Dijk	Amadeus	1995-2014
Norway	Bureau Van Dijk	Amadeus	1995-2014
Spain	Bureau Van Dijk	Amadeus	1995-2014
Sweden	Bureau Van Dijk	Amadeus	1995-2014
Switzerland	Bureau Van Dijk	Amadeus	1995-2014
United Kingdom	Bureau Van Dijk	Amadeus	1995-2014
United States	Census Bureau	Annual Survey of Manufacturers, Census of Manufacturers	1976-2014

Notes: Data includes revenue and number of employees, with varying coverage of capital stock (tangible fixed assets) and wage-bill. Amadeus data includes both manufacturing and services firms.

of these European countries contains nearly the full population of public and private firms.[11] Gopinath, Kalemli-Özcan, Karabarbounis and Villegas-Sanchez (2017) also use data from Amadeus and Alfaro and Chen (2018) use data from Orbis, a related firm dataset produced by BVD.

My sample covers both manufacturing and services firms in developed and developing countries. While the government surveys cover only manufacturing firms, the BVD data covers the entire spectrum of 2-digit industries. I report results for the pooled sample of all firms, separately for manufacturing firms, and separately for services firms, though the latter subset lacks

vintages for the same firm-year.

[11] Denmark, Ireland, and Portugal also have mandatory reporting requirements, but were unavailable to me due to data licensing restrictions and missing or outdated geographic identifiers.

developing country coverage.[12] BVD also reports additional branch locations and subsidiary ownership for many firms. I drop all firms that list subsidiaries or additional branches so that reported firm output aligns as closely as possible to my measure of temperature exposure at the main location. I also drop firms containing fewer than three observations and those with missing data for revenue or number of employees.

In total, the sample includes 17 countries that cover 59.4% of the world's manufacturing output and 51.1% of the global population.[13] The dataset also meets the globally representative criterion by spanning virtually the full range of climate and income levels in the global cross-section. According to the Penn World Tables, PPP-adjusted GDP per capita in my sample ranges from $1,137 in India in 1985 to $64,274 in Norway in 2014, which covers the 3rd to the 99th percentile of the global population in 2014. Similarly, country-level average daily maximum temperature in my sample ranges from 8.5 C°in Norway to 31.5 C°in India, covering the 1st to the 90th percentile of global population-weighted long-run temperature. Thus, to the extent that income and average temperature predict adaptation to extreme temperatures, my data is informative about the full range of heterogeneity in the global temperature-productivity relationship.

Climate Data

I use temperature data from Version 3 of the Global Meteorological Forcing Dataset (GMFD) produced at Princeton University. The data covers the entire world at a 0.25°by 0.25°grid for the years 1948-2016. GMFD is a reanalysis dataset that reconstructs historical temperature using a combination of observational data and local climate models. Following Graff Zivin and Neidell (2014) and other work on temperature and labor productivity, I use daily maximum temperature as my variable of interest to best approximate the temperature people experience during working hours.

I match firm and climate data at the county level. The government surveys provide county

[12] I drop firms marked mining, construction, utilities, and agriculture, though results are very similar when including these firms in the pooled sample.

[13] I cannot include the United States in my main pooled specification because I can only access the data at a secure government facility. I also exclude the data from China from my main specification for data quality reasons explained in Section 4.

location for each firm directly. The BVD data provides city name and zip code, which I match to the county-level using GeoPostcodes, a global geocoding dataset provided by GeoData Limited.[14] I apply nonlinear transformations to the GMFD temperature variable at the pixel level, and then average across pixels to the county level weighting by population.[15]

Other Data

I use purchasing power parity adjusted GDP per capita data from the Penn World Tables as a measure of the income level of each country-year in my sample.

3 Empirical Strategy

In order to quantify the effects of climate change on sectoral reallocation and aggregate productivity, my empirical results must execute three objectives. First, I need to estimate the causal effect of temperature on productivity in manufacturing and services. Second, I need to estimate the heterogeneity in that relationship such that I can predict the response to temperature for every country in the world. The model counterfactuals in Section 7 require an estimate of the response of manufacturing productivity to temperature in Algeria without having data from Algeria. Third, my estimates should incorporate the benefits and costs of adaptation. Future projections should reflect the fact that the effects of a given temperature realization will likely diminish as countries grow richer, firms improve technology, and agents adjust expectations to the shifting distribution of temperatures. To quantify the effects of climate change in Section 7, I need to make projections not just for Algeria today, but for future Algerian firms experiencing climate change in 2080.

[14] GeoData Limited estimates that their latitude and longitude coordinates for the center of each zip code are precise to within 100 meters. I independently verify a subset of observations in each country to ensure accuracy. I also hand-code a small number (under 1%) of unmerged observations using city name, and drop those unmerged observations for which the city name is non-unique within a country.

[15] For some countries, the administrative unit to which I aggregate is more comparable to a town than a county.

3.1 Conceptual Framework

To motivate my estimation strategy I start with a version of the production function from Burnside, Eichenbaum and Rebelo (1993) with variable labor effort:

$$Y = AK^{\alpha}(e * L)^{1-\alpha} \text{ with } 0 \leq e \leq 1 \tag{1}$$

The parameter e governs effective units of labor input. Intuitively, temperature could affect e through several channels. Extreme temperatures could cause illness or physical fatigue, impair cognitive function, or increase the disutility of labor such that workers reduce effort or minutes spent working.[16]

Rearranging the production function in terms of output per worker and taking logs gives:

$$ln\left(\frac{Y}{L}\right) = ln(e) + \left(\frac{1}{1-\alpha}\right)ln(A) + \left(\frac{\alpha}{1-\alpha}\right)ln\left(\frac{K}{Y}\right) \tag{2}$$

Equation 2 provides the basis for using output per worker as the dependent variable in my main specification. The change in output per worker equals the change in e when the firm's technology and capital-to-output ratio stay constant.[17] To gain further insight into the firm's optimal response to climate conditions, I model worker effort as a function of exposure to extreme heat (cooling degree days), extreme cold (heating degree days), and adaptation investments b_h and b_c:[18]

$$e^* = 1 - CDD * g_h(b_h) - HDD * g_c(b_c) \tag{3}$$

$$g \geq 0, g' < 0, g'' > 0$$

[16]The health effects of extreme temperatures have been widely documented, including in Deschênes and Greenstone (2011). Several laboratory experiments, including Seppanen, Fisk and Lei (2006) find evidence of reduced worker cognitive functioning. Graff Zivin and Neidell (2014) use time-use survey to show that people allocate less time to working in the presence of extreme temperatures.

[17]If capital is not adjustable in the short-run then short-run changes in $\frac{Y}{L}$ will slightly understate the change in e as $\frac{K}{Y}$ will also increase due to the fall in Y. In the long-run when the firm readjusts capital to its optimal level, the change in output per worker exactly equals the change in e.

[18]I define cooling degree days and heating degree days in Equation 6.

In this framework, the firm has access to separate technologies that mitigate the impact of extreme heat and extreme cold on worker effort with diminishing returns in each.[19] The first order conditions for a profit-maximizing firm yield the following expression for the firm's optimal investment in hot weather adaptation b_h:

$$-g'(b_h) = \frac{c_h * e}{p * MPL * L * CDD} \tag{4}$$

Since g is convex in b_h, Equation 4 predicts that firm adaptation investments will be increasing in the firm's exposure to extreme heat (CDD), the marginal product of labor, the firm's labor input, and the price of output, and decreasing in the cost of the adaptive technology, c_h, and the level of worker effort.[20] Thus, the firm's optimal condition predicts that worker effort will be less sensitive to temperature at more productive firms with more expected exposure to extreme temperatures, but that this reduced sensitivity comes at a cost.

To capture this heterogeneity, my empirical strategy focuses on modeling output per worker, and consequently e, as a function of temperature realizations, access to technology, and expectations over the distribution of temperature. By measuring the effects of climate change on e, I can use my estimates to project the change in the sector-by-country aggregate productivity parameters, Z_{jk}, that govern average output per worker in the model introduced in Section 5.

3.2 Causal Effect of Temperature

Following the framework outlined in Deryugina and Hsiang (2014), I start by noting that workers experience daily realizations of weather. San Francisco and Washington D.C. have similar annual temperatures, but very different exposure to extremes. To capture this logic, I treat daily output as a function of temperature on day d, $Y_d = f(T_d)$. To aggregate to annual output, the level of

[19]Zhang, Deschenes, Meng and Zhang (2018) mention that capital equipment could also perform poorly in extreme temperature conditions. If so, augmenting the production function with variable effective capital utilization, u, as in Burnside and Eichenbaum (1996), would capture this effect. In that case, the interpretation in Equation 2 would be that the reduction in $\frac{Y}{L}$ was attributable to a combination of declines in e and u.

[20]Optimal adaptation investment is decreasing in the level of worker effort because there are concave returns to effort.

my data, I sum daily outputs along with functions of daily temperature, $f(T_d)$, across all days experienced by firm i in year t:

$$Y_{it} = \sum_{d=1}^{365} Y_{id} = \sum_{d=1}^{365} f(T_{id}) = F(T)_{it} \tag{5}$$

Thus, I treat nonlinear transformations of daily temperature summed over the year as my primary independent variable of interest. Using annual data also has the important advantage of allowing for intertemporal substitution of labor. If workers produce less due to extreme temperatures on Tuesday but produce extra on Saturday instead, annual data captures the effects of temperature net of this reallocation.

For parsimony, my main specification uses a piecewise linear functional form for temperature, where output is allowed to vary linearly with daily maximum temperature above 30°C (CDD) and below 5°C (HDD):

$$f(T) = \begin{cases} \beta_1(5 - T_{max}) & \text{if } T_{max} < 5 \\ 0 & \text{if } 0 \leq T_{max} \leq 30 \\ \beta_2(T_{max} - 30) & \text{if } T_{max} > 30 \end{cases} \tag{6}$$

This formulation allows cold and hot temperatures to have separately estimated effects, β_1 and β_2, on productivity. I also conduct robustness checks with more flexible functional forms such as a polynomial of degree four and bins of daily maximum temperature.

Following other work in the climate impacts literature, I isolate the causal impact of temperature by exploiting interannual variation in weather. In line with the framework outlined in Section 3.1 my main specification models log output per worker at firm i in year t as a function of the vector of temperature effects, $\boldsymbol{\beta}$:

$$ln\left(\frac{Y_{it}}{L_{it}}\right) = \boldsymbol{\beta} F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \tag{7}$$

I control for permanent firm-specific features such as technology and management with firm fixed effects δ_i and for unobserved aggregate shocks such as technological progress and recessions with region (country or state) by year fixed effects κ_{rt}. I cluster my standard errors at the firm and county-by-year level to account for both serial and spatial correlation. Equation 7 allows for estimating the average treatment effect of temperature realizations, which fulfills part of the purpose of this section.

3.3 Heterogeneity and Adaptation

Following the strategy of Carleton et al. (2018), I allow for heterogeneity in the effect of temperature on output per worker by interacting the vector of temperature coefficients with income and long-run average temperature. This setup follows from the prediction in 3.1 that more productive firms in high-income countries and those that expect to experience extremes more frequently will be better adapted. I specify the interacted regression as follows:

$$\begin{aligned} ln\left(\frac{Y_{it}}{L_{it}}\right) &= \boldsymbol{\beta} F(T)_{it} + \boldsymbol{\gamma_1} ln(GDPpc)_{rt} \times F(T)_{it} \\ &+ \boldsymbol{\gamma_2} TMEAN_i \times F(T)_{it} + \delta_i + \kappa_{rt} + \epsilon_{it} \end{aligned} \tag{8}$$

The interaction variables in Equation 8 are country-level annual GDP per capita and long-run average daily maximum temperature in the county containing firm i.[21]

Estimating Equation 8 allows me to predict the treatment effects of extreme cold, β_1, and extreme heat, β_2, as a function of two factors - income and average climate. While there are certainly other variables that affect temperature sensitivity, this parsimonious specification makes it feasible to predict the treatment effects in any country for which I have data on GDP per capita and average temperature. Given the existence of this data for the full range of countries in the global cross-section, as well as of readily available plausible future projections of temperature change and economic growth, this approach allows me to project the effects of temperature

[21] I use country-level income because reliable data on subnational income is difficult to acquire. Average temperature is calculated as a 40-year average in the county of firm i, which is the same geographic scale at which contemporaneous temperature is measured.

both across space and over time. In line with the goals for this section, the interacted model allows me to predict the effects of temperature in Algeria today and in Algeria in 2080.

The coefficients on the interaction terms in Equation 8 are identified using cross-sectional, rather than panel, variation, but the identification assumption is also weaker. Estimating the main causal effect of temperature relies on the standard identification assumption - that the independent variable of interest is uncorrelated with omitted variables that affect output per worker conditional on the set of controls. For the interaction variables, however, I am interested in how income and climate *predict* temperature sensitivity, rather than in isolating their specific causal effect. Thus, the identification assumption is not that income and climate are uncorrelated with omitted variables affecting temperature sensitivity, but rather that this correlation remains constant across space and over time. Indeed, the aim is to use income and average climate as a proxy for the full suite of underlying mechanisms, and omitted variables, that govern adaptation. The cross-sectional approach will produce valid predictions if the effects of temperature realizations on output per worker in parts of the world with income levels and average temperatures similar to India are similar to the effects measured in India.[22]

Allowing the treatment effects of temperature to vary with long-run conditions also bridges the gap between weather and climate. A primary concern with using weather variation to inform estimates of the costs of climate change is that the estimated treatment effects may change as agents adjust their expectations in the long-run. I address this concern by explicitly modeling the treatment effects as a function of those expectations, as represented by long-run average temperature. In my formulation, climate is a distribution of temperatures and weather is a draw from that distribution. By allowing the treatment effect of a draw to depend on the distribution, my estimates for the effects of each draw remain valid as the distribution shifts. Intuitively, a hot day in Toronto could be more harmful than a hot day in Texas because it is more unexpected, but becomes less so as Toronto warms and its agents adapt. I capture this effect by assigning Toronto the estimated treatment effect of Texas once it has heated up to that long-run temperature in the

[22]Empirical estimation of adaptation in the climate impacts literature broadly relies heavily on cross-sectional variation because of the inherent difficulty in finding quasi-experimental variation in long-run conditions.

future.

4 Empirical Results

4.1 Main Regression Results

Table 2 contains the main results from estimating Equations 7 and 8. Column 1 displays the treatment effect of extreme temperatures for the average unit of output in the countries in my sample by weighting observations by country-level GDP and the inverse of each country dataset's sample size. While the estimated average treatment effects show that the effects of temperature are statistically different from zero, the magnitude of these coefficients is far too small to be economically significant. The estimates in Column 1 imply that a day with maximum temperature of either -5°C or 40°C would reduce annual output per worker by just 0.03% relative to a day in the moderate range of 5°C to 30°C.

Column 2 in Table 2 shows substantial heterogeneity in the effects of temperature on annual output per worker. Consistent with the approach taken in Carleton et al. (2018), I do not weight the regressions in which I model heterogeneity explicitly because the aim is to understand how the treatment effect varies across the full observed range of the interaction variables. The unweighted regression with differential sample sizes in different places also effectively allows areas with more data, and consequently more precise estimates of the effect of temperature, to contribute more to estimating the interaction terms.

The main effects of temperature in the unweighted interacted regression in Column 2 are large, negative, and precisely estimated, though the magnitudes cannot be interpreted without considering the interaction terms. The coefficients on both interaction terms for log GDP per capita are large and positive, indicating that richer countries are insulated from the effects of both extreme heat and cold. Consistent with intuition about adaptation to long-run conditions, the coefficient on the interaction term for average long run temperature is positive for hot extremes and negative for cold extremes, indicating that places are less susceptible to temperatures which they experience more frequently. All four interaction coefficients on income and average temperature are consistent with the predictions from Equation 4 - more productive

Table 2: Effects of Daily Temperature on Annual Revenue per Worker

	(1) Revenue/Worker	(2) Revenue/Worker	(3) Revenue	(4) Employment	(5) Revenue/Worker
TMax-30	-0.0000311	-0.00119	-0.00250	-0.00131	-0.00100
	(-2.29)	(-4.73)	(-6.80)	(-5.25)	(-4.03)
5-TMax	-0.0000315	-0.000956	-0.00180	-0.000842	-0.000452
	(-2.15)	(-2.15)	(-2.91)	(-1.92)	(-2.07)
(TMax-30) X log(GDPpc)		0.0000715	0.000178	0.000107	0.0000595
		(4.07)	(6.79)	(6.06)	(3.65)
(TMax-30) X $\overline{\text{TMax}}$		0.0000186	0.0000334	0.0000148	0.0000160
		(4.85)	(6.24)	(3.93)	(3.96)
(5-TMax) X log(GDPpc)		0.0000898	0.000167	0.0000769	0.0000416
		(2.14)	(2.85)	(1.85)	(2.02)
(5-TMax) X $\overline{\text{TMax}}$		-0.00000292	0.00000212	0.00000504	0.000000703
		(-1.54)	(0.93)	(2.85)	(0.59)
N	4125776	4125776	4125776	4125776	17938084
Manufacturing	X	X	X	X	X
Services					X
Firm FE	X	X	X	X	X
Country X Year FE	X	X	X	X	X
Inverse Sample Size Weights	X				
GDP Weights	X				
Countries Included	15	15	15	15	15

Notes: t-statistics in parentheses. Dependent variables all in logs. Standard errors are two-way clustered at the firm and county-by-year level. Column 1 shows the coefficients from estimating Equation 7 and Columns 2-5 show the results from Equation 8. Outcome variables come from the data sources listed in Table 1 and temperature data is from GMFD. Countries included are Austria, Belgium, Colombia, Finland, France, Germany, Greece, India, Indonesia, Italy, Norway, Spain, Sweden, Switzerland, and the United Kingdom. Section 4.3 shows results for the United States and Appendix C shows results for China.

firms with more exposure to given extremes invest more in adaptation.

Figure 3 shows the predicted effects of temperature from Column 2 of Table 2 at points across the distribution of observed income and climate levels in the world. Consistent with the results of the GDP-weighted regression in Column 1, the graphs show that temperature has little effect on productivity in rich countries (top row), with some effects from hot days in cold, rich places (top left cell) and mild effects from cold days in hot, rich places (top right cell).

Conversely, extreme temperatures have very large effects on productivity in poor countries

Figure 3: Predicted Heterogeneous Response of Annual Manufacturing Revenue per Worker to Daily Maximum Temperature

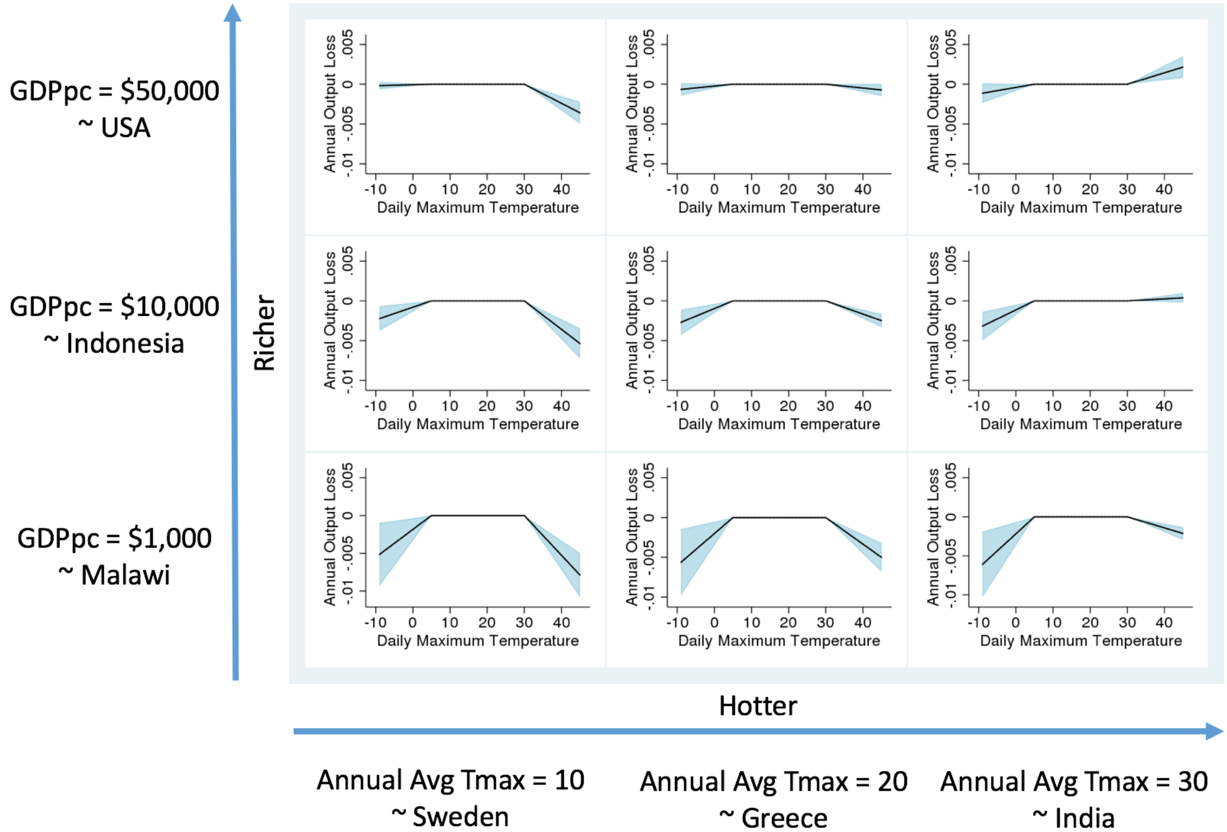

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 2 of Table 2.

(bottom row). Experiencing one day at -5°C or 40°C in a poor country with moderate long-run temperatures (bottom middle cell) reduces annual output per worker by about 0.4%. In a working year consisting of 50 work weeks of 5 days each, this is equivalent to each worker reducing production on that day to zero with no compensating substitution to other days. These effects in poor countries imply potentially large productivity costs from climate change in hot parts of the world in the absence of adaptation. In parts of Sub-Saharan Africa, climate change projections imply an increase in extreme heat on the order of moving 100 days per year from 30°C to 40°C by 2080, which would suggest substantial declines in manufacturing productivity in poor countries.

Columns 3 and 4 of Table 2 separately estimate the effects of temperature on revenue and employment. The effects of both hot days and cold days on revenue are substantially larger than those on revenue per worker because firms adjust employment in response to extreme temperatures. As shown in Appendix Figures A-1 and A-2, which again evaluate the predicted coefficients throughout the covariate space, these effects also primarily manifest only in poor countries. This finding is consistent with the firm's first order condition in the framework laid out in Section 3.1 - firms should be expected to reduce labor input in response to the fall in the marginal product of labor driven by a decline in e. However, it is perhaps surprising that firms in my sample do not face adjustment costs large enough to dissuade this adjustment in response to the short-run variation used to identify these effects.

Column 5 of Table 2 shows the effects of temperature on a pooled sample of manufacturing and services firms. The effects are very similar to the sample of only manufacturing firms in both magnitude and patterns of adaptation, with the exception of the finding that colder countries are less vulnerable to extremely cold temperatures. The sample size increases substantially in this specification because many of the firms in my data are services firms, though I do not have any services coverage in low-income countries.

4.2 Robustness

I conduct robustness checks with different ways to specify the functional forms of temperature. Appendix Figures A-3 and A-4 show the predicted effects from the main specification in Column 2 of Table 2 using bins and a polynomial of degree four in daily maximum temperature, respectively. The results are qualitatively very similar to the main specification.

I also show robustness to including more stringent state-by-year, rather than country-by-year, fixed effects. The results are very similar for specifications that use all the data (pooling manufacturing and services firms) with more flexible functional forms such as bins or a polynomial of degree four. These two specifications are shown in Appendix Figures A-6 and A-7. These results are sensitive to functional form, however. The more parsimonious functional forms with a single parameter each governing the response to cold days and hot days show muted effects, particularly in the specification with manufacturing firms only. This is consistent with the fact that considerably less variation in temperature realizations remains within states in a given year, so more data and flexible estimation is necessary to recover the underlying pattern.

Figure A-8 shows robustness to including controls for capital. While the standard errors for this specification are somewhat larger because I lack data on capital for approximately a quarter of the observations in the main specification, the pattern of predicted effects is very similar.

4.3 U.S. Results

In this section, I use separate estimates of the effect of extreme temperatures on manufacturing in the United States to externally validate the results in Section 4.1.[23] Predictions using the global interacted regression suggest that temperature has a negligible effect on annual manufacturing revenues in rich, temperate countries such as the U.S. (see the top middle cell of Figure 3). Figure 4 shows the treatment effect of temperature on annual manufacturing revenue per worker estimated on data from the U.S. Census Bureau:

Figure 4: Estimated Response of U.S. Annual Manufacturing Revenue per Worker to Daily Maximum Temperature

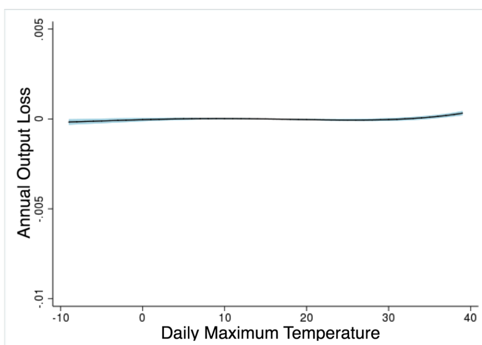

Notes: Figure shows the response of annual revenue per worker to a polynomial of degree four in daily maximum temperature estimated using Equation 7. Outcome variable data comes from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD. Standard errors are two-way clustered at the firm and county-by-year level.

[23]The results in Section 4.1 do not include data from the United States due to physical constraints on data access. Plant-level manufacturing data from the United States Census Bureau must be analyzed at restricted access Federal Statistical Research Data Centers (RDC).

Consistent with predictions from global data in Figure 3, I find a precisely estimated null effect of temperature on output-per-worker in the U.S.[24] The U.S. data also includes information on other inputs that I lack in my global sample, allowing me to directly observe some of the adaptation costs incurred by U.S. firms. Appendix Figure A-13 shows that the average U.S. plant increases expenditures on electricity and other fuels by several thousand dollars for each extremely hot and cold day, presumably for cooling and heating expenses.[25] These expenditures are small in the context of U.S. plant size, however, such that temperature still has a null effect on revenue total factor productivity, which accounts for expenditures on energy and materials, as shown in Figure A-12.

4.4 Projected Global Sensitivity to Extreme Temperatures

To connect the regression results from this section with the model presented in Section 5, I predict the effects of temperature in all 158 countries for which I will estimate the model. Figure 5 shows the predicted effects of a day with maximum temperature of 40°C on annual manufacturing revenue per worker and Figure 6 shows the effect of a -5°C day. Consistent with intuition about adaptation and the results displayed in Figure 3, poor countries and those which experience given temperatures less frequently are more susceptible to extreme realizations.[26]

Projecting the impacts of climate change also requires accounting for adaptation by adjusting the temperature sensitivities shown in Figures 5 and 6 to projected changes in long-run average temperature. The firm's optimal adaptation decision in Equation 4 implies that firms will increase investment in protection from extreme heat as the climate warms. I account for the benefits of these investments by reevaluating predicted heat sensitivity at projected end-of-

[24]The result displayed in Figure 4 uses a polynomial of degree four in daily maximum temperature, but the null result is robust to choice of functional form. Appendix Table A-1 shows a range of specifications, all of which are consistent with a null effect on output and employment.

[25]Total energy expenditures are defined as the sum of electricity expenditures and the cost of other fuels. Full results for this outcome variable are shown in Appendix Table A-2.

[26]Note that following Carleton et al. (2018), these predictions define full adaptation as productivity that is invariant to temperature, and thus do not allow the effect of extreme temperatures to go above zero. The effects of extreme temperatures are weakly negative in the range of incomes and climates in the sample used for estimation, and I maintain this pattern as incomes and temperatures go out of sample.

Figure 5: Predicted Effect of a 40°C Day on Annual Manufacturing Revenue per Worker

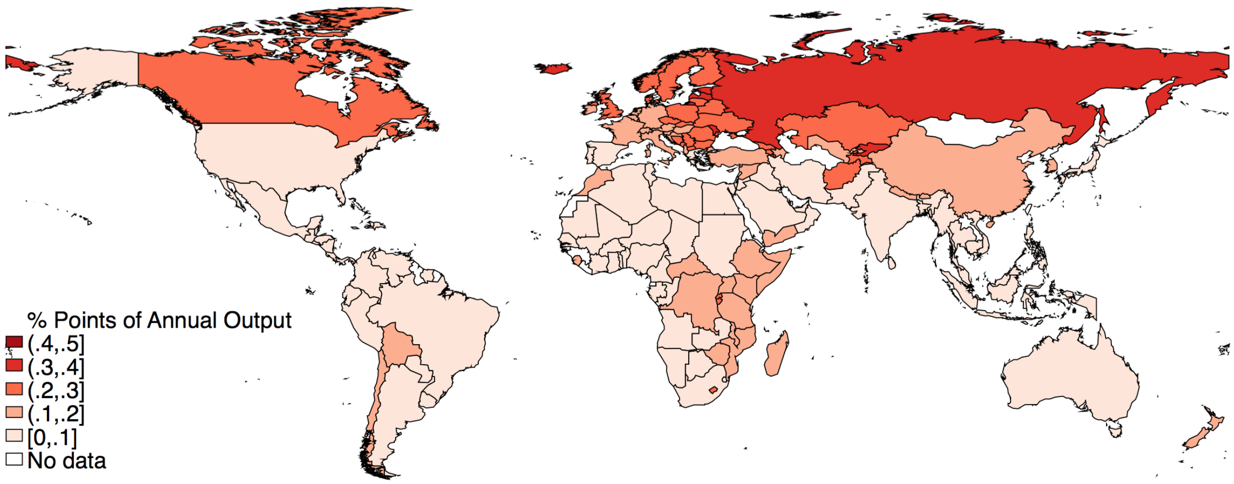

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's level of income and long-run average temperature.

Figure 6: Predicted Effect of a -5°C Day on Annual Manufacturing Revenue per Worker

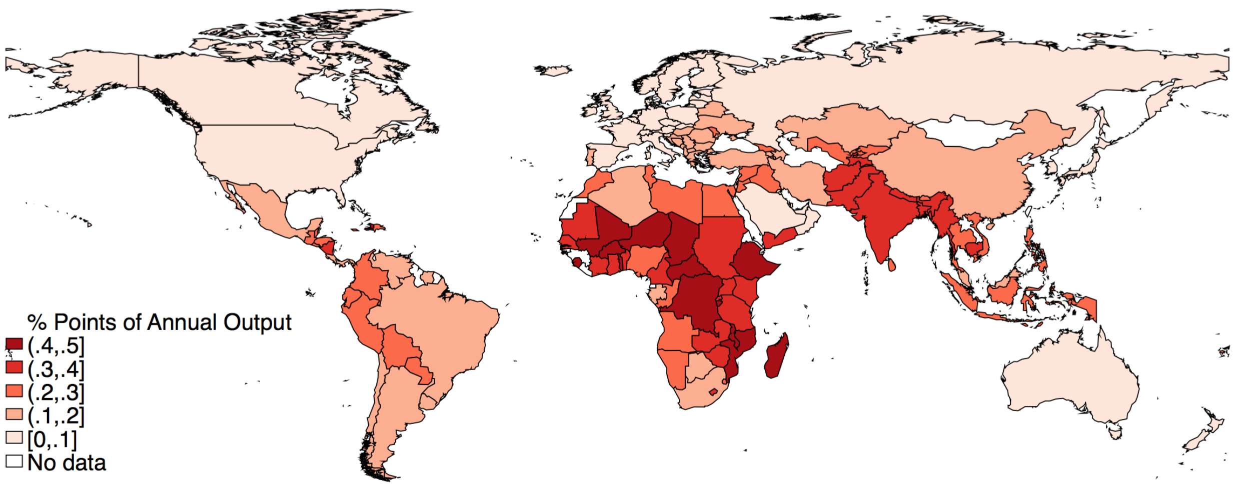

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a -5°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's level of income and long-run average temperature.

century temperatures in Appendix Figure A-17.[27] The results show noticeably muted effects when allowing for expectations to adjust to future temperatures. The mean global damage from a 40°C day is about 34% lower when evaluated at future temperatures (0.067% of annual revenues versus 0.1%) and firms in 67 countries become invariant to hot days compared with 39 countries at current temperatures.

The adaptation benefits of adjusting to extreme heat come at a cost. If it were costless to protect production from extreme heat, no firms would show effects of temperature on productivity. Instead, my results show that firms which experience given extremes infrequently find it optimal to invest less in adaptation, implying that the costs they would incur to achieve a marginal reduction in temperature sensitivity exceed the benefits. I leverage this intuition combined with the firm's first order conditions in Section 3.1 to infer a revealed preference measure of these adaptation costs following methods developed in Carleton et al. (2018). Appendix D covers the details of this calculation.

Quantifying the aggregate productivity consequences of climate change also requires projecting temperature sensitivity in services. I make projections for services using the pooled sample of manufacturing and services firms due to my lack of services data coverage in poor countries.[28] This choice follows from the estimated strong gradient of temperature sensitivity with respect to income but very similar coefficients between the manufacturing only and manufacturing/services pooled specifications in Columns 2 and 5 of Table 2.[29] Intuitively, my results suggest that manufacturing firms in India are a better proxy for services firms in India than services firms in Germany would be. Appendix Figures A-20 and A-21 show predicted current global

[27] End-of-century temperature projections are the 30-year average of annual average maximum temperature from the climate model predictions used in Section 7.1. In Section 7.6 I also allow for economic growth to make countries richer in the future, further reducing their temperature sensitivity.

[28] I show prediction results for regressions using only services firms in Appendix Figures A-9, A-10, and A-11. The results for extreme heat with more flexible functional forms such as a fourth degree polynomial are qualitatively similar to those of the pooled manufacturing and services regression, but these specifications are sensitive to functional form. Furthermore, the predictions in poor countries are extrapolating far out of the sample, which only includes European firms in a narrow range of high income levels.

[29] A formal test shows that coefficients for manufacturing and services firms in the pooled regression have statistically indistinguishable responses to extreme heat and marginally significant evidence that services firms are less susceptible than manufacturing firms to extreme cold.

sensitivity to hot and cold days in services using results from the pooled regression. I follow the same procedure to account for future adaptation benefits and costs as in manufacturing.

Overall, the results in this section allow me to predict the sensitivity of non-agricultural firm output per worker to extreme temperatures in every country in the world in the present and future. I use these results to project the impact of climate change on global comparative advantage between agriculture and manufacturing in Section 7.1, and to simulate the corresponding changes in sectoral allocation and aggregate productivity.

5 Model

This section lays out a static general equilibrium model of global production, consumption, and trade in agriculture, manufacturing, and services to analyze how changes in sectoral productivity affect sectoral specialization, trade flows, aggregate productivity, and welfare. I show that the model makes ambiguous predictions about how reductions in agricultural productivity affect the labor share of agriculture, and that openness to trade is a key determinant of the aggregate consequences of asymmetric sectoral productivity shocks.

The ingredients of the model are as follows:

5.1 Model Ingredients

Consumption Following the demand system specified in Comin, Lashkari and Mestieri (2015), consumers in each country gain utility from final goods in each of the three sectors - agriculture, manufacturing, and services - according to the following implicitly defined utility function:

$$\Omega_a^{\frac{1}{\sigma}} U^{\frac{\epsilon_a}{\sigma}} C_a^{\frac{\sigma-1}{\sigma}} + \Omega_m^{\frac{1}{\sigma}} U^{\frac{\epsilon_m}{\sigma}} C_m^{\frac{\sigma-1}{\sigma}} + \Omega_s^{\frac{1}{\sigma}} U^{\frac{\epsilon_s}{\sigma}} C_s^{\frac{\sigma-1}{\sigma}} = 1 \tag{9}$$

Here, $\{\epsilon_a, \epsilon_m, \epsilon_s\}$ are utility elasticities for each sector that allow for nonhomothetic preferences, $\{\Omega_a, \Omega_m, \Omega_s\}$ are fixed sectoral taste parameters, and σ is the cross-sector elasticity of substitution. I choose this nonhomothetic CES preference specification because it can closely match the observed pattern of smooth structural transformation out of agriculture.[30]

[30]Generalized Stone-Geary preferences, another common specification used to incorporate nonhomotheticity, diverge from the data by predicting sharp declines in the share of agricultural consumption

Households consume their full wage, w, which varies at the level of country k. The aggregate budget constraint, summed across the country-level population L_k, equates income to total expenditures across the three sectors:

$$P_{ak}C_{ak} + P_{mk}C_{mk} + P_{sk}C_{sk} = w_k L_k \tag{10}$$

Demand for the final good in sector j in country k is given by:

$$C_{jk} = \Omega_j \Big(\frac{P_{jk}}{w_k}\Big)^{-\sigma} U^{\epsilon_j} \tag{11}$$

Production

The final good in sector j in country k is a CES composite of intermediate varieties indexed by i:

$$Y_{jk} = \left(\int_0^1 y_{ijk}^{\frac{\eta-1}{\eta}} \, di \right)^{\frac{\eta}{\eta-1}} \tag{12}$$

Intermediate goods producers each receive a productivity draw, z_{ijk}, drawn from a Frechet distribution with sector-specific shape parameter θ_j and sector-country specific start value Z_{jk}. The production function for intermediate goods is linear in labor:[31]

$$y_{ijk} = z_{ijk} * l_{ijk} \tag{13}$$

$$z_{ijk} \sim F_{jk} \text{ where } F_{jk}(z_i) = exp(-Z_{jk} z^{-\theta})$$

$$\text{and } Z_{jk} = f(\mu_{jk}, T_{jk}, E(T_{jk})) \tag{14}$$

The sector-country specific aggregate productivity parameters, Z_{jk}, connect the model to my empirical results in Section 4. In particular, I allow Z_{jk} to be a function of temperature realiza-

at low income levels and converging toward homothetic preferences at middle and higher income levels. Though the model fit is weaker, my results and the implications for sectoral reallocation are qualitatively robust to using Stone-Geary preferences.

[31] Excluding capital from the model is implicitly equivalent to assuming freely mobile and undistorted capital markets around the world. In future drafts, I plan to conduct a robustness check with land included as an input.

tions, T_{jk}, expectations over temperature, $E(T_{jk})$, and a vector, μ_{jk}, of country-sector specific features such as technology, institutions, and human capital. In making future projections in Section 7, climate change enters the model by perturbing the vector of Z_{jk} with empirically estimated productivity impacts that vary at the country-sector level.

Trade

The trade portion of my model follows Eaton and Kortum (2002). When selling to foreign countries, intermediate goods producers face an iceberg trade cost, τ_{ijk}, that varies at the exporter-importer-sector level. So, intuitively, shipping food from Canada to Malawi incurs a different trade cost than shipping food from Malawi to Canada, and manufactured goods shipped between Canada and Malawi have two separate trade costs of their own. Services are nontradable.

Intermediate goods producers price at marginal cost. Since labor is the only input, the price of a domestically produced good in country k is given by $p_{ijk} = \frac{w_k}{z_{ijk}}$. When selling to foreign country n and incurring the cost of trade, the intermediate goods producer in country k prices as follows:

$$p_{ijk} = \frac{\tau_{jkn} w_k}{z_{ijk}} \tag{15}$$

This representation of trade incorporates Ricardian comparative advantage both within and across sectors. A producer's ability to sell competitively priced exports depends both on their productivity and on the domestic wage. Low productivity countries will have low wages in equilibrium, so their relatively productive producers will be able to export their products even if their absolute productivity is low. Thus, relative productivity between sectors is the key determinant of net imports and exports.

The final goods producer sources each variety from the lowest-priced producer. The sectoral final goods prices are given by the CES price index of all intermediate varieties used in that sector:

$$P_{jk} = \left(\int_0^1 p_{ijk}^{1-\eta} di \right)^{\frac{1}{1-\eta}} \tag{16}$$

Intuitively, the price of the final good in agriculture, P_{ak}, can be thought of as a price index for the complete basket of food items while the price of each individual variety, p_{iak}, is the price of one particular food, such as apples. η is the elasticity of substitution between varieties.

Equilibrium

The model has two equilibrium conditions. First, total income in country k is the sum of all domestic and foreign sales in all three sectors.

$$w_k L_k = \sum_{j=1}^{3} \left(\pi_{jkk} P_{jk} C_{jk} + \sum_{n \neq k}^{N} \pi_{jkn} P_{jn} C_{jn} \right) \tag{17}$$

Here, π_{jkn} is the share of varieties from sector j consumed in country n that country k produces. So country k receives income both from its production share of domestic consumption in sector j, and from the share of consumption in every foreign country comprised of its exports. Since consumption equals income in each country, this condition also ensures that trade balances.

The second equilibrium condition concerns the labor market. The total labor force is allocated across the three sectors:

$$L_k = L_{ka} + L_{km} + L_{ks} \tag{18}$$

In autarky, market-clearing requires that income equals expenditures in each sector, $P_{jk} C_{jk} = w_k L_{jk}$, which means that the labor share, l_{jk}, equals the expenditure share, X_{jk}. In the open-economy case, the labor share equals the production share of revenues in each sector, incorporating net exports. This gives the following equation from Uy, Yi and Zhang (2013):

$$l_{jk} = \pi_{jkk} X_{jk} + \sum_{n=1}^{N} \pi_{jkn} X_{jn} \frac{w_n L_n}{w_k L_k} \tag{19}$$

This condition illustrates the importance of both domestic consumer preferences and international trade in determining the sectoral allocation of labor. Intuitively, Equation 19 says that if country k has agricultural consumption worth 30% of spending and agricultural net exports worth 10% of GDP, then 40% of its labor force will be in agriculture.

Aggregate GDP Losses and Willingness-To-Pay

I calculate the willingness-to-pay to avoid climate change productivity impacts as equivalent variation using the nonhomothetic measure of utility from the Comin, Lashkari and Mestieri (2015) preference specification.

I also quantify the aggregate GDP effects of sectoral productivity changes by using a Törnqvist (1936) price index that uses sectoral expenditure shares from before and after the shock, (X_{jk0} and X_{jk1}), to construct an aggregate price index with which to deflate nominal income:

$$P_k^T = \prod P_{jk}^{(X_{jk0}+X_{jk1})/2} \longrightarrow GDP_k = \frac{w_k L_k}{P_k^T} \tag{20}$$

This captures the logic of Baqaee and Farhi (2017), who extend Hulten (1978) to show that the aggregate productivity impact of a sectoral shock is given by the weighted average of the pre and post-shock sectoral shares. The intuition here is simple. If productivity falls markedly in agriculture, the aggregate impact is accentuated if more of the economy moves into agriculture and tempered by reallocation to other sectors. Thus, quantifying the magnitude and direction of sectoral reallocation is a key part of estimating the aggregate productivity consequences of climate change.

5.2 Comparative Statics

I now use the model to characterize the factors that influence sectoral reallocation in response to climate change. Consider a country that suffers an agriculture-biased reduction in aggregate productivity, consistent with projections for hot parts of the world made in Section 7. To see how the labor share in agriculture changes in Equation 19, I first consider the impact on the agricultural expenditure share, X_{ak}. The expression for X_{ak} from solving the consumer's problem is as follows:

$$X_{ak} = \Omega_a \left(\frac{p_{ak}}{P_k}\right)^{1-\sigma} \left(\frac{w_k}{P_k}\right)^{\epsilon_a-(1-\sigma)} \tag{21}$$

Taking logs gives:

$$logX_{ak} = log(\Omega_a) + \underbrace{(1-\sigma)log\left(\frac{p_{ak}}{P_k}\right)}_{\text{Substitution Effect}} + \underbrace{(\epsilon_a-(1-\sigma))log\left(\frac{w_k}{P_k}\right)}_{\text{Income Effect}} \tag{22}$$

The agriculture-biased reduction in productivity has two effects that appear in Equation 22.[32] First, the reduction in productivity drives down the equilibrium real wage ($\frac{w_k}{P_k}$), making consumers poorer. If $(\epsilon_a - (1-\sigma)) < 0$, as is the case with the parameter estimates presented in Section 6, then the reduction in real wage drives up the expenditure share on food, X_{ak}. This is the effect of nonhomotheticity. Food is a larger share of consumption for poorer people, so climate change tends to drive up the share of agricultural consumption by making people poorer.

Second, the relative decline in agricultural productivity will increase the domestic price of agricultural goods relative to the aggregate price index ($\frac{p_{ak}}{P_k}$).[33] If $\sigma < 1$, as is also the case in Section 6, then the rising relative price of agricultural goods raises the expenditure share on agriculture. Intuitively, if food is not substitutable with other consumption, then its relative quantity falls less than the relative price rises, and the share of spending on food goes up. This is the same logic that underlies Baumol's cost disease (Baumol and Bowen, 1966), a theory that endeavors to explain why low-substitutability service sectors with relatively low productivity growth, such as health care and education, tend to rise as a share of expenditures over time.

Together, nonhomotheticity and low substitutability at the sector level combine to push up

[32]This equation also appears in Comin, Lashkari and Mestieri (2015). They estimate that nonhomotheticities (the income effect) account for about 75% of observed historical structural transformation, with changes in relative prices (the substitution effect) accounting for the rest.

[33]In a closed economy, relative sectoral prices are exactly proportional to sectoral productivities. In an open economy, the domestic relative price of agriculture responds to domestic agricultural productivity in proportion to the domestic share of consumption.

the expenditure share on agriculture in response to declines in agricultural productivity. The macro-development literature on structural transformation (see, for instance, Gollin, Parente and Rogerson (2007)) refers to these features of consumer preferences as 'the food problem' - the explanation given to the large share of the labor force in agriculture in most developing countries despite very low absolute and relative productivity.

These features of the model also explain why my model's predictions about the protective effects of reallocation diverge from those of Costinot, Donaldson and Smith (2016). Their paper finds that reallocating production across crops reduces the aggregate damages from climate change by two-thirds. To capture reallocation at the crop level, their model has no income effects and high substitutability across products.[34] This specification makes sense for capturing reallocation across crops, but does not generalize to the cross-sector case where income effects become important and the elasticity of substitution is very low. Intuitively, if the productivity of corn falls markedly relative to the productivity of wheat, consumers can respond by eating more wheat. If the productivity of producing food falls relative to the productivity of manufacturing, however, consumers cannot subsist by eating more manufactured goods.

In contrast to the food problem, the Ricardian comparative advantage effects of falling relative productivity in agriculture will tend to push labor into other sectors. Returning to Equation 19, shifting comparative advantage away from agriculture will tend to push up food imports (π_{akk} falls for country k) and push down food exports (π_{akn} falls). Equation 23 captures the horserace between the food problem and international trade that drives general equilibrium sectoral reallocation in response to climate change.[35]

$$l_{ak} = \underbrace{\pi_{akk}}_{\downarrow}\underbrace{X_{ak}}_{\uparrow} + \underbrace{\sum_{n\neq k}^{N} \pi_{akn} X_{an} \frac{w_n L_n}{w_k L_k}}_{\downarrow} \tag{23}$$

[34]They estimate an elasticity of substitution of 5.4 across varieties of the same crop and 2.82 across crops. I estimate an elasticity of 0.29 between sectors.

[35]The importance of trade for promoting structural transformation out of agriculture has been previously emphasized by Tombe (2015), Teignier (2018), and Uy, Yi and Zhang (2013).

In autarky, falling relative agricultural productivity would drive up the labor share in agriculture, exacerbating the aggregate productivity costs. In an economy with costless trade, climate change would dramatically shift the global geography of agricultural production and trade flows, substantially limiting the aggregate costs. To quantify the relative strength of these effects in practice, I need to estimate the parameters of the model and simulate the general equilibrium response to the estimated impacts of climate change on productivity at the country-sector level.

6 Model Estimation

6.1 Parameter Estimates

I estimate the model presented in Section 5 to match data from 158 countries on sectoral GDP shares, bilateral trade flows in agriculture and manufacturing, and value-added per worker. Table 3 shows a list of the target moments and data sources corresponding to each model parameter. For the trade data obtained from UN Comtrade, I classify HS 1988/92 codes 1-24 as agriculture and 28-97 as manufacturing to best approximate food and non-food imports.[36]

I estimate consumption parameters and trade costs using simulated method of moments.[37] To assign values to Z_{jk}, I choose country level relative sectoral productivities to match the ratio of value-added per worker in agriculture, manufacturing, and services, and adjust the overall level of $\{Z_{ak}, Z_{mk}, Z_{sk}\}$ to match country-level nominal GDP.[38] I calibrate the trade elasticities using the values estimated by Tombe (2015); $\theta_a = 4.06$, and $\theta_m = 4.63$.

Table 4 displays my estimates of the preference parameters for the nonhomothetic CES utility specification. Two points about these estimates are worth noting. First, I estimate a cross-

[36]Since trade data is reported in gross output terms but GDP is in value-added, I deflate the trade data by country-sector-level value-added to output ratios obtained from the United Nations Statistical Division. Following recommendations from UN Comtrade documentation, I use importer-reported trade data where possible, but default to exporter-reported data for smaller developing countries with large discrepancies between importer and exporter reported data.

[37]To simulate the model, I directly draw productivities from the Frechet distributions for 20,000 varieties for each sector for each country. I assign the production of each variety in each country to the lowest cost producer based on wages, trade costs, and productivity. I then find the vector of wages under which the equilibrium condition holds and national income equals national spending for every country. I estimate the consumption parameters to match sectoral share data using the patternsearch algorithm in Matlab, and choose bilateral trade costs to match the data on bilateral trade flows by sector.

[38]Since trade flows are in nominal terms, I match nominal GDP in the model for consistency. The nonhomothetic price index deflates nominal income to a measure of welfare.

Table 3: Model Parameters and Target Moments

Parameters	Data Moment	Data Source
σ	Sectoral GDP Shares	World Bank
$\Omega_a, \Omega_m, \Omega_s$	Sectoral GDP Shares	World Bank
$\epsilon_a, \epsilon_m, \epsilon_s$	Sectoral GDP Shares	World Bank
θ_a, θ_m	Calibrated from Tombe (2015)	
τ_{jkn}	Trade Flows	UN Comtrade
Z_{jk}	Sectoral Value-Added per Worker	World Bank
L_k	Population	World Bank

Notes: Table shows the data sources for moments targeted in my simulated method of moments procedure to estimate parameters for the model presented in Section 5. Data is for the global cross-sectoin in 2011, accessed from the World Bank Databank.

sector elasticity of substitution, $\sigma = 0.27$, of substantially less than one, indicating that the expenditure share in a sector sharply increases with its relative price. My estimate of σ to target the global cross-section of sectoral shares matches up well with that of Comin, Lashkari and Mestieri (2015), who use various historical panel datasets to estimate σ between 0.2 and 0.6. Second, I estimate that $\epsilon_a - (1-\sigma) = -0.44$, which implies from Equation 22 that the consumption share of agriculture is strongly diminishing in real income. Thus, my parameter estimates imply clearly that a decline in aggregate productivity concentrated in agriculture will raise the expenditure share of agriculture through both the income and substitution effect.

6.2 Model Fit

The model closely matches the features of the data most relevant to the counterfactual simulations of the impacts of climate change. Table 5 summarizes the correlation between key simulated moments in the model and their empirical counterparts.[39] I match the income level of each

[39] A coefficient of 1 with $R^2 = 1$ would constitute a perfect fit. The fit for other moments in the model is displayed in Appendix Figures A-27 to A-32.

Table 4: Parameter Estimates

Parameter	Description	Estimate
σ	Cross-Sector Elasticity of Substitution	0.27
ϵ_a	Agriculture Utility Elasticity	0.29
ϵ_m	Manufacturing Utility Elasticity	1
ϵ_s	Services Utility Elasticity	1.15
Ω_a	Agriculture Taste Parameter	11.73
Ω_m	Manufacturing Taste Parameter	3.70
Ω_s	Services Taste Parameter	10

Notes: Parameters estimated using simulated method of moments. Ω_s is normalized to 10 as only relative values of Ω_j affect consumer choices. Since the focus is on cross-sector reallocation, I set the elasticity of substitution across varieties, η, equal to 1 for tractability so that varieties have equal revenue shares.

country almost exactly by scaling the country-level aggregate productivity parameters. Similarly, my simulations closely match the domestic production share of agricultural consumption since I choose exporter-importer-sector-specific trade costs, τ_{jkn}, to match all observed bilateral trade flows.[40] As shown in Appendix Figure A-33, most developing countries import little of their food. In the data, the average person in the poorest quartile of the world consumes 91.3% domestically produced food (89.4% in the simulation) compared to 45.1% in the richest quartile (52.4% in the simulation). I present suggestive evidence on some of the underlying causes of these high barriers to trade in poor countries in Section 9.

My model also explains most of the variation in the global agriculture share of GDP. I slightly under-predict agricultural shares on average, but overall the model explains 60.3% of the variation in the data. This is a relatively strong fit considering that only the seven free parameters in Table 4 were chosen to match 316 independent target moments consisting of GDP shares for agriculture, manufacturing, and services in 158 countries. As shown in Figure 7, the nonhomo-

[40]The simulated domestic production shares of expenditures have no systematic bias, but explain only 81% of the variation in the data because some countries have imbalanced trade.

Table 5: Summary of Model Fit

	(1) Data log(GDP per capita)	(2) Data Ag Share of GDP	(3) Data π_{akk} (Ag Domestic Production Share)
Simulated log(GDP per capita)	1.006		
	(0.00251)		
Simulated Ag Share of GDP		0.866	
		(0.0563)	
Simulated π_{akk}			1.009
(Ag Domestic Production Share)			(0.0392)
Observations	158	158	158
R^2	0.999	0.603	0.809

Notes: Table shows the results from regressing empirical moments in the data on their simulated counterparts. Data on nominal income levels and the agriculture share of GDP are from the World Bank. Data on the domestically produced share of expenditures in agriculture is constructed using Comtrade data.

Figure 7: Agriculture Share of GDP - Data vs. Simulation

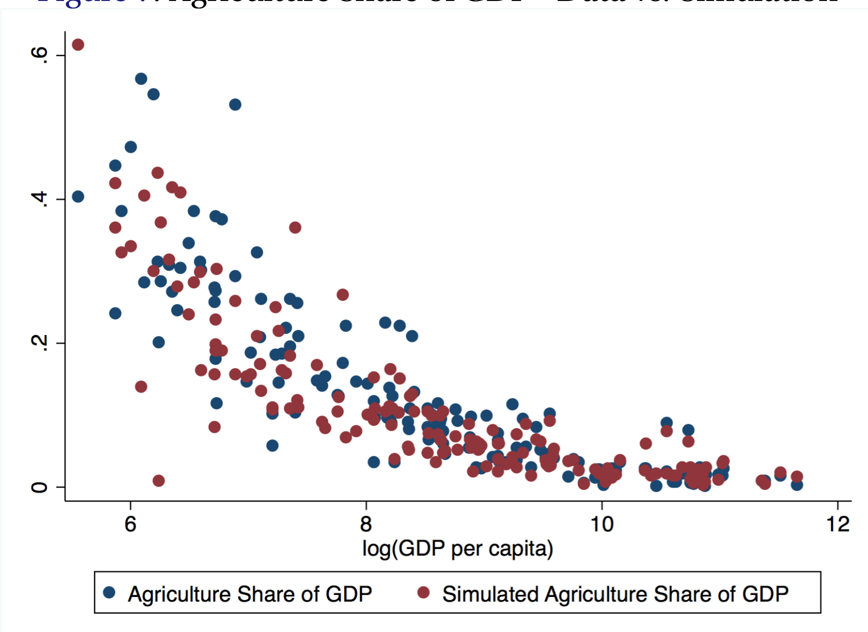

Notes: Graph shows the fit of simulated agriculture share of GDP in the model to data from the World Bank. The simulation explains over 60% of the variation in the data, and reproduces the smooth pattern of non-homotheticity observed in the empirical relationship between agriculture shares and income.

thetic CES demand specification enables the simulation to closely mirror the smooth decline of agricultural GDP with log income per capita.[41]

The model also reproduces the general pattern of high relative prices for agricultural consumption in poor countries - a moment I do not target in my estination. In Figure 8, I compare the simulated pattern of the relative price of agricultural and manufacturing consumption, P_{ak} and P_{mk}, to an empirical analogue constructed using aggregate sectoral price indices from the World Bank's International Comparison Program. While the simulated and empirical price indices have different units that prevent direct comparison, they share the same pattern of high relative prices for food in developing countries with low relative agricultural productivity.

Figure 8: Relative Price of Food - Data vs. Simulation

The graph on the left shows the ratio of a country-level food price index to an aggregate price index using data from the International Comparison Program. The graph on the right shows an analogous moment in the model - the ratio of the aggregate agricultural and manufacturing price indices, P_a and P_m. The model reproduces the empirical relationship that poor countries tend to have higher relative prices for food - a moment I do not target in my estimation.

Overall, the model matches the existing global pattern of sectoral specialization through a combination of consumer preferences and barriers to trade. Low incomes and the high relative

[41]For comparison, the best fit using a Stone-Geary utility specification has an R^2 of 0.43 and predominantly underpredicts the agriculture share as shown in Appendix Figure A-34.

price of food drive up agriculture's share of expenditures in poor countries through the nonhomotheticity and low elasticity of substitution in the preference specification. High estimated trade costs chosen to rationalize observed trade flows tightly link domestic consumption to domestic production, causing many developing countries to specialize in agriculture despite its low relative productivity.[42] In the next section, I use the model to investigate projected sectoral reallocation and its welfare consequences in response to climate change.

7 Model Counterfactuals

This section uses the estimated model to project the impacts of climate change on trade flows, sectoral specialization, prices, GDP, and welfare.

7.1 Estimated Productivity Impacts

I start by projecting the impacts of climate change on country-sector level productivity. For agricultural productivity effects, I use the estimates from Cline (2007) displayed in Figure 1. This analysis uses micro-data from 18 countries in Africa, North and South America, and Asia representing over 35% of the world's agricultural production to estimate Ricardian cross-sectional regressions of agricultural output (in dollars) from grains, fruits, vegetables, and livestock as a function of temperature, precipitation, and irrigation. Because we expect farmers to have optimized crop choice and land use decisions in response to local long-run climate conditions, I interpret the estimated effects of temperature and precipitation from these cross-sectional regressions as net of adaptation through choice of crops and livestock. Projections using the empirical estimates are averaged with projections from leading crop models from agronomy, which also account for adaptation through crop-switching and adjusted farming techniques.[43] I use Cline (2007) in my analysis because it uses globally representative data to produce results broadly consistent with the literature on climate and agricultural production, and represents the most comprehensive available source of global impact estimates that account carefully for

[42]As discussed in Section 5, this explanation is consistent with the work of Tombe (2015), Gollin, Parente and Rogerson (2007), and the broader literature on structural transformation.

[43]The crop model projections in Cline (2007) account for reallocation across crop types within country, shifting planting dates, and increased irrigation and fertilizer use. None of the estimates in the analysis account for any response of international trade.

adaptation within the agricultural sector.

Figure 9: Projected Impact of Climate Change on Manufacturing Productivity

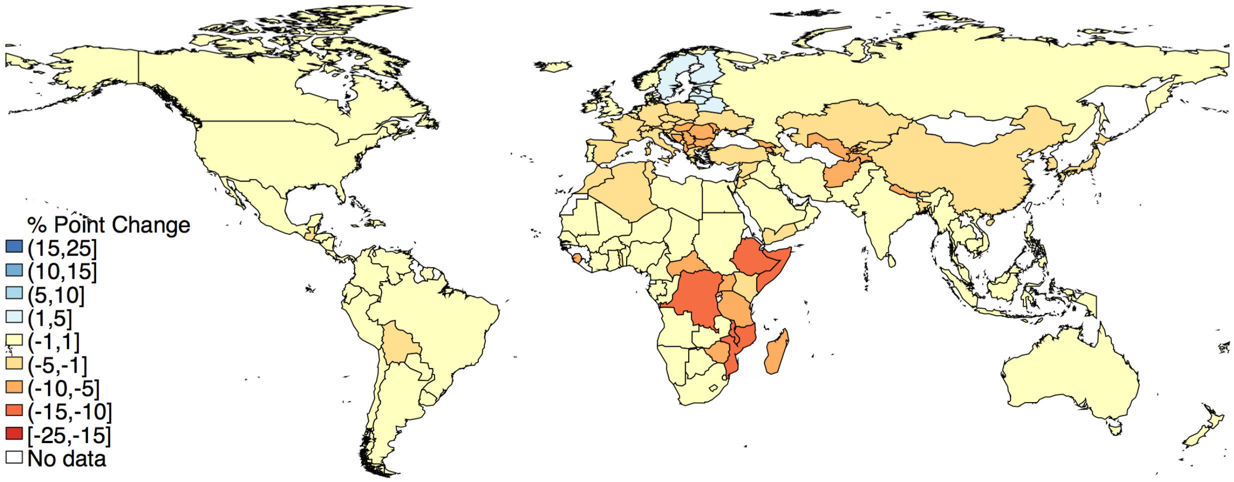

Notes: Map shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 2 of Table 2 at each country's income and end-of-century long-run average temperature.

Figure 10: Projected Impact of Climate Change on Agricultural Relative Productivity

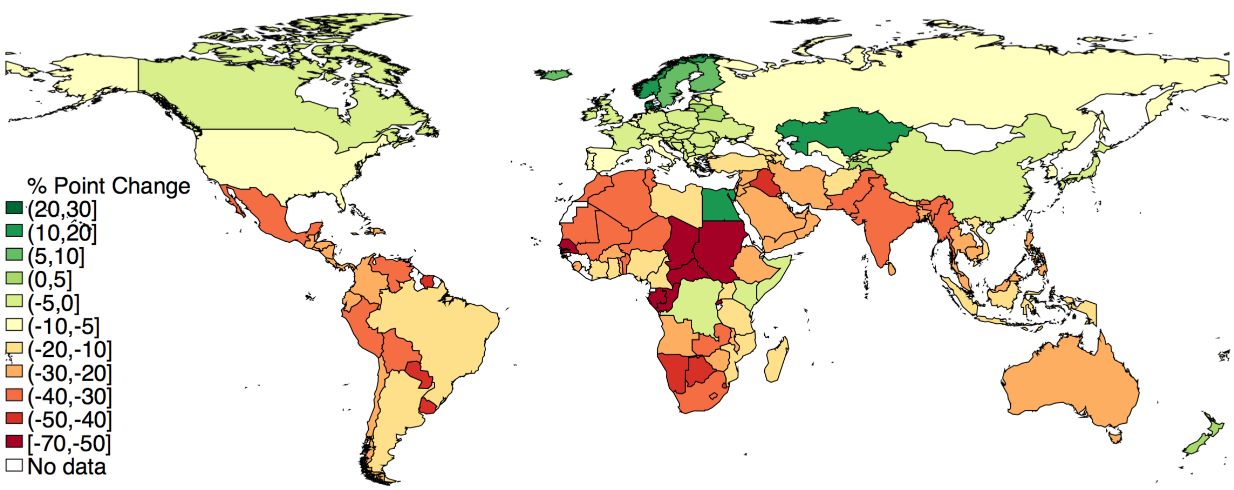

Notes: Map shows the change in agricultural productivity from Cline (2007) minus my estimate of the change in manufacturing productivity, shown above, in percentage points.

To project the impact of climate change on productivity in manufacturing and services, I combine the country-sector specific temperature sensitivities estimated in Section 4.4 with projections of the future distribution of temperature in 2080-2099.[44] I obtain future temperature predictions from the CSIRO-MK-3.6.0 model produced by Jeffrey et al. (2013), one of the climate models used by Cline (2007), for consistency with the projected changes in agricultural productivity.[45] The projected changes in manufacturing and services productivity are shown in Figure 9 and Appendix Figure A-24 respectively. Figure 10 brings together the estimated impacts on agricultural productivity from Cline (2007) with my estimates of the change in manufacturing productivity to show the change in the relative productivity of agriculture among the tradable sectors for each country in the world.

The pattern in Figure 10 shows clearly that climate change shifts comparative advantage in agriculture toward colder countries far from the equator on average. While the negative effects of climate change on manufacturing productivity are concentrated in similar parts of the world to agricultural productivity, they are generally smaller in magnitude. Every country in Africa, South Asia, and Latin America (with the exception of Egypt) has larger estimated productivity losses in agriculture than manufacturing. Thus, to the extent that specialization follows Ricardian comparative advantage, we would expect to see agricultural production move toward colder places away from the equator in response to climate change.

I integrate these empirically estimated into the model by applying them to the sector-country specific aggregate productivity Z_{jk} and recalculating equilibrium wages, prices, and trade flows.

[44] I use the estimates that allow for firms to adjust adaptation investments to their end-of-century temperatures. I account for the costs of this adaptation in Section 7.6.

[45] My estimates from the interacted model in Section 4 give me an estimate of the reduction in annual manufacturing and services output per worker for each degree-day above 30°C and below 5°C. The CSIRO model projections give me population-weighted change in degree-days above 30°C and below 5°C for every country in the world in 2080-2099, which are shown in Appendix Figures A-22 and A-23. I multiply the country-level coefficients by the projected changes in hot and cold temperatures to get the impacts shown here.

7.2 Comparative Advantage and Trade

Figure 11 shows the projected equilibrium change in agricultural net exports in response to climate change. Consistent with the estimated change in comparative advantage, the predominant pattern is that hotter countries experiencing large declines in agricultural productivity import more food, while cooler countries with neutral or improving agricultural productivity export more food. For instance, Denmark and Canada roughly double agricultural net exports, from 1.9% to 3.8% and 0.5% to 1.2% of GDP respectively. Conversely, most of Sub-Saharan Africa and South Asia increase imports of food. The few exceptions to this finding are those hot countries for whom the change in agricultural productivity is not large relative to the change in manufacturing productivity, particularly in relation to their close trading partners.

Figure 11: Projected Impact of Climate Change on Agricultural Net Exports

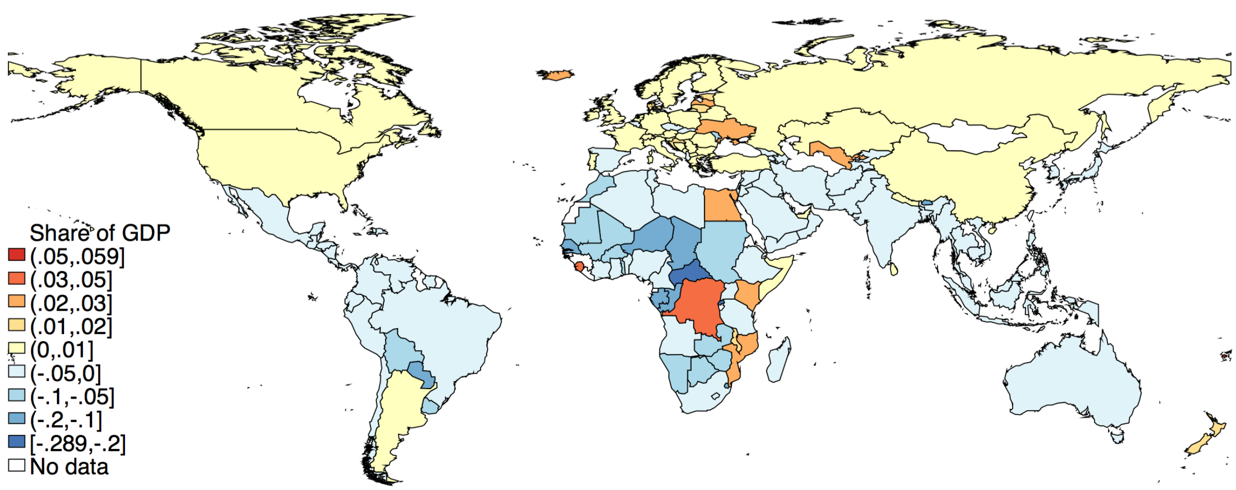

Notes: Map shows model simulations of the change in agricultural net exports as a share of GDP driven by the effects of climate change on sector-level productivity and comparative advantage shown in Figure 10. The full set of country-level results shown in this map are listed in Appendix Tables A-3 to A-20.

The magnitudes of the projected change in trade flows are generally modest as a share of the economy. No country increases agricultural net exports by more than 6% of GDP, and only 12 out of 158 countries decrease agricultural net exports by more than 10% of GDP. The full set of country-level changes in net exports and the domestically produced share of agricultural consumption, (π_{akk}), are shown in Appendix Tables A-3 to A-20.

7.3 Sectoral Reallocation

As shown in Section 5.2, the change in trade flows is only a partial summary of the change in sectoral specialization. Agriculture's share of GDP (and consequently the labor force) depends on both the change in net exports and the change in the expenditure share on food. I reproduce Equation 23 summarizing labor reallocation in response to an agriculture-biased decline in productivity here for convenience:

$$l_{ak} = \underbrace{\pi_{akk}}_{\downarrow} \underbrace{X_{ak}}_{\uparrow} + \underbrace{\sum_{n=1}^{N} \pi_{akn} X_{an} \frac{w_n L_n}{w_k L_k}}_{\downarrow}$$

The change in net exports shown in Figure 11 captures the first and third effects in the above equation. Given the strong nonhomotheticity and low cross-sector elasticity implied by the estimates of ϵ_a and σ in Section 6, the change in the agriculture expenditure share, X_{ak}, is also likely to be substantial.

The horserace between these two competing effects - comparative advantage and 'the food problem' - that govern sectoral reallocation in response to climate change plays a critical role in the aggregate productivity and welfare consequences. As discussed in Section 5, the simple logic formalized by Baqaee and Farhi (2017) is that production moving toward the sector suffering a larger decline in productivity exacerbates the aggregate consequences of a given shock.

I decompose the competing effects of climate change on the agriculture share of GDP by running separate counterfactuals with and without trade. In autarky, the change in a sector's relative price equals the change in that sector's productivity. Thus, I start by applying country-sector level price changes equal to the inverse of the projected change in productivity and calculating the change in expenditure shares. This gives me the change in X_{ak}, which in autarky equals the change in agriculture's share of GDP. In contrast, the standard counterfactual incorporating trade gives me the full effect of both types of reallocation. Table 6 displays the baseline, autarky

Table 6: Counterfactual Ag GDP Shares - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	.065	.067	.07
Brazil	-.169	0	.063	.071	.068
Canada	-.022	-.007	.019	.019	.026
China	-.072	-.036	.064	.068	.074
Denmark	.109	.006	.033	.032	.051
Ethiopia	-.313	-.102	.359	.437	.409
India	-.381	0	.161	.224	.194
Kenya	-.054	-.044	.156	.16	.185
Mozambique	-.217	-.104	.367	.426	.451
Rwanda	-.601	-.058	.409	.678	.351
United States	-.059	.003	.023	.024	.028
Zambia	-.396	0	.36	.496	.41
Poorest Quartile	-.319	-.02	.199	.256	.227
World	-.101	-.01	.038	.044	.043

Notes: Table shows model simulations of the change in agriculture share of GDP driven by the effects of climate change. The full set of country-level results shown in this map are listed in Appendix Tables A-21 to A-29.

counterfactual, and trade-inclusive counterfactual agriculture shares of GDP for a selection of countries, and Appendix Tables A-21 to A-29 contain these results for all 158 countries.

The results in Table 6 show that the consumption response and trade response both have substantial effects on specialization in agriculture, with significant heterogeneity across countries. In Ethiopia, India, and Zambia, the 'food problem' effect dominates and the agriculture share of GDP rises in response to climate change despite large relative declines in agricultural productivity. In contrast, the trade effect dominates in Rwanda, where the domestic share of agricultural expenditures falls from 85% to 54%. Other countries, such as Canada, Denmark, and Kenya see an increase in agricultural specialization because of increased exports driven by improvements in relative agricultural productivity compared to their trading partners.

Figure 12 shows the full worldwide change in agriculture's share of GDP. On average, the global agriculture share of GDP rises from 3.8% to 4.3% because agricultural productivity falls

Figure 12: Projected Impact of Climate Change on Agricultural GDP Share

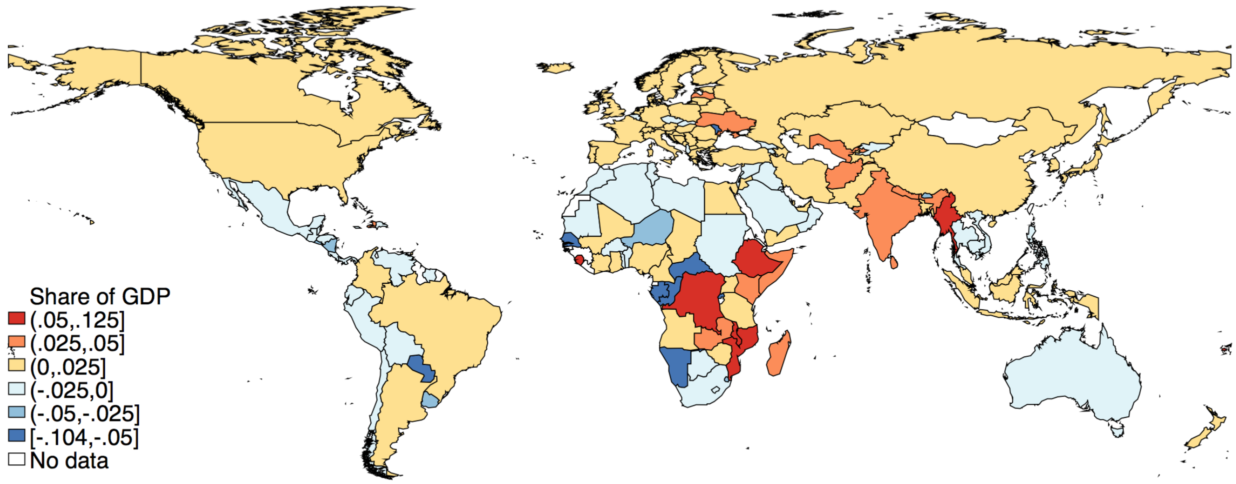

Notes: Map shows the model simulations of the change in the agriculture share of GDP driven by climate change. Appendix Tables A-21 to A-29 contain the full set of country-level results pictured here.

in more places than it rises, raising X_{ak}, and net exports for the world are zero. More specifically, the 'food problem' effect particularly dominates on average in those countries suffering large relative declines in agricultural productivity. The average change in the agriculture share of GDP for countries facing a 10% or larger decline in relative agricultural productivity, weighting by their share of agricultural workers, is +2.1 percentage points from an initial share of 17.3%.

7.4 Aggregate Productivity and Willingness-to-Pay

The estimated sectoral productivity effects combined with the changes in sectoral specialization map directly into changes in aggregate productivity. Table 7 shows the change in real GDP for each counterfactual in select countries, deflating nominal income at the country level using the Tornqvist price index from Equation 20. The results for the full set of countries are shown in Appendix Tables A-30 to A-38.

The results make clear that projected reallocation *exacerbates* the impact of climate change on aggregate productivity in most countries, as well as globally on average. Global GDP declines 1.9% in the counterfactual that holds sectoral shares fixed, but 2.1% when allowing for reallocation. GDP in the poorest quartile of countries falls by 8.3% in the no reallocation counterfactual,

Table 7: Counterfactual GDP Losses (Share of GDP) - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.002	0	.001
Brazil	-.169	0	-.01	-.015	-.013
Canada	-.022	-.007	-.018	-.018	-.016
China	-.072	-.036	-.043	-.045	-.045
Denmark	.109	.006	0	0	.005
Ethiopia	-.313	-.102	-.163	-.218	-.217
India	-.381	0	-.074	-.131	-.127
Kenya	-.054	-.044	-.037	-.038	-.034
Mozambique	-.217	-.104	-.14	-.192	-.199
Rwanda	-.601	-.058	-.334	-.557	-.508
United States	-.059	.003	0	0	.001
Zambia	-.396	0	-.175	-.328	-.314
Poorest Quartile	-.319	-.02	-.083	-.132	-.126
World	-.101	-.01	-.019	-.023	-.021

Notes: Table shows model simulations of the change in GDP driven by the effects of climate change. The full set of country-level results are shown in Appendix Tables A-30 to A-38.

and 12.6% with reallocation. This happens for two reasons. First, as discussed in Section 7.3, the 'food problem' pushes up the labor share of agriculture in many countries while agricultural productivity declines dramatically. Second, as Dingel, Meng and Hsiang (2019) have shown, the spatial correlation of the productivity impacts heighten their importance. Since food prices in Rwanda are a function of agricultural productivity in Rwanda and its closest trading partners, the losses to Rwanda intensify when accounting for the full general equilibrium effects, including those of shocks that hit their neighbors.

How can reallocation that worsens aggregate productivity and measured GDP be consistent with optimizing behavior? In Table 8, I calculate the willingness-to-pay (WTP) to avoid climate damages under each counterfactual as the equivalent variation loss in income at the baseline equilibrium set of wages and prices. The results show that the full reallocation counterfactual mitigates the welfare consequences of climate change, as captured by willingness-to-pay, even while increasing the impact on GDP. The WTP under the no reallocation counterfactual is particularly dramatic because it forces agents to deviate from optimal consumer behavior. This highlights that the no reallocation counterfactual is, in some sense, an unrealistic straw man. In the presence of very large projected increases in food prices, keeping fixed the expenditure share on food would require declines in the quantity of food consumed that are strongly inconsistent with the observed low substitutability between food and non-food. To summarize the intuition, people are willing to sacrifice income (GDP) to reallocate expenditures toward food when food prices rise because they need food to survive.

Table 8: Equivalent Variation Willingness-to-Pay (Share of GDP) - Selected Countries

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.008	-.002	0
Brazil	-.169	0	-.041	-.01	-.008
Canada	-.022	-.007	-.018	-.016	-.014
China	-.072	-.036	-.057	-.04	-.04
Denmark	.109	.006	.003	0	.005
Ethiopia	-.313	-.102	-.364	-.171	-.169
India	-.381	0	-.311	-.085	-.082
Kenya	-.054	-.044	-.052	-.035	-.031
Mozambique	-.217	-.104	-.279	-.143	-.147
Rwanda	-.601	-.058	-.725	-.434	-.387
United States	-.059	.003	-.002	0	.001
Zambia	-.396	0	-.481	-.208	-.199
Poorest Quartile	-.319	-.02	-.277	-.092	-.088
World	-.101	-.01	-.04	-.018	-.017

Notes: Table shows model simulations of the willingness-to-pay to avoid the effects of climate change. The full set of country-level results are shown in Appendix Tables A-39 to A-47.

Figures 13 and 14 show the global distribution of willingness-to-pay to avoid climate change, and the change in food prices, P_{ak}, which comprise a key driver of the welfare losses. Food prices rise in 156 of the 158 countries, and rise by at least 25% in 41 countries containing over 32% of the world's population.[46] Climate change does net damage as measured by WTP in 150 countries,

[46] The large changes in food prices also imply that the incidence of these losses may fall on urban consumers as much or even more than on farmers suffering lost productivity.

Figure 13: Willingness-to-Pay to Avoid Climate Change

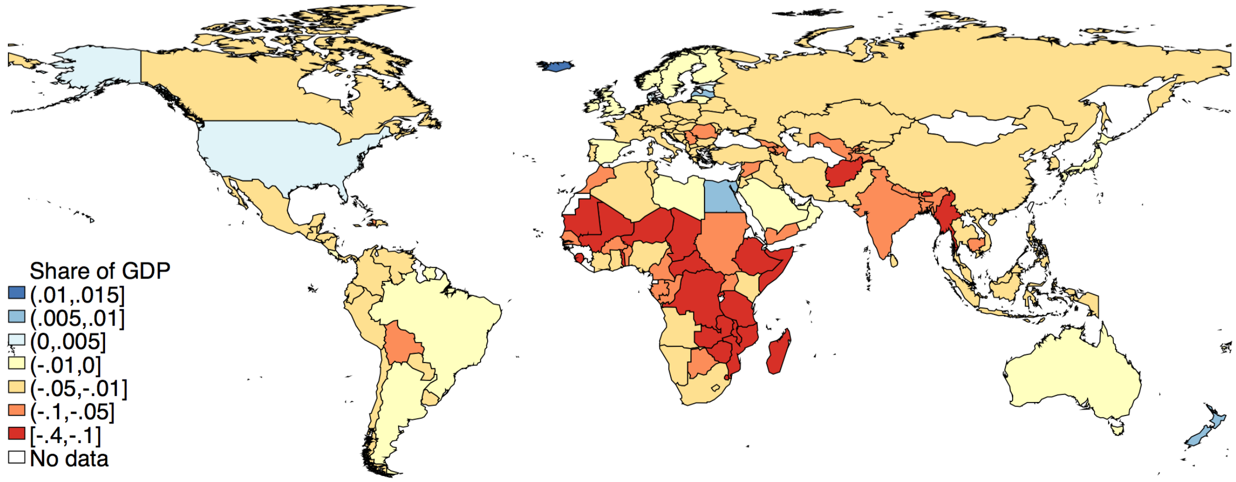

Notes: Map shows model simulations of the willingness-to-pay to avoid the effects of climate change as a share of GDP. The full set of country-level results shown in this map are listed in Appendix Tables A-39 to A-47.

Figure 14: Projected Percentage Change in Food Prices

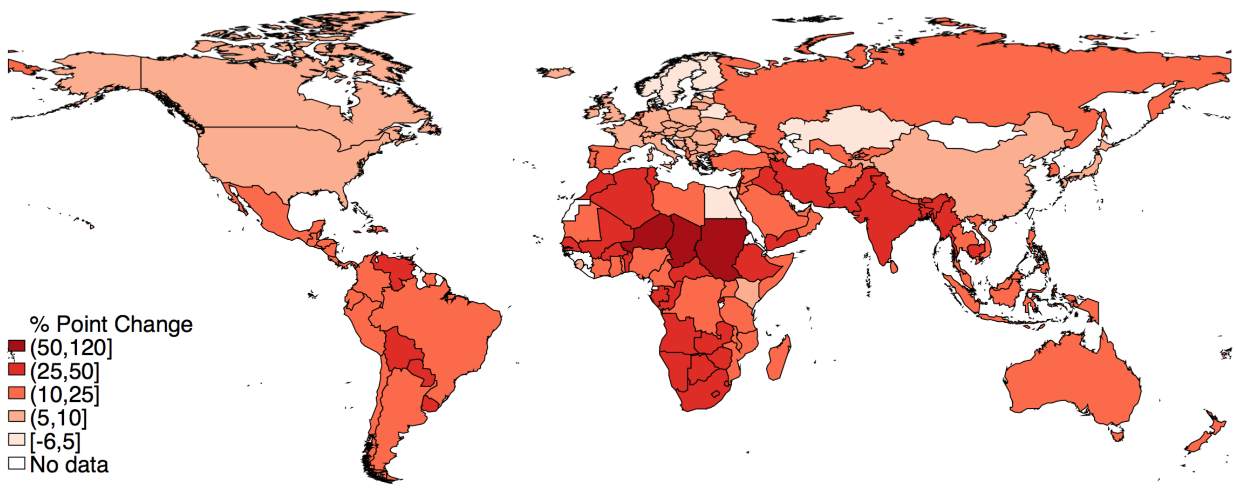

Notes: Map shows model simulations of the change in food prices driven by climate change. The full set of country-level results shown in this map are listed in Appendix Tables A-48 to A-56.

and causes welfare losses exceeding 8% of GDP in 32 countries covering 27% of the world's population. Because the losses are concentrated in poor countries, global willingness-to-pay is only 1.7% of GDP. However, the population-weighted average global losses are 4.7% of GDP, and the population-weighted average for countries in the bottom quartile of income is 8.8% of GDP. The interpretation of this number is that climate change will cost the average person in the poorest quartile of the world nearly 9% of their income. Note that these results account neither for the costs of firm-level adaptation investments nor for the benefits of anticipated economic growth, both of which will be included in Section 7.6.

7.5 Low Trade Cost Counterfactual

The analysis of sectoral reallocation and aggregate productivity in Sections 7.3 and 7.4 demonstrates that openness to trade mitigates the harm from climate change by counteracting 'the food problem.' To further investigate the magnitude to which facilitating trade could contribute to climate change adaptation, I run an additional counterfactual exercise in which I replace the estimated matrix of bilateral trade costs, τ_{jkn}, with a uniform low value representing increased openness to trade. In particular, I set the cost of all bilateral trade for both manufacturing and agriculture at 100%. I choose this number rather than 0% to acknowledge the fact that some level of shipping costs, regulatory discrepancies, and language barriers are inherent to cross-country trade, so no amount of policy intervention could make trade perfectly costless. A 100% tariff-equivalent trade cost is toward the low end of the estimated distribution - approximately equal to the cost I estimate for shipping food from Belgium to Australia. I choose this value to represent an ambitious, yet realistically feasible, change in global trade policy.

To disentangle the benefits of trade for climate change adaptation from the more general gains from trade, I rescale each country's vector of sectoral productivity parameters, Z_{jk}, such that I continue to match the baseline levels of GDP per capita in the initial equilibrium. Note, however, that without the estimated high barriers to trade in developing countries the model can no longer match the observed global pattern of the agriculture share of GDP. In this hypothetical world of increased openness, developing countries import substantially more food from richer countries with high relative productivity in agriculture even in the absence of climate change.

Table 9 shows the WTP to avoid climate change under different trade cost scenarios for a select subset of countries especially vulnerable to climate change. Appendix Tables A-57 to A-65

Table 9: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Rwanda	-.434	-.387	-.086
Central African Republic	-.428	-.356	-.037
Chad	-.25	-.226	-.032
Malawi	-.225	-.225	-.119
Zimbabwe	-.223	-.212	-.074
Zambia	-.208	-.199	-.001
Ethiopia	-.171	-.169	-.091
Sierra Leone	-.13	-.164	-.105
India	-.085	-.082	-.013
World	-.018	-.017	-.013
Poorest Quartile	-.092	-.088	-.029

Notes: Table shows model simulations of the willingness-to-pay to avoid the effects of climate change under different scenarios - autarky, estimated global barriers to trade, and an alternative scenario at which all bilateral trade costs are set at a low level typical of OECD countries. The full set of country-level results are shown in Appendix Tables A-57 to A-65.

show these counterfactuals for the full range of countries. Two things about these results are worth noting. First, as shown in Table 9, reducing trade barriers dramatically reduces the costs of climate change in the hardest-hit countries. Overall, the WTP for the average person in the lowest quartile of global income is only 2.9%, relative to 8.8% in the estimated trade cost case.

Second, the effects of openness to trade vary substantially across countries. For 40 countries representing 15.1% of the global population, WTP to avoid climate change as a share of GDP is *higher* in the low trade cost scenario.[47] The intuition for this result is as follows. When trade barriers are high and local consumption depends mostly on local production, the effects of deteriorating productivity are also concentrated locally. Conversely, more trade makes the world more interdependent and dilutes the effects of a local shock across many countries. If consumption in Austria is more linked to production in Zimbabwe, then Austrian consumers suffer more from shocks that hit Zimbabwe. Conversely, Zimbabwean consumers insulate themselves from

[47] To be clear, these countries still experience overall gains from trade. But once those general gains are netted out, they suffer larger climate change damages in this scenario.

the local shock by consuming a more diversified global portfolio of products.

Overall, trade reduces the aggregate global willingness-to-pay to avoid climate change by 7.4% relative to autarky under existing global trade policy, and by 30.7% under the specified alternative assumption of freer trade. This pattern holds much more starkly in poor countries. For the average person in the poorest quartile of the world, trade reduces WTP by 4.5% relative to autarky under existing policy, but by 68.2% under freer trade. I discuss possible policy mechanisms to realize these gains in Section 9.

7.6 Future Projections

The results in Sections 7.1 to 7.5 use projections for future temperature change, but hold the baseline global economy fixed at the present day equilibrium. In this section, I endeavor to better represent the future baseline in 2080 by allowing global income levels to evolve according to projections from the Shared Socioeconomic Pathway (Scenario Three) developed by Cuaresma (2017) of the International Institute for Applied Systems Analysis.[48]

Allowing for economic growth to take place has two important effects on the aggregate consequences of climate change. First, the agriculture share of GDP declines as countries grow richer due to nonhomothetic preferences for food, reducing the aggregate consequences of agriculture-specific productivity shocks. I capture this effect in the model by applying projected income growth to 2080 as sector-neutral increases in the baseline values of Z_{jk}. Second, my results from Section 4 imply that sensitivity to temperature for manufacturing and services firms declines markedly as countries become richer. I capture this by re-evaluating the sensitivity to temperature shown in Figures 5 and 6 at 2080 levels of log GDP per capita. Appendix Figures A-25 and A-26 show that the effects of temperature on non-agricultural productivity accounting for adaptation are substantially muted, even in this relatively low growth scenario that projects only slightly more than a doubling of global income between 2015 and 2080.

Table 10 shows the impact of expected economic growth on the agriculture share of GDP and expected willingness-to-pay. Appendix Tables A-66 to A-74 show the results for all countries.

[48]Use of the Shared Socioeconomic Pathways in future projections of climate change damages follows from the work of Carleton et al. (2018).

Table 10: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Central African Republic	1.47	.299	.287	.094	-.436	-.316
Rwanda	1.14	.409	.39	.322	-.394	-.366
Zimbabwe	4.17	.302	.122	.14	-.248	-.111
Malawi	2.84	.436	.309	.357	-.244	-.167
Zambia	1.52	.36	.28	.318	-.233	-.181
Chad	1.13	.257	.213	.243	-.226	-.221
Sierra Leone	1.49	.139	.177	.173	-.204	-.146
Ethiopia	1.23	.359	.333	.376	-.19	-.182
India	3.24	.161	.087	.106	-.082	-.045
Poorest Quartile	3.05	.199	.126	.144	-.1	-.062
World	2.2	.038	.025	.028	-.027	-.015

Notes: Table shows model simulations of the effects of projected economic growth on the agriculture share of GDP and the willingness-to-pay to avoid climate change in select countries. Economic growth projections come from Cuaresma (2017). The full set of country-level results are listed in Appendix Tables A-66 to A-74.

The willingness-to-pay numbers in Columns 5 and 6 of Table 10 also incorporate the firm-level adaptation costs shown in Appendix Figure A-18, thus accounting more comprehensively for the anticipated costs as well as benefits of adaptation.[49] This particular future scenario includes little to no projected growth for many currently poor countries, allowing for contrast with those that grow faster. This comparison shows the importance of economic growth in mitigating the harm from climate change. Table 10 shows that Zimbabwe and Malawi get substantially richer in this projection, and their agriculture share of GDP and climate change damages decline markedly. In contrast, climate change continues to be very harmful to countries that grow slowly, such as Rwanda and Chad.

[49] I exclude this revealed preference measure of firm-level adaptation costs from Tables 8 and 9 because they are calculated as a share of manufacturing and services output, which vary dramatically as a share of baseline total output in the low trade cost scenario, thus complicating the comparison of climate change damages between the estimated and low trade cost counterfactuals.

The results in Table 10 show that the aggregate global WTP for climate change is 2.7% of GDP at current global income levels and 1.5% at future projected incomes. The average WTP for a person in the bottom quartile of the world is 10.0% from the present baseline and 6.2% from the future baseline. To summarize the importance of the distributional consequences of climate change, I follow Jones and Klenow (2016) to calculate the willingness-to-pay of a Rawlsian social planner taking the certainty equivalent of being any person in the world with random probability.[50] The Rawlsian welfare losses from climate change are 6.2% of global GDP from the present income baseline and 3.6% of global GDP from the future baseline, more than twice as high as the aggregate willingness-to-pay calculated by summing across agents.

8 Supporting Empirical Evidence

In this section, I present country-level panel regression evidence consistent with the model counterfactuals. In particular, my results in Section 7 suggest that the 'food problem' outweighs the trade response, on average, in driving sectoral reallocation due to climate change. This finding is supported by the simulated method of moments inference that underlies my parameter estimates, is consistent with both cross-sectional and historical patterns of sectoral specialization in the world, and is further bolstered by existing empirical evidence that aims to isolate the causal effect of agricultural productivity on structural transformation. In particular, Gollin, Hansen and Wingender (2018) proxy for improvements in agricultural productivity using variation in the development, diffusion, and climatic suitability for high-yielding crop varieties and Bustos, Caprettini and Ponticelli (2016) study the introduction of genetically engineered soybean seeds in Brazil. Both papers find that rising agricultural productivity drove labor out of agriculture and into industry. Here, I present evidence suggestive of the converse more representative of climate change - that declines in agricultural productivity increase the agriculture share of GDP and labor on average relative to the counterfactual.

Table 11 summarizes the data sources used in this part of my analysis.[51] Following Schlenker

[50] Following Jones and Klenow (2016) I use log utility in this calculation.

[51] I use BEST temperature data with a 1°global grid in this specification because aggregating GMFD temperature data from a 0.25°grid for every country worldwide exceeds my available computational resources.

Table 11: Country-Level Panel Data

Variable	Data Source
Temperature	Berkeley Earth Surface Temperature Dataset
Ag Share of GDP	World Bank
Ag Share of Labor Force	International Labour Organization
Food Share of Imports	UN Comtrade
GDP	World Bank

Notes: Data covers 164 countries from 1960-2012 with varying coverage by country and dataset. Economic data from all sources above are retrieved from the World Bank Databank.

and Roberts (2009), I use "growing degree days" (GDD) between 0°C and 29°C and "killing degree days" (KDD) above 29°C as temperature transformations representing positive and negative shocks to agricultural productivity. I aggregate GDD and KDD to the country level for each year weighting by each pixel's share of cropland.[52]

I estimate the following panel regression with observations at the country-year level for four separate outcome variables - log GDP, food share of imports, agricultural share of GDP, and agricultural share of labor:

$$Y_{it} = \beta_1 GDD_{it} + \beta_2 KDD_{it} + \delta_i + \kappa_t + \epsilon_{it} \tag{24}$$

The regression exploits idiosyncratic variation in weather controlling for country fixed effects, δ_i, and year fixed effects, κ_t to estimate the plausibly causal effect of shocks to agricultural productivity. I weight observations by their share of the global agricultural labor force to recover expected reallocation for the average farm worker in the world.

The results in Table 12 are broadly consistent with my model simulations in Section 7. The

[52]Following standard procedure in estimating temperature effects on agricultural productivity, degree days are calculated by fitting a sinuisoidal curve through daily minimum and maximum temperature, and then integrating the proportion of each day above a certain threshold.

Table 12: Country-Level Panel Regression

	(1) log(GDP)	(2) Food Share of Imports	(3) Ag Share of GDP	(4) Ag Labor Share
KDD X 100	-0.121	0.00258	0.00875	0.00991
	(-2.31)	(0.64)	(1.08)	(1.55)
GDD X 100	0.0505	-0.00429	-0.00140	-0.00138
	(1.64)	(-2.45)	(-1.54)	(-0.38)
Observations	3602	2916	3171	3715
Country FE	X	X	X	X
Year FE	X	X	X	X
Ag Labor Weights	X	X	X	X

Notes: t-statistics in parentheses. Reported Driscoll and Kraay (1998) standard errors are robust to heteroskedasticity, spatial correlation, and autocorrelation of up to 5 lags. Results come from estimating Equation 24 with crop-area weighted growing and killing degree days. Data covers 164 countries from 1960-2012 with varying coverage by country and outcome variable. Economic data from all sources above are retrieved from the World Bank Databank.

composition of imports shifts toward food in response to negative agricultural productivity shocks (KDD), and away from food in response to positive shocks (GDD), but the magnitudes of these changes are small. Consistent with an important role for 'the food problem,' the agriculture share of GDP and labor rise with KDD and fall with GDD, with magnitudes roughly similar to those in the model. Here, the agriculture share of GDP rises by slightly under 1 percentage point for an agriculture-biased shock that reduces GDP by 12%. To construct a corresponding level of reallocation in my model simulations, I calculate that the agricultural population-weighted average change in the agricultural share of GDP for those countries suffering large declines in agricultural productivity (<10 percentage points) is +2.1 percentage points from an average agricultural productivity fall of 29.5%.[53]

[53]An additional feature of the regression that supports the approach taken in the model is the very similar coefficients estimated for the agriculture share of GDP and the labor force. These two shares are equivalent in my model because I allow wages to equalize across sectors, but the agriculture share of labor is generally higher in the data since agricultural wages tend to be lower. The similar coefficients in Table 12 suggest that projecting reallocation in ag GDP is informative for understanding labor reallocation even if the levels of these two variables differ.

The results from the country-level regressions are imprecise and insufficient in isolation to make full general equilibrium projections or welfare calculations relating to sectoral reallocation in response to climate change.[54] Taken together with the analysis in Sections 6 and 7 and the existing body of evidence, however, these results reinforce the important role of the 'food problem' in mediating the aggregate consequences of climate-driven agricultural productivity shocks.

9 Policy Implications

This paper has three sets of implications relevant to policy on climate change and development. First, the results inform cost-benefit analysis on policies to reduce greenhouse gas emissions and avoid damage from climate change. These results are not a comprehensive evaluation of the costs of climate change - I omit international migration, uncertainty, health effects, and non-temperature effects such as storms and sea-level rise, among other topics, from my analysis. I do, however, address an existing challenge in the literature by estimating global reductions in aggregate productivity and calculating their welfare consequences in a framework that accounts for reallocation of economic activity between agriculture and non-agriculture.

Second, my results inform decisions about the best way to channel efforts to adapt directly to the consequences of climate change. If it were true that agricultural activity is likely to shift substantially away from hot developing countries, optimal investments in adaptation might focus on retraining farm workers to transition to non-agricultural occupations. Instead, my finding that climate change is more likely to increase specialization in agriculture in hot countries underscores the urgent need to reduce the temperature-sensitivity of production through technology, irrigation, heat-resistant crop varieties, or other means. The agricultural productivity consequences projected by Cline (2007) will take place gradually and worsen far into the future, and need not be invariant to efforts to reduce them.

Third, and perhaps most importantly, my results speak to the importance of reducing barriers to trade in developing countries as a mechanism for climate change adaptation. The results

[54] I show results for the unweighted regressions in Appendix Table A-75. I gain precision in the unweighted specification because the agriculture labor share weights are missing for a nontrivial share of the observations, but have a less interesting interpretation of the coefficients as effects on the average country in the world rather than on the average unit of agricultural labor.

in Section 7.4 suggest that the costs of climate change to the average person in the poorest quartile of the world could be reduced by more than half in a world with a plausible increase in openness to trade. Reducing tariffs would be one place to start, but tariffs account for a relatively small proportion of estimated trade costs. As Tombe (2015) documents at length, red tape barriers appear to be a far more important deterrent in many places. Figures 15 and 16 show data from the World Bank Ease of Doing Business Indicators on fees and delays associated with importing a container.

Figure 15: Direct Costs to Import a 20-Foot Long Container (USD)

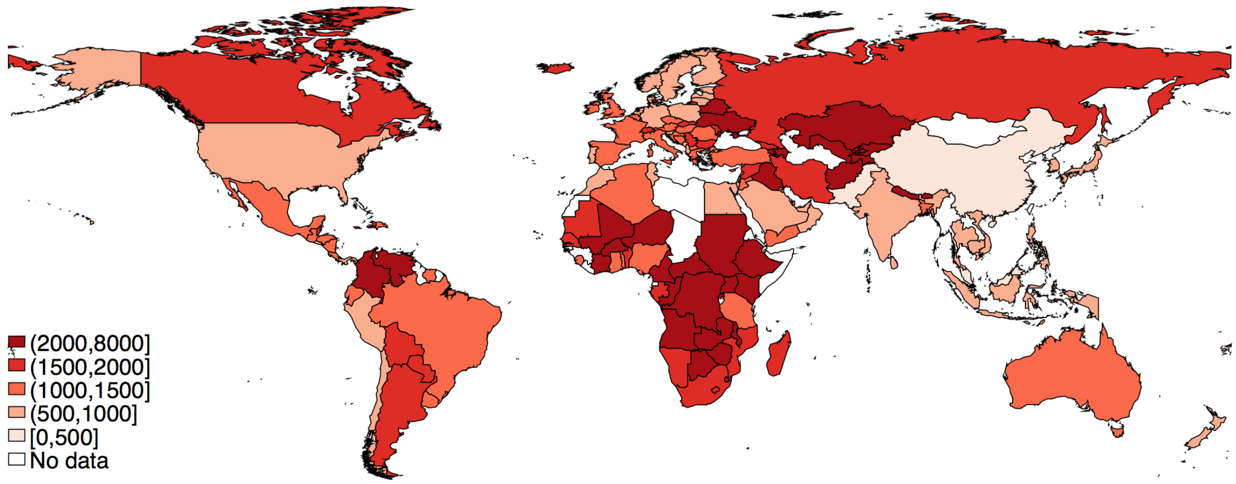

Notes: Figure shows the direct cost to import one container of goods. Costs include documents, administrative fees for customs clearance, terminal handling charges, and inland transport, but not tariffs or taxes. Data comes from the World Bank Ease of Doing Business index.

The average country in Sub-Saharan Africa requires 9 documents and over $2700 in fees for customs clearance, document processing, customs brokerage, terminal handling, and inland transport to import a 20-foot container of goods, exclusive of tariffs and unofficial payments. Importing a shipment to Sub-Saharan Africa also requires waiting an average of 37 days upon arrival at the border for compliance with customs clearance, inspection procedures, and document preparation, likely a prohibitive length of time for many food imports. These types of trade barriers do not involve international negotiations or physical constraints to shipping over long distances, and thus could be a relatively tractable place to target reforms that could make a

Figure 16: Days to Import a Container

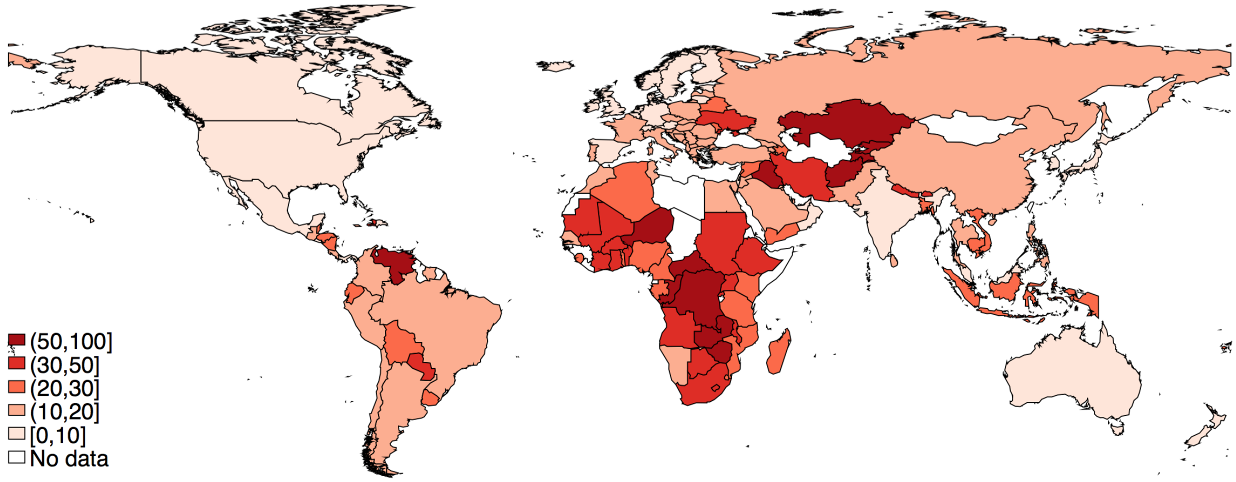

Notes: Figure shows the average number of days required to import a container. Delays include customs clearance, government inspection procedures, and documentary compliance requirements. Data comes from the World Bank Ease of Doing Business Index.

substantial impact on climate change adaptation.

10 Conclusion

The standard intuition in economics is that reallocation improves outcomes. Falling productivity raises prices and encourages substitution to other products. But this logic does not hold for broad categories of necessary consumption, such as food. If a fall in productivity causes the price of corn to rise sharply, people can adapt by eating more rice. But when people become poorer and the relative price of food rises, they cannot compensate by substituting away from food.

This paper investigates the importance of subsistence requirements for food for the general equilibrium and aggregate productivity effects of climate change. I show that climate change predominantly shifts comparative advantage in agriculture away from the equator as the effects of extreme temperatures on non-agricultural productivity are generally smaller than those on agriculture. On average, however, the effect of a large decline in productivity concentrated in agriculture moves specialization toward, rather than away from, agriculture because of the special properties of consumer preferences for food. Countries with large climate change productivity impacts in agriculture that are more open to trade suffer less because they are more able to increase imports of food and shift production toward other sectors. Overall, reducing barriers to trade could reduce the losses from climate change by more than half for the poorest quarter of the world's population.

I conclude with several suggestions for future research. First, while my work is informative about the cost-benefit analysis of climate change mitigation, additional effort is required to integrate these general equilibrium effects directly into calculations of the social cost of carbon. Second, while my analysis shows that reducing barriers to trade is a necessary condition to induce sectoral reallocation to curtail the costs of climate change, I cannot conclude that it is sufficient. A low trade cost counterfactual in which specialization in agriculture shifts away from the equator still relies on uncertain assumptions about diminishing returns to expanding production of tradable manufactured goods in developing countries, as well as on the availability of complementary inputs such as soil quality and arable land in cold countries experiencing improved temperature suitability for agriculture. A final topic concerns the political economy of trade policy regarding food. Policymakers often prioritize "food security" as a stated aim, implying a preference for domestic food production secure from interference by foreign countries. To the extent that this goal conflicts with adaptation to climate change in light of large declines in agricultural productivity in certain regions, it may be worth examining this tradeoff more closely, both in practice and in perception.

Chapter 2: Do Renewable Portfolio Standards Deliver Cost-Effective Carbon Abatement?

with Michael Greenstone

Abstract

Renewable Portfolio Standards (RPS) are the largest and perhaps most popular climate policy in the US, having been enacted by 29 states and the District of Columbia. Using the most comprehensive panel data set ever compiled on program characteristics and key outcomes, we compare states that did and did not adopt RPS policies, exploiting the substantial differences in timing of adoption. The estimates indicate that 7 years after passage of an RPS program, the required renewable share of generation is 1.8 percentage points higher and average retail electricity prices are 1.3 cents per kWh, or 11% higher; the comparable figures for 12 years after adoption are a 4.2 percentage point increase in renewables' share and a price increase of 2.0 cents per kWh or 17%. These cost estimates significantly exceed the marginal operational costs of renewables and likely reflect costs that renewables impose on the generation system, including those associated with their intermittency, higher transmission costs, and any stranded asset costs assigned to ratepayers. The estimated reduction in carbon emissions is imprecise, but, together with the price results, indicates that the cost per metric ton of CO2 abated exceeds \$115 in all specifications and ranges up to \$530, making it least several times larger than conventional estimates of the social cost of carbon. These results do not rule out the possibility that RPS policies could dynamically reduce the cost of abatement in the future by causing improvements in renewable technology.[55]

[55]We thank Frank Wolak, Ken Gillingham, Nancy Rose, Matt Zaragoza-Watkins, Chris Knittel, Dick Schmalensee, James Bushnell, Koichiro Ito, and participants at the MIT Public Finance and Industrial Organization Lunches, UC Berkeley Energy Camp, Harvard Environmental Economics Lunch, and MIT CEEPR meetings for their comments. We also thank Catherine Che and Henry Zhang for providing excellent research assistance.

1 Introduction

Even as evidence mounts on the costs of climate change, the United States has had great difficulty developing significant and enduring policy responses, particularly in the power sector which is a primary source of greenhouse gas emissions. One major exception has been renewable portfolio standards (RPS) that require that a certain percentage of electricity supply in a state is met by generation from sources that are designated as renewable. The first RPS was passed in Iowa in 1991 and since then others have followed suit. As of 2015, RPS policies have been enacted in 29 states and the District of Columbia, which together account for 62% of electricity generation.[56] Further, the ambition of these policies has grown dramatically. In the early years of implementation, RPS policies typically required increases in the renewables share of electricity of a couple of percentage points, but states have greatly ramped up their ambitions, with, for example, 2030 targets of 41% (Massachusetts), 44% (Connecticut), 50% (New York), and 60% (California). Indeed, RPS have been credited with greatly expanding the penetration of renewable technologies, most frequently wind and solar, which rose from 0.1% of all generation in the United States in 1990 to 5.3% in 2015 . Further, their penetration rate has increased greatly in recent years and indeed they accounted for approximately half of the new installed capacity since 2010.[57]

Despite the popularity of these policies, there is little if any systematic evidence on RPS' impacts on electricity prices or carbon emissions. A common approach to estimating their costs is to calculate the difference in costs associated with a RPS that is, compare the costs of a renewable plant with the costs of a fossil fuel plant that it replaces. This type of calculation entails comparing the levelized cost of energy (LCOE), calculated by dividing the total direct costs associated with investment in new capacity by expected total lifetime energy production. The latest data from the Energy Information Administration's Annual Energy Outlook (EIA (2019)) suggests that solar and wind plants can produce electricity at about 6 cents per kWh, while a natural gas combined cycle plant produces at roughly 4 cents per kWh. Since to date RPS

[56] An additional seven states enacted non-binding targets under similar programs.

[57] This fact comes from Bushnell, Flagg and Mansur (2017).

policies have only increased renewable penetration by a few percentage points, it is this type of comparison of LCOEs that lead observers to suggest that RPS policies have had only a minimal impact on electricity prices; one recent study found that they increase retail electricity bills by about 2% (see, e.g., Barbose (2018)).

However, this comparison of LCOEs misses three key ways in which renewables impose costs on the electricity generation system that need to be covered and are reflected in *retail* prices but can be difficult to observe directly or measure systematically. First, and most obviously, renewables by their very nature are intermittent sources of electricity. Solar plants cannot provide power when the sun doesn't shine and wind plants cannot provide it when the wind isn't blowing. On average, utility scale solar plants have a capacity factor (i.e., average power generated divided by its peak potential supply over the course of a year) of about 25% and wind plants are not much higher at 34% according to the EIA. This means that a comparison of LCOEs between these intermittent sources and "baseload" technologies that "always" operate (e.g., natural gas combined cycle plants have capacity factors of 85%) is very misleading with respect to total system costs, because they do not account for the additional costs necessary to supply electricity when they are not operating. For example, given current cost structures, the installation of renewables are frequently paired with the construction of natural gas "peaker" plants that can quickly and relatively inexpensively cycle up and down, depending on the the availability of the intermittent resource.

Second, renewable power plants require ample physical space, are often geographically dispersed, and are frequently located away from population centers, all of which raises transmission costs above those of fossil fuel plants. A literature review of transmission cost estimates for wind power by the Lawrence Berkeley National Laboratory (LBNL) finds a median estimate of about $300 per kW, or about 15% of overall wind capital costs (Mills, Wiser and Porter, 2009). This is approximately equivalent to adding 1.5 cents per kWh to the levelized cost of generation for wind. More generally, a separate analysis by the Edison Electric Institute in 2011 found that 65% of a representative sample of all planned transmission investments in the US over a ten-year period, totaling almost $40 billion for 11,400 miles of new transmission lines, were primarily directed toward integrating renewable generation.[58] The highly disproportionate share of transmission requirements for renewables relative to their share of generation highlights the importance of accounting for the associated costs as part of the total cost of renewable energy.

Third, RPS driven increases in renewable energy penetration can also raise total energy system costs by prematurely displacing existing productive capacity, especially in a period of flat or declining electricity consumption. Adding new renewable installations, along with associated flexibly dispatchable capacity, to a mature grid infrastructure may create a glut of installed capacity that renders some existing baseload generation unnecessary. The costs of these stranded assets do not disappear and are borne by some combination of distribution companies, generators, and ratepayers. Thus, the early retirement or decreased utilization of such plants can cause retail electricity rates to rise even while near zero marginal cost renewables are pushing down prices in the wholesale market. The incidence of excess capacity costs on ratepayers is likely greater in regulated markets with vertical integration, although even in deregulated markets there are various mechanisms for direct payments to producers unconnected to actual generation that can contribute to the rates consumers face.[59] Overall, there exists no comprehensive source of data on payments to displaced electricity producers, and even the availability of such information would not provide an obvious path to attributing these costs to the integration of renewables. Like many of the other ancillary costs of renewable energy integration, directly observing the total costs associated with stranded capacity is unlikely to be feasible.

As an alternative to what we believe is the nearly impossible task of directly measuring each of the mechanisms by which RPS policies influence costs, this paper compares states that did and did not adopt RPS policies, using the most comprehensive panel data set ever compiled

[58]The Edison Electric Institute collected a representative sample of transmission projects totaling over $61 billion from their members, who cover about 70% of the total US electricity market. See EEI (2011) and Mills, Wiser and Porter (2009).

[59]For instance, ISO New England made over $1 billion of capacity market payments unconnected to actual generation in 2013, comprising 12% of their total wholesale market expenditures. Over 95% of these payments supported existing, rather than new, capacity. The Independent System Operator for New England covers production in Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont. They publish capacity market information in their annual market report: https://www.iso-ne.com/static-assets/documents/2015/05/2014-amr.pdf.

on program characteristics and key outcomes from 1990-2015. Importantly, there is variation in the timing of the adoption of RPS programs across states, which lends itself to powerful event-study style figures that reveal no meaningful evidence of pre-existing different trends in outcomes between adopting and non-adopting states. Further, we are able to control for a series of potentially confounding electricity policies.

There are three key findings. First, RPS policies' statutory requirements for renewable generation frequently overstate their net impact on generation, because they often include generation that existed at the time of the policy's passage. For example, six years after Minnesota adopted its RPS policy, its statutory or total requirement was that renewables account for 14.2% of generation. Yet at the time of adoption, renewables already accounted for 5.3% of generation. So, its net requirement in this year was 8.9%. Due to the substantial heterogeneity in the form and structure of RPS policies, it is challenging to estimate the net requirements and there is no common source for this information. For a handful of states in our sample, even the gross requirement differs across data sources. Nevertheless, our best estimates are that 7 years after adoption the average adopting states' net requirement was 1.8% of generation and 12 years after it was 4.2%.

Second, electricity prices increase substantially after RPS adoption. The estimates indicate that in the 7th year after passage average retail electricity prices are 1.3 cents per kWh or 11% higher, totaling about $30 billion in the RPS states. And, 12 years later they are 2.0 cents, or 17%, higher. The estimated increases are largest in the residential sector, but there are economically significant price increases in the commercial and industrial sectors too. These estimates are robust to controlling for local shocks to electricity costs in a variety of ways. Given the price increases, we also test for impacts on economic activity and fail to find any impact on electricity consumption or state level employment. There is some evidence of a decline in manufacturing employment, but it would not be judged statistically significant by conventional criteria.

Third, the estimates indicate that passage of RPS programs leads to reductions in the generating mixs carbon intensity, although these estimates can be noisier and more sensitive to specification than is ideal. The estimated decline in emissions intensity implies a reduction of

71-250 million metric tons of CO2 across the 29 RPS states 7 years after passage. When the CO2 emissions estimates are combined with the estimated impact on average retail electricity prices, the cost per metric ton of CO2 abated exceeds $115 in all specifications and can range up to $530, making it at least several times larger than conventional estimates of the social cost of carbon (Greenstone, Kopits and Wolverton (2013); EPA (2016)).

Our paper builds on previous work in the economics and engineering literatures that considers the costs and benefits of renewable electricity generation and the impact of RPS programs in particular. One significant line of existing research investigates how baseload, dispatchable, and intermittent resources interact on the grid and how this affects the value of generation from the respective sources and renewables in particular (Denholm and Margolis (2007); Borenstein (2008); Lamont (2008); Joskow (2011); Cullen (2013)). Recent work by Gowrisankaran, Reynolds and Samano (2016) has made particular progress in quantifying the costs of intermittency, and their model resembles the one we present in Section 3. This line of research in economics runs parallel to an engineering literature that uses an energy systems modeling approach to evaluate similar questions (Milligan, Ela, Hodge, Kirby, Lew, Clark, DeCesaro and Lynn (2011); Jacobson, Delucchi, Cameron and Frew (2015)).

The literature on RPS program impact in particular has thus far largely consisted of ex-ante impact estimation. Fischer (2010) and Schmalensee (2012) document the conceptual issues underlying the costs of these programs and Chen, Wiser and Bolinger (2007) survey pre-program prospective assessments, often commissioned by states considering adoption. The median estimate projected that RPS standards would raise retail prices by 0.7%, though the range of projections included significant heterogeneity. The authors also note the importance of underlying assumptions, which focus on capital infrastructure and fuel input costs. A limited body of post-implementation evaluations of certain RPS programs has found slightly larger costs of approximately 2-4% (Heeter, Barbose, Bird, Weaver, Flores-Espino, Kuskova-Burns and Wiser (2014); Tuerck, Bachman and Head (2013)), although this literature has largely taken place outside peer-reviewed journals and generally does not account for all the ways these programs can affect system costs. An important exception to this is Upton and Snyder (2017), who use a difference-in-difference synthetic controls framework to show that RPS programs substantially raise electricity prices and modestly reduce emissions at the state-level.[60]

The paper proceeds as follows. Section 2 provides background on RPS policies and their typical implementation. Section 3 constructs a model to explicate the channels through which integrating renewable generation could raise costs. Section 4 outlines our data sources and presents summary statistics on the electricity sector prior to RPS passage. Section 5 describes our empirical strategy, and Section 6 presents and discusses the results. The paper then finishes with Interpretation and Conclusion sections.

2 Renewable Portfolio Standards

By 2009, 29 states and the District of Columbia had adopted mandatory portfolio standards, while an additional seven states had passed optional standards.[61] These programs currently cover 62% of electricity generation in the US. Figure 17 is a map of the United States that indicates which states have enacted RPS programs, with the colors indicating the years of enactment. Most RPS programs require that retail electricity suppliers meet a percentage of demand with energy from renewable sources.[62] Once in place, the standard typically increases along a predefined schedule until a specified fraction of generation is achieved. For example, California's policy specifies a goal of 33% retail sales from renewables by 2020, with interim targets of 20% by 2013 and 25% by 2016. While the standards sometimes exempt certain providers, most often smaller municipal or cooperative suppliers, they cover 82% of electric load in a state on average.[63]

[60]Tuerck, Bachman and Head (2013) and associated work by those authors also constitute exceptions to this pattern. They account for intermittency and other associated costs using techniques such as engineering estimates, and produce somewhat higher cost estimates of close to 5% of retail prices, though these are still smaller than the effects implied by our estimates.

[61]West Virginia also passed an *Alternative and Renewable Energy Portfolio Standard* in 2009 with characteristics similar to an RPS but which we do not consider. While renewables received some preference in this program, a much broader set of generation sources qualified, including "Advanced Coal Technology," and there was no guaranteed compliance from renewable sources. This program was also repealed before its first binding requirement came into effect.

[62]Our data classify qualifying generation as one of wind, solar, biomass, geothermal, landfill gas, or ocean power, with some states also allowing small hydroelectric.

[63]The statistic on load covered comes from the North Carolina Clean Energy Centers Database of State Incentives for Renewables Efficiency (DSIRE).

Figure 17: RPS Passage by State

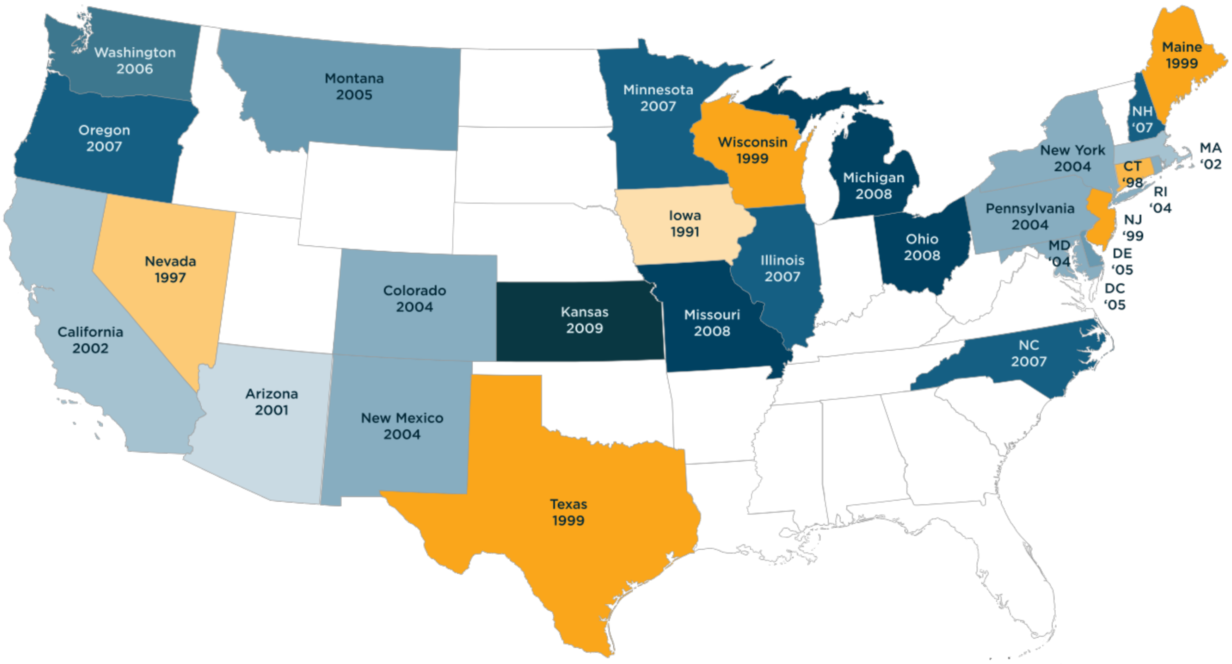

Notes: States that have adopted any RPS policy are colored according to the year in which RPS legislation was passed. Data comes from the US Department of Energy and state government websites.

The key feature of RPS programs is that compliance requires production from a set of designated technologies with the frequent motivation of aiming to help spur innovation that lowers those technologies' costs over time. In practice, the list always includes wind and solar, but whether other technologies are included differs from state to state. Nuclear power is excluded from the policy in all but two states (Massachusetts and Ohio), although it is also a zero carbon energy source.

Electricity providers must demonstrate compliance with the program through Renewable Energy Credits, or RECs, which certify that a given unit of electricity production qualifies to meet the standard. Most RECs are awarded by various regional authorities encompassing several states, which issue unique serial numbers for every megawatt hour of generation produced by registered generators. The approximate coverage of these systems is shown in Appendix Figure A-35. This independent tracking seeks to prevent double counting of generation used for RPS compliance. While there is some scope for transferring RECs between regional systems, in practice most RPS compliance occurs within tracking regions, a fact we will return to later on when considering the impact of RPS on generation outcomes.

Once awarded, credits can be sold separately from the underlying electricity, enabling flexible transfer of the rights to environmental benefits and providing additional revenue to renewable suppliers.[64] In most cases, individual generators must be further approved by the state office administering the RPS to assure that they comply with the specific requirements for generators set forth by that state. In restructured markets, retail providers then purchase RECs generated by these approved facilities, either via brokers or directly through individual contracts. In non-restructured markets, retail providers may also use RECs generated by their own renewable facilities. The serial numbers of the RECs obtained are filed for compliance and their retirement verified with the relevant tracking system. Depending on program rules, excess RECs may also be "banked" for use in later years, though there are typically vintage restrictions requiring relatively recent credits be used. Therefore, REC prices reflect the marginal costs of *producing* electricity from one of the designated technologies, relative to the least expensive alternative, but they do not capture the systemwide costs of *supplying* that electricity, which additionally reflect the costs associated with intermittency, transmission, and compensating owners of stranded assets.

Most RPS programs enforce compliance using a system of Alternative Compliance Payments (ACPs), which effectively fine retail providers for failing to acquire sufficient RECs to cover their sales. These payments are large, averaging about $50 per MWh.[65] Such penalties are substantial, representing about half of the average revenue per MWh observed in 2011. In addition to a penalty, ACPs also provide an effective cost-ceiling for the REC market, as they provide an outside option for compliance. While in practice few retail suppliers fulfill their requirements through ACP payments, REC markets in some states have periodically traded at the ACP level,

[64]A minority of RPS programs have the more stringent requirement that credits be "bundled" with electricity delivered into the state, as demonstrated by transmission to a state balancing authority.

[65]In the case of mandates for generation specifically from solar energy, they can climb even higher, sometimes exceeding $400 per MWh.

suggesting that marginal sources of compliance can be relatively high cost.

While statutory requirements like Maine's 40% target appear quite large, they often ramp up gradually from lower levels and may not reflect the amount of marginal generation actually mandated by RPS policies. Intuitively, if an RPS requirement were entirely covered by existing sources at its inception, in a competitive market we would expect producers to bid down the price of RECs to zero. Distinguishing the amount of new renewable generation required to comply with RPS policy is quite difficult in practice, since covered sources of generation vary from state to state even within narrowly defined categories. For instance, some states allow small-scale hydropower but not large-scale hydropower to qualify for their RPS. Further, six states, including Maine, explicitly mandate that part of their RPS be met using newly constructed renewable capacity. Our best estimate of the net requirement imposed by RPS policies takes the gross amount of MWh required for RPS compliance, as reported by LBNL, and subtracts existing generation from the broad categories of covered sources in the year prior to RPS passage.

Figure 18 reports each states's total and net requirements as of seven years after the state passed RPS legislation, ordering states by the calendar year in which they first adopted an RPS. While these numbers do not fully account for the complications described above, they do show a clear pattern of statutory requirements overstating the amount actually necessary to achieve compliance. For instance, seven event years after passage, the gross requirement in Michigan is 6.2%, but the net requirement after subtracting existing generation in the year of passage is only 2.6%. On average, seven event years after RPS passage, RPS states have a total requirement of 5.1%, but a substantially lower net requirement of 1.8%. In the remainder of the paper, we primarily focus on estimates of net requirements, described in greater detail in Section 4.

Figure 19 plots the number of RPS programs passed into law in each year.[66] The majority of programs were not passed until after 2000. While a number of states adopted RPS policies during, or subsequent to, broader electricity market restructuring, RPS programs have also been

[66]Iowa was the first state to establish a binding standard in 1991, requiring the states's two investor-owned utilities to build or contract for 105 MW of renewable capacity. Although Iowa originally enacted an *Alternative Energy Law* in 1983, the program wasn't given a concrete goal or made compulsory until a revision in 1991, so we consider that the first year of passage.

Figure 18: RPS Passage by State

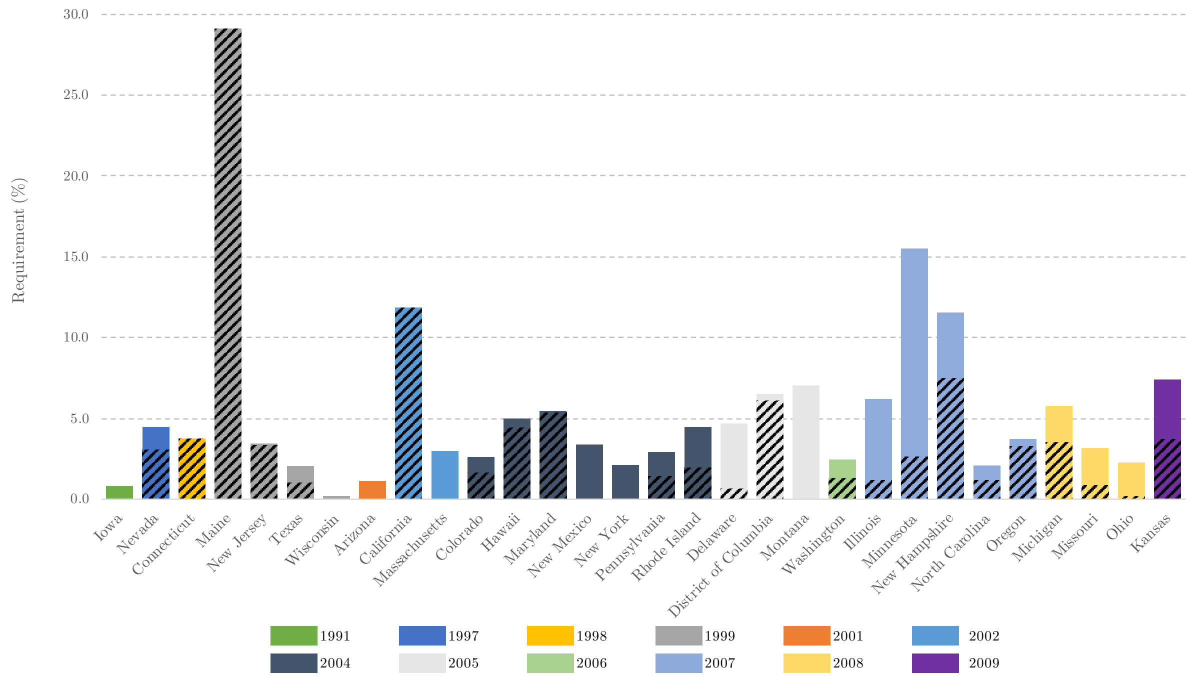

Notes: States are sorted by the year in which their RPS policies were first passed. The bars are colored according to RPS passage year. The total height of each bar denotes the gross RPS requirement at $\tau = 6$; the non-patterned portion of each bar denotes net requirement at $\tau = 6$. The data for gross RPS requirements are from the LBNL, in MWh, and are converted to percentages by dividing by contemporary generation at $\tau = 6$. Note that these percentages do not exactly equal the prescribed statutory percentages in the regulation.

adopted in a number of traditionally regulated markets. Figure 19 also plots real national average retail electricity prices (right y-axis) which declined from about 12 cents per kWh to 10 cents per kWh from 1990 through 2002 but by the end of the sample in 2015 returned to 12 cents per kWh.[67] This break in the decline in prices and subsequent upwards turn loosely corresponds with the number of states that passed RPS programs in those years. Whether this relationship is causal will be examined in much greater detail below.

[67]All monetary figures are reported in January 2019 dollars.

Figure 19: RPS Passage by State

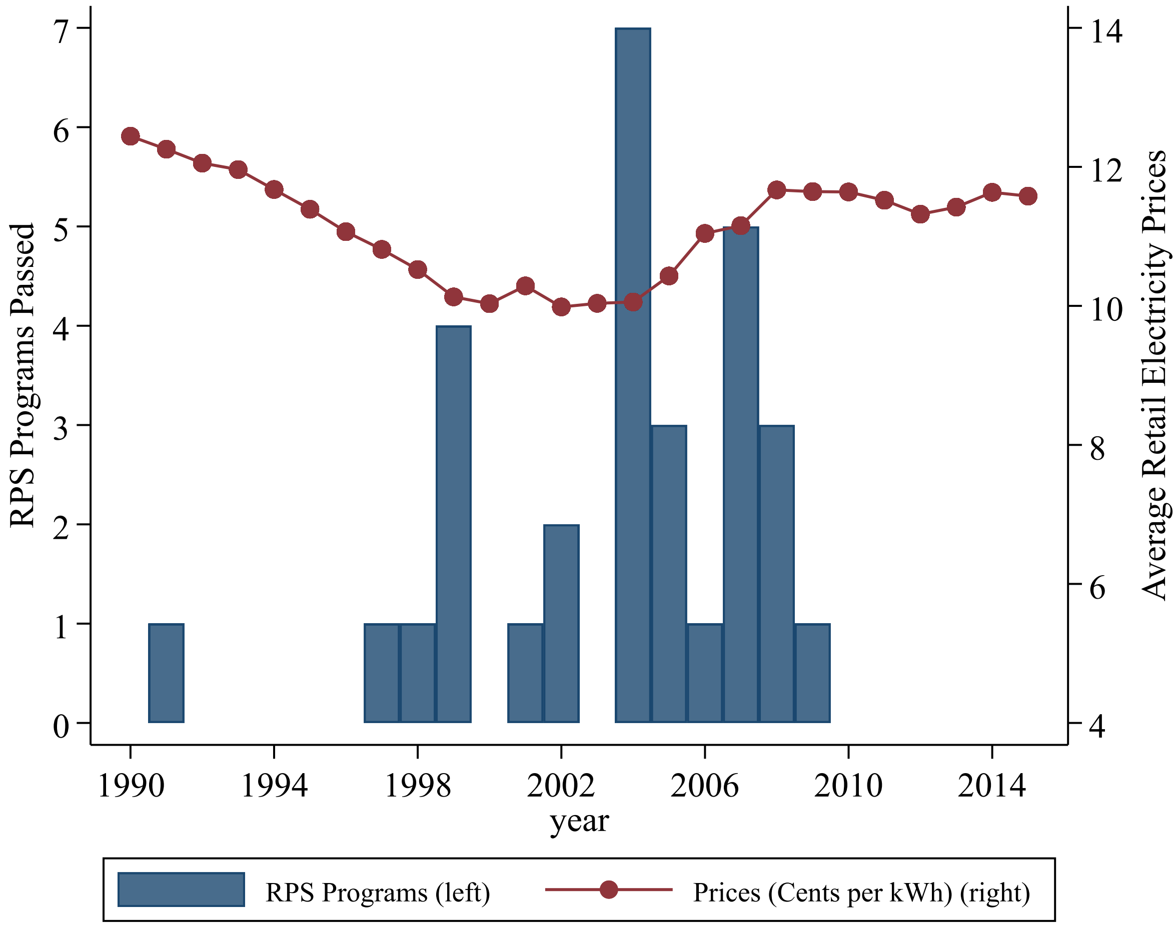

Notes: Year of RPS passage comes from the U.S. Department of Energy and state government websites. Price data is from the EIA.

3 Conceptual Framework

As discussed above, standard LCOE estimates measuring the direct capital and maintenance costs of various generation sources provide an incomplete summary of the impact of transitioning electricity production to renewable sources on consumer prices. We set out a simplified model of the decision-making process of a retail electricity provider to illustrate the mechanisms through which renewable integration can raise costs, and consequently retail prices. The model demonstrates how intermittency, transmission, and the displacement of existing capacity infrastructure interact to raise the total costs incurred by a utility. Notably, the model highlights the wide range of parameters and nontransparent data inputs that would be required to calculate these costs directly. The paper's empirical procedure sidesteps this difficulty by summarizing the aggregate effect of these mechanisms through the reduced-form impact of RPS standards on retail electricity prices.

For simplicity, the model assumes a vertically integrated setting with a single utility responsible for both power capacity and retail provision. The intuition from this framework translates straightforwardly to a deregulated setting with a retail provider purchasing electricity from competing generators, except for the assumption that ratepayers always pay the full cost of installed capacity. As discussed below, the extent to which owners of capital bear the losses from excess capacity stranded by integrating renewable sources will be one factor that contributes to the overall effect on retail prices.

3.1 Representative Utility Model

A representative utility chooses capacity investments and daily generation sources to fulfill two requirements: ensuring that they meet the full electricity demand of their customers every day and that their annual electricity production meets the RPS requirement.[68] Utilities have three types of production capacity available with which to meet daily electricity demand: renewables, R, baseload power, B, and dispatchable peaker plants, D, the latter two of which we assume come from non-renewable sources. Baseload generation produces a constant daily amount governed by annual capacity, B_t, and cannot be adjusted in response to daily demand. Renewable generation is stochastic and drawn from a distribution $F(R)$, with mean, $\tilde{R}$, standard deviation, σ_R, and support $[\underline{R},\ \overline{R}]$. $F(R)$ is a function of installed renewable capacity, R_t. The daily demand for electricity is also drawn from a distribution, $G(E)$, with mean $\tilde{E}$, standard deviation σ_E, and support $[\underline{E},\ \overline{E}]$. So given the available capacity of B_t, R_t, and D_t in year t, the utility observes the daily draws of E_s and R_s and chooses the level of dispatchable power, D_s, to satisfy customer

[68]The period in which instantaneous demand must be met can equivalently be thought of as an hour rather than a day.

demand.

$$E_s = B_t + R_s + D_s, \tag{25}$$
$$E_s \sim G(E_t),\ R_s \sim F(R_t)$$

With knowledge of this daily optimization problem, the utility chooses investment in new capacity at the beginning of each year. Total capacity in period t consists of the depreciated capital from last period plus new investments in each of the three categories of electricity sources:

$$C_t = B_{t-1}(1-\delta_B) + R_{t-1}(1-\delta_R) + D_{t-1}(1-\delta_D) + I_B + I_R + I_P \tag{26}$$

The utility chooses annual investments in new capacity to fulfill its two primary requirements. First, the RPS requirement dictates the proportion of annual electricity production that must come from renewables. For mandated renewable percentage, M, the utility must satisfy the following:

$$\frac{\sum_{s=1}^{365} R_s}{\sum_{s=1}^{365} E_s} \geq M \tag{27}$$

Under RPS requirements, failure to meet this condition will cost the utility a per-unit fine, f, for the amount by which renewable generation falls below the threshold. To avoid paying the fine, utilities must have enough installed renewable capacity, R_t, to produce enough electricity to meet this requirement. Determining what constitutes enough renewable capacity also may not be straightforward. If draws from the $F(R)$ distribution are correlated across days, simply ensuring that $\frac{E[R_s]}{E[E_s]} = M$ might not be sufficient to ensure compliance with the RPS mandate in a year with systematically low realizations for renewable generation. The utility will trade off the cost of increasing renewable capacity, R_t, with investments, I_R, against the fine for noncompliance when making their choice over optimal R_t.

Second, the utility must ensure it can supply enough energy every day of the year. We assume there is an infinite penalty for failing to meet demand. Since both energy demand and renewable production are stochastic, the utility must have enough dispatchable generation available to fill the largest possible daily need. In particular, the utility chooses D_t such that it can meet total electricity needs on a hypothetical day with the highest possible demand draw, $\overline{E}$, and the lowest possible renewable generation draw, $\underline{R}$.

$$D_t = \overline{E} - B_t - \underline{R} \tag{28}$$

In addition to choosing investment, the utility also has the option to prematurely retire capacity at the beginning of each period. The carrying costs of retired capacity are lower and for simplicity we assume that capacity that has not been retired will be run. Under certain conditions, they may choose to retire baseload capacity because too much baseload generation could prevent the utility from meeting the RPS requirement. If $\frac{B_t}{E[E_s]} > 1 - M$, for instance, then renewable production would be expected not to meet its mandate even without any dispatchable production. To ensure compliance with the RPS mandate, the utility must estimate the amount of dispatchable production necessary during the year and then scale back B_t such that the expected sum of baseload and dispatchable generation does not exceed $1 - M$ as a proportion of all production.

Total costs for the utility include the fixed costs of installed capacity, associated transmission and distribution requirements, and the variable costs associated with each type of power. The utility finances new investments such that they make a constant annual payment over a horizon of T years. The annualized prices of installed capacity, p_B, p_R, and p_D, incorporate differences in the cost per MWh for baseload, dispatchable, and renewable sources, as well as any differences in financing costs or investment tax incentives. New transmission investments in each period, which are also financed over a T-year horizon with annualized payment p_T, are a function of new installations across the three categories and depreciation of the existing transmission capital stock, with geographically dispersed renewable installations such as wind and solar likely having greater associated requirements. Since renewables require no fuel inputs, they incur no variable costs whereas baseload and dispatchable power have average costs ac_B and ac_P for each unit generated. For the purposes of this model, these average costs capture not only the cost of fuel inputs, but also any startup and shutdown costs associated with the operation of these generating sources. Thus, the utility's total costs in period t are as follows:

$$\begin{aligned} TC_t = &\sum_{k=t-T}^{t} p_{Bk} I_{Bk} + \sum_{k=t-T}^{t} p_{Dk} I_{Dk} + \sum_{k=t-T}^{t} p_{Rk} I_{Rk} \\ &+ \sum_{k=t-T}^{t} p_{Tk} Tr(I_{Rk}, I_{Bk}, I_{Dk}) + 365 B_t ac_B + \sum_{s=1}^{365} D_s ac_D \end{aligned} \tag{29}$$

The retail rate is given by total costs in period t divided by total kilowatt-hours of energy produced plus a markup, μ, assigned by the regulator. Thus:

$$\text{Retail Rate in Year t} = (1+\mu)\frac{TC_t}{\sum_{s=1}^{365} E_{st}} \tag{30}$$

3.2 Empirical Requirements for Estimating the Full Costs of RPS

This framework illustrates the major practical difficulties involved in developing the costs of RPS programs piece-by-piece. This simplified model reveals that even if renewable and non-renewable production have the same LCOE, defined by the prices of installed capacity and fuel inputs, transitioning a mature grid infrastructure could increase costs through a wide variety of channels. The list of excess costs includes:

- investments in new dispatchable capacity to protect against shortfalls of intermittent renewable generation,
- investments in new transmission infrastructure to accommodate the geographic locations of new renewable capacity,
- premature retirements of baseload capacity and/or transmission infrastructure that serves nonrenewables to reduce nonrenewable production enough to meet RPS mandates.

Further, the incidence of this last category between ratepayers and owners of capital is unclear ex ante, although ratepayers seem more likely to bear the costs in traditional regulated cost-plus markets, compared to restructured ones. Regardless of the ultimate incidence, these costs are part of the full costs of the introduction of a RPS program. However, it is worth noting that this last category category is "transitional" in nature, while the first two are permanent features of increasing renewables' share of production.

It is instructive to consider the challenges with constructing a bottom-up or structural estimate of the costs of an RPS policy. First, it would require data or estimates of several moments from the distributions of daily energy demand, $G(E_t)$, and daily renewable generation, $F(R_t)$, the pre-existing level of installed capacity by generation type, B_t,D_t,R_t, the respective depreciation rates, investment prices, and fuel input prices for each of these three energy categories, and the transmission investments necessary to incorporate renewable capacity. Second, the estimates would need to make a series of assumptions for how utilities project electricity demand, renewable intermittency, the need for dispatchable generation to protect against insufficient or excess supply, as well as the decision criteria for retiring baseload generation. Third, estimating the model would require going beyond the representative utility setup and incorporating interactions between heterogeneous generators and retail providers in restructured and non-restructured markets; these interactions have proven to be quite complex to model as they involve questions of market power and doing so in this context would undoubtedly be both a great research topic and a difficult problem to solve. Fourth, the incidence of these costs between ratepayers and owners of capital is also a complicated question and, as we noted above, is likely affected by the regulatory environment.

Our approach circumvents this complex interplay of underlying mechanisms with a reduced-form approach that captures the costs imposed on ratepayers due to all potential mechanisms through which RPS policies raise costs. If generators or distributors bear part of the costs, our approach will not capture the full social costs of RPS policies.

Finally, we note that coincident to the increase in the number of RPS programs and the scope of their ambitions, there have been important changes in the operation of electricity markets. As one example, several Regional Transmission Organizations began holding centralized auctions for capacity market payments in the mid-2000s as RPS programs began to proliferate.[69] Since their initiation, these payments comprise a substantial fraction of overall market costs - reaching 9-28% of total costs in ISO New England, the New York ISO, and the PJM Interconnection between 2008 and 2016 (GAO, 2017). These three RTOs cover all or part of 15 of the 29 RPS-adopting states in our main sample, and the Mid-continent Independent System Operator (MISO) added a capacity market auction covering 4 more RPS states in 2013. Further, Bushnell, Flagg and Mansur (2017) document that similar payments to maintain "Resource Adequacy" take place in other locations as well and that they likely also preceded the centralized auctions in those four RTOs. We take the significant share of these types of payments for generator availability after RPS implementation, be it through auctions or resource adequacy payments, as suggestive evidence that RPS program's mandated increase in intermittent renewable generation imposes systemwide costs on electricity markets, very likely due to these technologies' intermittency.

Thus, it is at least plausible that an important share of RPS programs' total costs comes from the indirect costs that they impose on the electricity supply system. These costs are not evident from a simple comparison of LCOEs or RECs prices. Of course, the qualitative evidence about the growth of capacity markets or resource adequacy payments is not decisive and could be due to other factors, so the remainder of the paper exploits a differences in differences research design that is generated by the staggered adoption of RPS programs by some states and the non-adoption by other states.

4 Data Sources and Summary Statistics

In order to assess the retail price and other impacts of RPS programs, we construct a state-level panel from 1990 to 2015 with data on RPS programs, electricity prices, generation capacity and outcomes, and CO_2 emissions. We believe this is the most comprehensive data set ever compiled on RPS program characteristics and potential outcomes. This section describes each data source and presents some summary statistics describing the context of the policy.

[69]ISO-NY began their current system of auctions in 2003, PJM in 2004, NE-ISO in 2007, and MISO in 2013.

4.1 RPS Program Data

Since 1990, 29 states and the District of Columbia have adopted RPS programs. We construct indicators for the year in which legislation for a mandatory RPS program first passed in each state, compiled using state legislative documents, state government websites, and summaries from the U.S. Department of Energy. While there is typically a few years' lag between policy enactment and the onset of binding mandates for renewable generation, costs to electricity providers, and consequently customers, are likely to begin accruing when they start planning for and investing in the required future capacity. Data from the Lawrence Berkeley National Laboratory (LBNL) also include information about qualifying renewable sources under each program, including whether there are specific requirements for solar generation.

To better characterize each state's implementation, we also collect more detailed information on year-by-year requirements. Most RPS programs require an increasing percentage of electricity sales to come from renewable sources, leading to increased stringency over time.[70]

However, as mentioned earlier, the statutory percentage requirement may overstate the additional generation required if a large number of existing generators are eligible for compliance. To account for this, we construct a measure of RPS net requirements as the difference between statutory requirements and pre-existing renewable generation. We collect data from LBNL on total generation required from renewables in each RPS state in each year of enforcement (Barbose, 2018), and define pre-existing compliance as total generation from qualifying categories of renewables in the year before RPS legislation was passed. The difference is the amount by which each state had to expand renewable generation to comply with the policy - our measure of net requirements.[71]

Recall, Figure 18 highlighted the substantial differences between the total and net require-

[70]Iowa and Texas have fixed capacity requirements for new renewable generation, which will tend to decrease stringency over time if demand is increasing.

[71]Some states include waste-to-energy and similar forms of power generation in their RPS, but we do not have data for these sources, so we cannot account for these in our net generation estimates. In addition, there are 3 states - California, Montana, and Minnesota - for which the gross MWh requirements reported by LBNL differ by more than 3 percentage points from the statutory percentages reported by DSIRE.

ments. In addition to data on RPS programs, we also collect information from the North Carolina Clean Energy Center's Database of State Incentives for Renewables Efficiency (DSIRE) on the presence of other state programs that may influence the amount of renewable generation and the retail price of electricity (Barnes, 2014). We have information on the implementation dates of three types of programs: net metering, which pays consumers for electricity they add to the grid with distributed generation such as solar PV, green power purchasing, which gives consumers the option of paying to have renewable energy account for a certain percentage of their consumption, and public benefits funds, which place a surcharge on retail electricity prices to fund programs such as research and development, energy efficiency investments, and low-income energy assistance. This information is used to account for the presence of potentially confounding programs.

4.2 Electricity Sector

Information on electricity sector variables is drawn from Energy Information Administration (EIA) survey forms. Electricity prices are computed from EIA Form 861, a mandatory census of retail sales by electric power industry participants.[72] Respondents report sales and revenues separately for commercial, industrial, and residential sectors. Average price is then computed based on average revenue per megawatt-hour sold for each sector and for total retail sales.

Electricity generation by state and fuel source is compiled from EIA forms 906, 920, and 923, which concern power plant operations. This data is broken down by fuel type, ensuring plants with multiple fuel sources are accurately reflected in aggregate numbers. Generating capacity by state and fuel source is compiled from EIA Forms 860 and 867, along with starting year and month and location. These surveys cover all grid-connected generators larger than 1 MW in capacity currently able to deliver power. For simplicity, we aggregate the EIA's fuel type categories, measuring generation by hydroelectric, solar, wind, coal, natural gas, nuclear, other renewables, and other fuels.[73]

[72] The 3,300 respondents cover essentially the universe of retail suppliers, including electric utilities, energy service providers, power marketers, and other electric power suppliers.

[73] "Other Renewables" includes biomass, geothermal, and wood-based fuels, while Other covers remaining sources, including pumped storage, blast furnace gas, and other marginal fuels. See the

To measure CO_2 emissions, we use estimates derived by the EIA from power plant operations data taken from forms 767, 906, and 923. Their estimation process involves converting fuel use to BTUs to provide a common comparison measure. Next, fuel uses that do not generate emissions are subtracted out. Finally, source-specific carbon emission coefficients are used to convert to metric tons of carbon.[74] The result is a yearly panel of state emissions from electricity generation.

As part of our analysis, we also attempt to look at the difference between RPS impacts in regulated versus deregulated markets. Using data compiled for an earlier paper by Fabrizio, Rose and Wolfram (2007), we code an indicator for whether or not a state ever deregulates their electricity market, defined by retail market access for non-utility-owned generation plants.[75]

4.3 Manufacturing Employment

If RPS programs do in fact raise electricity prices, there may be downstream impacts on industries for which energy is a large input to production. To assess this, we construct a panel of employment in each state by industry code using data from the County Business Patterns (CBP). One issue with these data is that employment statistics are often suppressed when the industry code and establishment size potentially disclose information about a specific business. Following previous papers, we apply an imputation procedure to estimate employment for these cells, using the national average for the industry in that cell size. To allow comparisons across years, we recode NAICS industry codes used in later years to SIC industry codes, redistributing employment proportionally based on concordances provided by the census.[76] We then calculate total and manufacturing employment for each state in each year.

Electric Power Monthly published by the EIA for a full accounting of possible disaggregated fuel sources.

[74]More details on this process, including the conversion factors used, can be found in "Methodology and Sources" section of the *Monthly Electric Review* published by the EIA.

[75]We thank Fabrizio, Rose, and Wolfram for generously sharing this data.

[76]For further details, and code used, see Autor, Dorn and Hanson (2013a) and the accompanying data files. For 2012 and 2013, where official concordances are unavailable, we allocate employment proportionally based on 2011 employment using the official code mapping 2012 to 2007 NAICS.

4.4 Summary Statistics

Before describing our empirical approach in detail, we briefly present some summary statistics from the data and report on some comparisons of treatment and control states in the year prior to RPS passage. Table 13 presents summary statistics for treatment states, defined as those in which legislation passes in the following year, and control states, which consist of states that did not pass RPS by 2015. The summary statistics for control states are averaged across the set of control states that correspond to each RPS state's year of passage.

The statistics in Table 13 show some level differences between RPS states and control states in the year prior to legislation. RPS states tend to have somewhat more expensive electricity - 11.4 cents per kWh versus 9.4 in control states - larger populations, and better resources for producing solar and wind energy. Such level differences do not threaten the identification of our difference-in-differences design, but may be informative about the degree to which our results would be representative of the impact of a national RPS policy. The RPS states in our analysis are also more likely to have other simultaneous programs affecting renewable energy, including public benefit funds, net metering, and green power purchasing programs. We control for the time-varying passage of these programs at the state by year level in our analysis. Finally, we note that the pre-existing trends of electricity prices in treatment and control states are similar, with an average six-year decrease in electricity prices of 0.6 cents per kWh in both RPS states and control states prior to the year of passage. Our analysis in the next section will control for differences in pre-trends, but the similarity of these trends lends validity to the key identification assumption of equal trends in electricity prices in RPS and non-RPS states in the years before RPS passage.

5 Empirical Strategy

Our empirical approach begins with an event study-style equation:

$$y_{st} = \alpha + \sum_{\tau=-19}^{18} \sigma_\tau D_{\tau,st} + X_{st} + \gamma_s + \mu_t + \epsilon_{st} \tag{31}$$

Table 13: Summary Statistics

	Mean RPS (1)	Mean Control (2)	P-value RPS vs Control (3)
Price (2018 Cents/kWh)			
Total	11.4	9.4	0.01
Residential	13.4	11.3	0.01
Commercial	11.8	9.8	0.01
Industrial	8.5	6.9	0.01
Price Change τ = -1 vs -7 (2018 Cents/kWh	-0.6	-0.6	0.92
Total Sales (TWh)	76.2	64.3	0.38
Population (Millions)	7.0	4.7	0.11
CO_2 Emissions (Million mt)	48.0	49.2	0.90
Renewable Potential (PWh)			
Solar	9.1	6.6	0.34
Wind	1.1	0.9	0.40
Generation			
Total (TWh)	80.5	73.3	0.64
RPS Eligible (TWh)	8.9	5.9	0.36
RPS Eligible (% of Total)	13.5	13.0	0.89
Generating Capacity			
Total (GW)	20.3	18.4	0.60
RPS Eligible (GW)	2.5	1.6	0.36
RPS Eligible (% of Total)	14.2	14.3	0.99
Other Programs (%)			
Public Benefit Funds	0.41	0.11	0.00
Net Metering	0.66	0.45	0.04
Green Power Purchasing	0.07	0.02	0.29
Energy Efficiency	0.03	0.03	0.91
Restructuring	0.59	0.25	0.00

Notes: “Mean RPS” is for RPS states in the year prior to RPS passage. A control is defined for each RPS state as the mean across non-RPS states and RPS states that have yet to pass RPS, in the year prior to the reference RPS states RPS passage. “Mean Control” is the average across these controls. Column (3) reports p-values from a two-sample t-test between Column (1) and (2) that allows for unequal variances across groups. Iowa is excluded from these summary statistics due to the particularly early passage of its RPS.

where y_{st} is an outcome of interest in state s in year t. We include state fixed effects γ_s to control for any permanent, unobserved differences across states. Year fixed effects, μ_t, non-parametrically control for national trends in retail prices. The variables $D_{\tau,st}$ are separate indicators for each year τ relative to the passage of a RPS law, where τ is normalized to equal zero in the year that the program passed; they range from -19 through 18, which covers the full range of values of the τ's.[77] For states that never adopt an RPS program, all $D_{\tau,st}$ are set equal to zero. As non-adopters, they do not play a role in the estimation of the τ's but they aid in the estimation of the year effects, μ_t, as well as the constant, α.

The σ_τ's are the parameters of interest as they report the annual mean of the outcome variable in event time, after adjustment for state and year fixed effects. An appealing feature of this design is that because states passed RPS programs into law in different calendar years, it is possible to separately identify the σ_τ's and the year fixed effects μ_t. In the remainder of the analysis, we will particularly focus on the σ_τ's that range from -7 through 6. This is the maximum range for which the σ_τ's can all be estimated from a fixed set of states. Restricting the treatment period in this way holds the advantage of eliminating questions about the role that differences in the composition of states identifying the various σ_τ's plays. This range is determined by Nevada, which passed its law in 1997 on one side of the range, and Kansas, which passed its law in 2009 on the other side of the range. We will present event-study figures that plot the estimated σ_τ's against τ. These figures provide an opportunity to visually assess whether there are differential trends in the outcome variables prior to RPS passage, which help to assess the validity of the difference in differences identification strategy. The event-study figures also demonstrate whether any impacts on outcomes emerge immediately or over time, which will inform the choice of specification to summarize the average effect of RPS policies.

To summarize the information contained in the event-study plots and formally assess the program impact, we estimate two equations. In the first, we assume that the difference in differences' identification assumption of parallel trends holds and allow for RPS programs to have

[77]Iowa adopted a RPS in 1991, which means that only one pre-RPS year is available. Consequently, we drop Iowa from the primary sample although its inclusion does not alter the qualitative findings.

only a mean-shift effect on retail electricity price:

$$
\begin{aligned}
y_{st} = \delta_0 + \delta_1 1(-19 \leq \tau \leq -8)_{st} * 1(\text{RPS} = 1)_s + \delta_2 1(7 \leq \tau \leq 18)_{st} * 1(\text{RPS} = 1)_s \\
+\delta_3 1(0 \leq \tau \leq 6)_{st} * 1(\text{RPS} = 1)_s + X_{st} + \gamma_s + \mu_t + \epsilon_{st}
\end{aligned}
\tag{32}
$$

Here, the parameter of interest is δ_3, which measures the mean of the outcome variable in the first 7 years after the passage of RPS policies, in RPS states, relative to the preceding 7 years, after adjustment for state and year fixed effects. The coefficients δ_1 and δ_2 measure the mean of the outcome in the unbalanced samples in the years before and after the 14 year period where the sample is balanced, in RPS states. These are nuisance parameters.

Most RPS programs have requirements that increase gradually over time after legislation is passed, so it is likely that the impact on electricity prices will increase correspondingly. Therefore, a specification like a trend break model seems better equipped to summarize the effect of RPS programs on outcomes because it allows the programs' effect to grow over time. Further, specifications that allow for the possibility of differences in pre-adoption trends require weaker assumptions to produce valid estimates of the impact of RPS programs. For these reasons, we also fit an equation that allows for differential trends before and after RPS programs are passed into law:

$$
\begin{aligned}
y_{st} = \delta_0 + \delta_1 1(-19 \leq \tau \leq -8)_{st} * 1(\text{RPS} = 1)_s + \delta_2 1(7 \leq \tau \leq 18)_{st} * 1(\text{RPS} = 1)_s \\
+\delta_3 1(0 \leq \tau \leq 6)_{st} * 1(\text{RPS} = 1)_s + \beta_0 \tau_{st} + \beta_1 1(-19 \leq \tau \leq -8)_{st} * 1(\text{RPS} = 1)_s * \tau_{st} \\
+\beta_2 1(7 \leq \tau \leq 18)_{st} * 1(\text{RPS} = 1)_s * \tau_{st} + \beta_3 1(0 \leq \tau \leq 6)_{st} * 1(\text{RPS} = 1)_s * \tau_{st} \\
+X_{st} + \gamma_s + \mu_t + \epsilon_{st}
\end{aligned}
\tag{33}
$$

To summarize the cumulative effects, we calculate and report the impact seven years after RPS passage, which is given by δ_3+$6\beta_3$. Finally, we report standard errors that are clustered by state from the estimation of Equations 32 and 33 to allow for correlation in the errors within state over time.

6 Results

6.1 Net RPS Requirements and Retail Electricity Prices

We begin with an examination of the net RPS requirements. Figure 20 plots the event-year means of net RPS requirements against τ. Recall that event time is normalized so that the program passage year occurs at $\tau = 0$ and the vertical line at $\tau = -1$ indicates the last year before program passage. It is apparent that the RPS programs' passage into law leads to increases in the required use of the RPS technologies that begin almost immediately and increase every year. Seven years after passage, the average RPS state's net requirement is 1.8 percentage points of generation. It is noteworthy that this is substantially smaller than the increase in total or gross requirement which is 5.1%. through the end of the balanced sample which is 7 years later.

Figure 20: Estimated Effects of RPS Programs on Net Renewable Requirements

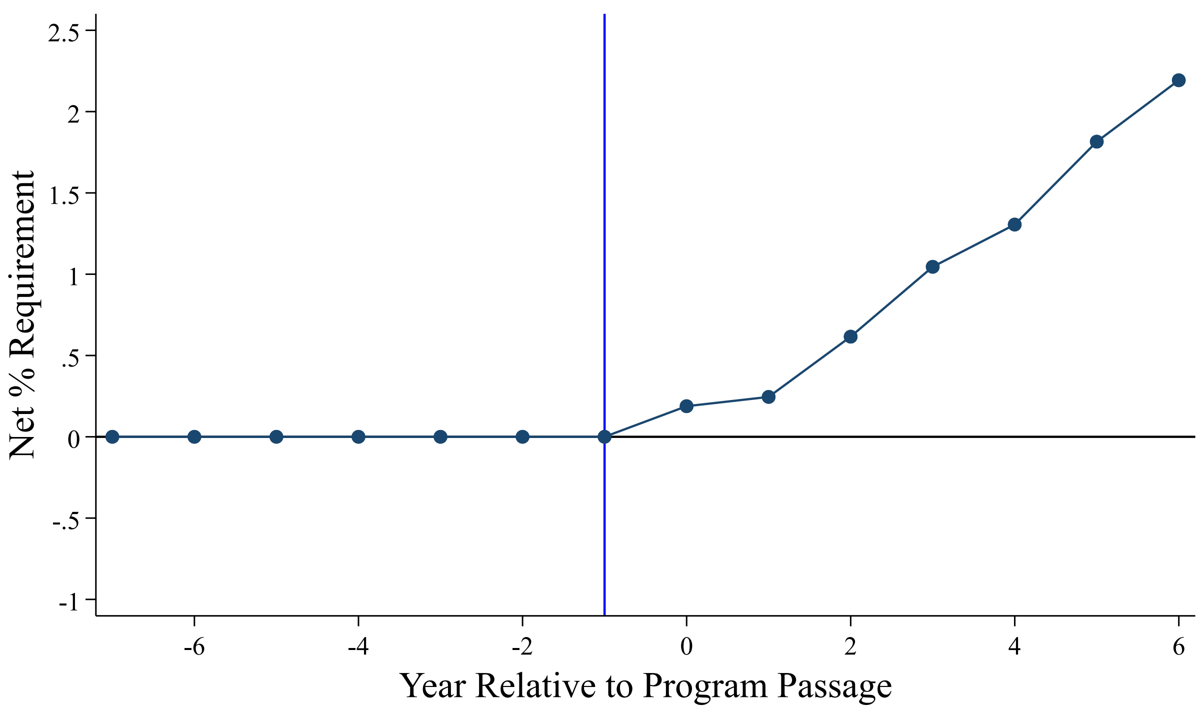

Notes: Graph shows the mean net RPS requirement percentage for event years τ = -7 to τ = 6. Gross RPS requirement data are from the LBNL. RPS program passage dates and requirements are from the Department of Energy and state government websites.

Figure 21: Estimated Effects of RPS Programs on Retail Electricity Prices

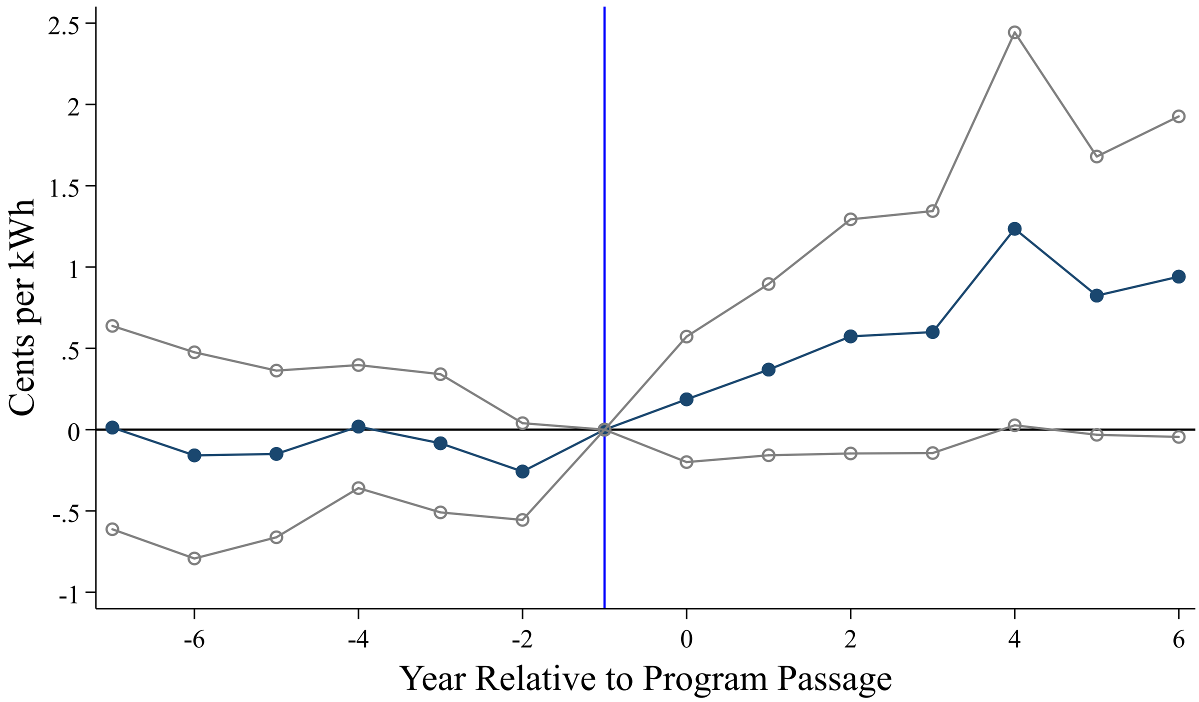

Notes: Graph shows coefficients for σ_τ for τ = -7 to τ = 6 from the event study specification in Equation (7) for retail electricity prices on indicator variables for years relative to program passage, controlling for state, year, and other programs fixed effects. Blue lines show the point estimates and gray lines contain the 95% confidence interval. Electricity price data and electricity generation for calculating net requirement are from the EIA. RPS program passage dates and requirements are from the Department of Energy and state government websites. Standard errors are clustered at the state-level.

Figure 21 reports on the estimation of equation 31 for the average retail price. where prices are normalized so that they equal zero at $\tau = -1$. Recall, the estimated σ_τ's are adjusted for state and year fixed effects. There are two primary points that emerge. First, there is no evidence of a meaningful difference in the trends of prices, either upwards or downwards, among adopting states in the six years preceding RPS programs becoming law, from $\tau = -7$ to $\tau = -1$. Thus, for example, there doesn't appear to be any evidence that prior to RPS passage, adopting states were differentially passing other policies that influence electricity prices positively or negatively or facing differential cost shocks. More broadly, this figure supports the validity of the difference in differences research design. Second, it is apparent that retail prices increased after program passage, but not all at once; the figure suggests that a model that allows for a trend break describes the data well. It is striking that the trend in prices appears to very closely shadow the trend in net RPS requirements.

Columns (1a) and (1b) in Panel A of Table 14 present results from the estimation of equations 32 and 33 that confirm the visual impression that retail electricity prices increase after RPS programs become law. The mean-shift specification suggests that RPS programs raised prices by 0.5 cents on average in their first 7 years. In the mean shift and trend-break model, the estimates indicate that retail prices in RPS states rise by roughly 0.16 cents each year post-passage, with statistically insignificant pre-trends and post-passage mean-shift.

Given these results and the visual event-study evidence suggesting that RPS programs affect the trend in prices, we treat Equation 33 as our primary specification. We focus on the effect 7 years after RPS passage, which is calculated as $\delta_3 + 6\beta_3$. Overall, the estimates from this regression suggest that RPS policies have increased retail electricity prices by about 1.3 cents per kWh seven years after passage. This increase is statistically significant and economically substantial, representing an increase of about 11.1% over the mean retail price at $\tau = -1$. Such a large increase in the retail price of electricity is striking, given the modest net requirements 7 years after passage. Further, these estimates are much larger than LCOE differences alone would suggest, indicating that the indirect costs of RPS mandates are an important component of their total costs.

Table 14: Estimates of RPS Impact on Retail Electricity Prices

	Average Retail Price		Average Retail Price Residential		Average Retail Price Commercial		Average Retail Price Industrial	
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)	(4a)	(4b)
Panel A: 7 Post-Passage Years, Balanced Sample								
Mean Shift (δ_3)	0.69*	0.37	0.57	0.21	0.73*	0.40*	0.83*	0.81*
	(0.41)	(0.24)	(0.43)	(0.24)	(0.41)	(0.23)	(0.45)	(0.47)
Trend Break (β_3)		0.15*		0.22**		0.09		0.02
		(0.09)		(0.09)		(0.09)		(0.10)
Effect of RPS 7 years after passage		1.24**		1.55**		0.92		0.94*
($6\beta_3 + \delta_3$)		(0.59)		(0.64)		(0.62)		(0.50)
Panel B: 12 Post-Passage Years, Unbalanced Sample								
Mean Shift (δ_3)	0.75*	0.40	0.66	0.27	0.77*	0.41	0.86*	0.68
	(0.45)	(0.28)	(0.47)	(0.28)	(0.44)	(0.27)	(0.48)	(0.41)
Trend Break (β_3)		0.14*		0.20**		0.09		0.08
		(0.07)		(0.08)		(0.08)		(0.08)
Effect of RPS 12 years after passage		1.94**		2.50**		1.41		1.57*
($11\beta_3 + \delta_3$)		(0.81)		(0.93)		(0.92)		(0.84)
Mean at τ = -1	11.4	11.4	13.4	13.4	11.8	11.8	8.5	8.5
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Programs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	1300	1300	1300	1300	1300	1300	1300	1300

Notes: Columns (1a) through (4b) show estimates from Equations (8) and (9), where (a) columns correspond to Equation (8) and (b) columns correspond to Equation (9), with total retail electricity price and sector-specific retail electricity prices as the response variables. In *Panel A*, coefficient estimates are for states with data 7 years before and 7 years after RPS passage. In *Panel B*, coefficient estimates are for states with data 7 years before and 12 years after RPS passage. Using Equation (9) notation, the effect of RPS 7 years after passage is $6\beta_3 + \delta_3$, and the effect of RPS 12 years after passage is $11\beta_3 + \delta_3$. Standard errors are clustered at the state-level. Asterisks denote p-values: < 0.10 (*), < 0.05 (**), < 0.01 (***).

We next consider whether RPS policies exhibit heterogeneous effects by the category of customer. The EIA divides retail sales among three sectors, residential, commercial, and industrial, that together account for total retail sales.[78] Residential is the largest sector for most years in our data, comprising about 37% of sales in 2015.[79] On a per-customer basis, though, the commercial and industrial consume significantly more. A typical commercial customer uses nearly seven times the typical residential consumption, while the typical industrial customer uses more than 120 times the typical residential consumption. As noted in Table 13, retail rates also vary among these groups, with residential customers paying the highest rates while industrial customers pay the lowest. This differentiated pricing may reflect demand elasticities that are correlated with usage, leading utilities to price discriminate by charging lower prices to their most intensive, and therefore price sensitive, customers (Bjørner, Togeby and Jensen, 2001).

The event-study figures derived from the fitting of Equation 31 for these outcomes are presented in Appendix Figure A-37. There is little evidence of difference in trends between adopting and non-adopting states prior to RPS passage. Industrial prices appear to shift upwards substantially in the first year after passage, while the commercial and residential sectors adjust more gradually. Overall, changes by sector track closely with net requirement changes, though perhaps with a slight lag.

The statistical sectoral price analyses for the balanced sample are reported in columns (2) - (4) of Panel A in Table 14. As in our analysis of total prices, sectoral prices appear best captured by the mean-shift and trend-break model, so we focus on estimates from Equation 33 in the (2b), (3b), and (4b) columns. In all three sectors, the point estimates represent substantial price increases in the first 7 years after RPS passage; they are 12.8% for residential, 7.7% for commercial, and 9.2% for industrial, although only the residential one would be judged statistically significant by conventional criteria.

[78]According the EIA, the sectors are composed of: Residential: “living quarters for private households,” Commercial: “service-providing facilities and equipment of: businesses; Federal, State, and local governments; and other private and public organizations,” Industrial: “all facilities and equipment used for producing, processing, or assembling goods.” For complete definitions, see the EIA’s Electric Power Monthly: http://www.eia.gov/electricity/monthly/.

[79]Authors’ calculation, from the EIA Electricity Data Browser.

The appeal of the Panel A results is that there is a balanced sample for all event years, but this sample restriction limits the number of post-years. In Panel B, we extend the post-period through $\tau = 11$ which allows us to estimate the effect of the RPS programs through 12 years after passage. However, the number of RPS states that reach $\tau = 11$ in the sample declines from 29 in the balanced sample to 16, so the cost is that there is not a constant sample of states for all event years.

The Panel B results tell much the same story of higher prices. As RPS programs are in force longer here, their net requirements increase and their impact on electricity prices also increases. The column (1b) estimates indicate that at twelve years after passage, the average retail price has increased by 2.0 cents per kWh or 17% and at the same point net RPS requirements have risen to 4.2 percentage points of generation (although gross or total RPS requirements are higher at 11.1 percentage points).[80] The remaining columns reveal that over this longer time horizon the higher electricity costs remain evident in all three sectors, with the residential sector experiencing the largest increase.

Table 15 explores the robustness of the Table 14 Panel A results to a variety of changes in Equation 33. In column (1), we drop the two states with nuclear in their original RPS (i.e., Massachusetts and Ohio) as these states' policies are closer to a zero carbon energy standard and in column (2) we drop Hawaii due to its unique geography. Neither of these sample restrictions meaningful change the qualitative findings. The remaining columns aim to adjust for the possibility of local shocks to electricity prices that might confound the adoption of RPS programs; specifically, columns (3) and (4) include year by census region and year by census division fixed effects, respectively. There are 4 Census regions and 9 Census divisions. The estimated increases in electricity prices are modestly smaller here than in Table 2, however the differences are small compared to the standard errors. Our conclusion is that these models that handle local shocks more flexibly leave the qualitative findings unchanged.

[80]Appendix Figures A-38 and A-40 present the accompanying extended period figures for net requirements and average retail prices. See Appendix Figure A-39 for a plot of gross, i.e. total, RPS requirements.

Table 15: Robustness Checks for RPS Impact

	Retail Electricity Price				
	(1)	(2)	(3)	(4)	(5)
Panel A: Total					
Effect of RPS 7 years after passage	1.37**	1.10*	1.19*	1.08**	0.96*
$(6\beta_3 + \delta_3)$	(0.60)	(0.61)	(0.60)	(0.53)	(0.52)
Panel B: Residential					
Effect of RPS 7 years after passage	1.80***	1.38**	1.49**	1.50***	1.38**
$(6\beta_3 + \delta_3)$	(0.67)	(0.66)	(0.64)	(0.53)	(0.52)
Panel C: Commercial					
Effect of RPS 7 years after passage	1.04	0.68	0.85	0.76	0.76
$(6\beta_3 + \delta_3)$	(0.63)	(0.62)	(0.62)	(0.53)	(0.54)
Panel D: Industrial					
Effect of RPS 7 years after passage	1.00*	0.87	0.90*	0.54	0.58
$(6\beta_3 + \delta_3)$	(0.51)	(0.53)	(0.50)	(0.56)	(0.62)
Other Programs		X	X	X	X
Other Programs and Energy Eff. Expenditures	X				
Excludes States with Nuclear in Original RPS		X			
Excludes Hawaii			X		
State FE	X	X	X	X	X
Year FE	X	X	X		
Year-Region FE				X	
Year-Division FE					X
N	1200	1248	1274	1300	1300

Notes: The (a) columns report the aggregate effect 7 years after RPS passage from the mean-shift model given by Equation (8). The (b) columns report the same effect from the trend-break model given by Equation (9). Coefficient estimates are for states with data 7 years before and 7 years after RPS passage. Year-Region fixed effects are for all combinations of years and Census regions. Year-Division fixed effects are for all combinations of years and Census divisions. The two states with nuclear in their original RPS are Massachusetts and Ohio. One specification excludes Hawaii due to its geographic isolation and thus its inability to trade electricity across state borders. Standard errors are clustered at the state-level. Asterisks denote p-values: < 0.10 (*), < 0.05 (**), < 0.01 (***).

6.2 Heterogeneity in RPS Price Effects

To this point, we have assumed that the effect of RPS programs on average retail prices are constant across states. However, there are several important characteristics that might differ across states and could affect the magnitude of the impact of RPS programs or their incidence

Table 16: Heterogeneous Effects of RPS Programs on Retail Electricity Prices

	Total	Residential
Panel A: Late Adopters		
Effect of RPS 7 years after passage	1.20	1.40
($6\beta_3 + \delta_3$)	(0.81)	(0.92)
(Effect of RPS)*Late	-0.22	0.18
	(1.47)	(1.50)
Panel B: Ever Restructured		
Effect of RPS 7 years after passage	2.04*	2.35*
($6\beta_3 + \delta_3$)	(1.14)	(1.20)
(Effect of RPS)*Restructured	-0.92	-0.85
	(1.32)	(1.42)
Panel C: Has Solar Set-Aside		
Effect of RPS 7 years after passage	0.84	1.13
($6\beta_3 + \delta_3$)	(0.78)	(0.97)
(Effect of RPS)*Solar Set-Aside	0.96	1.04
	(1.23)	(1.26)
Panel D: Heavy Coal States		
Effect of RPS 7 years after passage	0.83	1.26
($6\beta_3 + \delta_3$)	(0.86)	(1.01)
(Effect of RPS)*Heavy Coal	0.80	0.64
	(1.21)	(1.30)
State FE	Yes	Yes
Year FE	Yes	Yes
Other Programs	Yes	Yes
N	1300	1300

Notes: The coefficients give the aggregate effect of RPS programs on total and residential retail prices 7 years after passage estimated from the trend-break model. The top row in each panel shows the coefficient for the subset of states *not* in the given category and the bottom row shows the difference in the coefficient for the given subset. All coefficient estimates are for states with data 7 years before and 7 years after RPS passage. Using Equation (9) notation, the effect of RPS 7 years after passage is $6\beta_3 + \delta_3$. Standard errors are clustered at the state-level. Asterisks denote p-values: < 0.10 (*), < 0.05 (**), < 0.01 (***).

on ratepayers versus owners of capital. This subsection explores this possibility by taking the trend and mean shift model (i.e., Equation 33) and fully interacting it with an indicator for membership in a subsample of interest. Table 16 presents the results from this exercise for average retail prices and residential retail prices by reporting the effect of RPS programs among states not in the subsample and the marginal effect for the subsample. The latter estimate tests whether the seven year effect differs in the subgroup of interest and the full effect for this group is the sum of the two reported estimates.

Panel A examines whether RPS program effects differ for late adopters, defined as those with laws that were passed after 2004, the median year of passage in the data. This specification tests the hypothesis that the costs of RPS programs might be lower in the later years of the sample, perhaps due to decreasing costs for renewable energy or learning about how to more efficiently integrate renewables into the grid. Panel B explores differential impacts among states that have restructured electricity markets. Panel C examines the effect of setting specific requirements that can only be fulfilled by solar energy, which restricts flexibility to use the cheapest available renewable resource. Panel D estimates effects for heavy coal states, defined as those above median percentage coal generation in 1990, to test whether these states encounter higher costs to incorporating renewables.

This type of subgroup analysis is very demanding of the data, but some intriguing, albeit suggestive, patterns emerge. There is little evidence to support the hypothesis that the costs for ratepayers were lower in late (i.e., post-2004) adopting states. The point estimates in Panel B indicate that the impact on prices is smaller in states where electricity markets have been restructured, which is consistent with the possibility that it is easier to pass on the costs of stranded assets to ratepayers in vertically integrated non-restructured settings. However, the magnitude of the standard errors warrants caution in drawing strong conclusions. The point estimates suggest that solar set asides substantially increase prices, which is consistent with the fact that solar REC prices can be several times larger than general REC prices, but here too the imprecision of the estimates tempers the strength of any conclusions. Finally, the costs appear to be higher in heavy coal states, but the same problem of noisy estimates is evident.

Table 17: RPS Effect on Sales and Employment

	Sales		Employment			
	Total		Total	Total	Manufacturing	Manufacturing
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Mean Shift (δ_3)	-0.01		0.008		-0.026	
	(0.02)		(0.015)		(0.023)	
Effect of RPS 7 years after passage		0.01		0.036		-0.014
($6\beta_3 + \delta_3$)		(0.03)		(0.024)		(0.039)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Programs	Yes	Yes	Yes	Yes	Yes	Yes
N	1300	1300	1200	1200	1200	1200

Notes: The dependent variable in Columns (1a) and (1b) is the log of total sales in MWh. The dependent variable in Columns (2a) and (2b) is the log of total employment in each state; in Column (3a) and (3b) is log manufacturing employment. The (a)-columns show the mean-shift estimates from Equation (8) for sales or employment. The (b)-columns report the aggregate effect 7 years after program passage from the trend-break model given by Equation (9). Coefficient estimates are for states with data 7 years before and 7 years after RPS passage. Using Equation (9) notation, the effect of RPS 7 years after passage is $6\beta_3 + \delta_3$. Standard errors are clustered at the state-level. Asterisks denote p-values: < 0.10 (*), < 0.05 (**), < 0.01 (***).

6.3 Economic Activity

Since the estimates suggest that RPS programs lead to substantial increases in electricity prices, it is natural to examine whether there are impacts on the real economy. We begin by testing whether electricity consumption responds to these increases. Some previous studies suggest that consumers typically appear to be responsive to the average prices, which is our variable of interest, rather than marginal prices, potentially due to inadequate real-time information about current consumption (Borenstein (2009); Ito (2014)). In columns (1a) and (1b) of Table 17, there is little evidence of a change in electricity consumption.

The remaining columns of Table 17 report on the estimation of the same equations for total employment and manufacturing employment. Energy costs are a relatively high share of total costs in manufacturing. There is little evidence of an impact on overall employment as would be expected. The estimates suggest roughly 2% to 4% declines in manufacturing employment but neither would be judged statistically significant by standard criteria.

6.4 Generation

A number of previous papers have examined the impact of RPS programs on state renewable generation (see Shrimali, Jenner, Groba, Chan and Indvik (2012) for an excellent overview of the varied findings). In general, they find that program heterogeneity appears to have some impact, while requirement stringency generally does not. Considering individual state responses, however, is likely confounded by spillovers, as most RPS programs allow out-of-state resources within the REC region to comply.[81] To allow for these spillovers, we aggregate our state-level data to the REC region by taking state-level measures of technology-specific generation shares, taking CO2 emissions intensity (in metric tons per MWh) and whether an RPS program was law, and calculating a weighted average at the REC region level where the weight is the MWh of generation in the relevant state by year observation. REC permits can be traded within a REC region and the ten REC regions are shown in Appendix Figure A-35. We then estimate versions of Equation 33, except now an observation is at the region by year level, rather than state by year level.

Table 18 presents estimates for generation sources observed in the EIA data. There is a case for estimating unweighted (Panel A) and weighted (Panel B) versions of Equation 33 here. The case for the unweighted regression is that the data generating process takes place at the REC region level, with substantial cross-state spillovers due to the tradable REC permits. The case for weighting by the number of states in a REC region depends on whether one wants to count more heavily regions that are comprised of more states. Since the analysis of RPS on retail prices takes place at the state level, weighting REC regions by the number of states to recover the effect on the average state provides the most directly comparable results for the impact of RPS on prices, generation, and carbon intensity.

[81] Johnson (2014) find that future RPS levels are associated with current regional capacity additions.

Table 18: Estimates of RPS Impact on Generation and CO_2 Emissions (Trend Break)

	Hydro (1)	Solar (2)	Wind (3)	Other Renewables (4)	Coal (5)	Natural Gas (6)	Petroleum (7)	Nuclear (8)	CO_2 intensity (9)
Panel A: Unweighted									
Effect of RPS 7 years after passage	2.87	-0.08	1.22**	0.53	-1.27	1.26	-1.68	-2.78	-0.029
$(6\beta_3 + \delta_3)$	(2.46)	(0.13)	(0.56)	(0.48)	(3.70)	(3.80)	(1.39)	(3.03)	(0.034)
Effect of RPS 12 years after passage	4.87	0.04	2.84***	0.26	3.37	-2.17	-3.83	-5.62	-0.043
$(11\beta_3 + \delta_3)$	(3.80)	(0.18)	(1.05)	(0.87)	(7.70)	(8.32)	(3.54)	(4.94)	(0.057)
Panel B: Weighted									
Effect of RPS 7 years after passage	9.97	-0.39	0.98	0.86	-6.28	-4.98	-1.42	0.95	-0.101
$(6\beta_3 + \delta_3)$	(6.19)	(0.29)	(1.35)	(0.77)	(4.38)	(7.14)	(2.63)	(4.12)	(0.062)
Effect of RPS 12 years after passage	16.33	-0.23	0.93	1.24	-5.90	-9.89	-4.47	1.14	-0.149
$(11\beta_3 + \delta_3)$	(9.06)	(0.28)	(1.77)	(1.35)	(6.44)	(14.92)	(5.75)	(7.06)	(0.089)
Mean at τ = -1	4.53	0.00	0.07	2.42	48.28	14.60	7.56	21.98	0.641
Region FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other Programs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	260	260	260	260	260	260	260	260	260

Notes: Columns (1) through (8) show estimates from Equation (9), each with a specific generation source as the dependent variable. Column (9) also shows estimates from Equation (9), but uses the CO_2 emissions intensity as the dependent variable. Coefficient estimates are either for states with data 7 years before and 7 years after RPS passage, or for states with data 7 years before and 12 years after RPS passage. Using Equation (9) notation, the effect of RPS 7 years after passage is $6\beta_3 + \delta_3$, and the effect of RPS 12 years after passage is $11\beta_3 + \delta_3$. *Panel A* is a region-level generation-weighted average of the states in the region, unweighted by count of states in each REC-region. *Panel B* additionally includes state count as regression weights. Standard errors are clustered at the REC region-level. Asterisks denote p-value < 0.10 (*), < 0.05 (**), < 0.01 (***).

The table reports separate estimates both 7 and 12 years after RPS passage for generation shares of hydro, solar, wind, other renewables, coal, natural gas, petroleum, and nuclear, as well as CO2 intensity. Although these technology share regressions are noisy, a few interesting findings emerge. First, RPS passage is associated with substantial increases in the wind and hydro shares of generation. The increase in wind generation is consistent with anecdotal evidence about wind playing an important role in RPS compliance. However, we underscore that the wind estimates are statistically significant in some specifications but certainly not all. Second, the estimates suggest that RPS programs displace petroleum as its share declined meaningfully, although again the standard errors preclude definitive conclusions. Third, although RPS programs likely have countervailing effects on natural gas generation - with renewables likely to displace baseload generation but require an increase in the use of peaker plants as backup for their intermittent production - our estimates are too noisy to provide a test with real empirical content.

Column (9) reports on specifications where CO_2 intensity is the dependent variable. Just as with the generation outcomes, the REC-level value of this variable is calculated as the weighted average of state CO_2 intensity, where the weight is the MWh of generation in the relevant state by year. The mean of this variable in the year prior to program passage is 0.64. The Panel A estimates indicate modest declines in CO2 intensity that have associated t-statistics below 1. In Panel B, the estimated emissions intensity declines by about 16% (=.101/.641) seven years after RPS passage and by 23% twelve years after passage. Both of these estimates are close to being statistically significant at the 10% level.[82]

Overall, the table reveals that RPS programs are associated with changes in the generation mix that are admittedly sensitive to specification and often imprecise. The most consistent evidence appears to be that RPS programs led to reductions in the CO_2 intensity of generation, although the imprecision of these estimates also remains a source of concern. In the next subsection, we combine the emissions intensity results with the price effect results to learn about

[82]See Appendix Figure A-41 for event-study figures associated with these four estimates of the impact of RPS programs on CO_2 emissions intensity that illustrate the source of the column (9) estimates.

the costs per metric ton of CO_2 abated.

7 Interpretation

Our estimates suggest that RPS passage has imposed substantial costs on consumers of electricity to date. To make this concrete, we calculate the higher charges that electricity customers paid during the first 7 years after RPS passage in the 29 adopting states. This is calculated as the product of the estimated increase in prices in each post-passage year (from the fitting of Equation 33) and total electricity consumption in the 29 RPS states in the analysis. The other side of the ledger is the reduction in CO_2 emissions in the 29 RPS states. This is calculated as the product of the estimated effect of RPS passage on CO2 intensity and electricity generation separately for each year post-passage. Recall, the estimated reduction in emissions intensity is about 3.5 times larger in Panel B, compared to Panel A, of Table 18, so the results will be sensitive to the decision of whether to weight observations on REC regions.

A natural summary statistic of RPS programs efficacy is the cost per metric ton of CO2 abated and Table 19 uses this papers estimates to develop several of these measures. Specifically, the first row of each panel reports the cumulative effect of RPS programs in their first 7 years after passage, using the estimated impact on electricity prices in Table 14 and the unweighted (Panel A) and weighted (Panel B) regressions for CO_2 intensity from Table 18. Without discounting, the total additional RPS costs over the first 7 years are about $125 billion in the 29 adopting states. The cumulative reduction in CO2 emissions over the first 7 years after passage is 240 million metric tons in Panel A and 1,010 million metric tons in Panel B.

Column (3) reports the cumulative estimated costs per ton of CO2 abated during the first 7 years after passage and they are $530 and $124 in the two panels, with the wide range underscoring the sensitivity of the estimate to the estimated impact of RPS programs on CO_2 intensity. The second and third rows of each panel report on the cost per metric ton of CO_2 abated in the 7th and 12th years after passage. The cost per ton abated increases modestly between years 7 and 12 in both panels.

Overall, the estimates of the cost per metric ton of CO_2 abated are high by almost any metric.

Table 19: Estimated Cost of Abating CO_2 Emissions from RPS

	CO_2 Reduction (mm ton) (1)	Cost to Consumers (bn \$) (2)	Cost per Ton Reduced (\$) (3)
Panel A: Unweighted			
Cumulative Effect of RPS (for first 7 years after passage)	15.7	127.7	8,156
Effect of RPS 7 years after passage ($6\beta_3 + \delta_3$)	40.1	28.7	716
Effect of RPS 12 years after passage ($11\beta_3 + \delta_3$)	47.2	27.6	585
Panel B: Weighted			
Cumulative Effect of RPS (for first 7 years after passage)	896.9	127.7	142
Effect of RPS 7 years after passage ($6\beta_3 + \delta_3$)	229.7	28.7	125
Effect of RPS 12 years after passage ($11\beta_3 + \delta_3$)	204.0	27.6	135
State Count 7 years after passage	29	29	29
State Count 12 years after passage	16	16	16

Notes: Column (1) shows estimates from Equation (9) estimated at the REC level, where Panel A excludes and Panel B includes state-count weights. Column (2) shows estimates from Equation (9) estimated at the state level, so no state-count weights are used in either panel. Column (3) is the ratio of column (2) to (1). The cumulative effect of RPS is the sum of the year-by-year effects for $\tau = 0$ through $\tau = 6$ inclusive.

For example, the Obama Administration pegged the social cost of carbon (i.e., the monetized damages from the release of an additional ton of CO2 in the year 2019) at roughly $51 in current dollars (Greenstone, Kopits and Wolverton (2013); EPA (2016)). Thus, it appears that the costs of RPS programs exceed their carbon reduction benefits (again, these benefits would be larger if these programs reduce the future cost of renewable technologies that end up being deployed). Further, they exceed the price of a permit to emit a ton of CO_2 in all the major cap-and-trade markets globally by more than an order of magnitude. For example, the current prices in the CA, Regional Greenhouse Gas Initiative, European Union ETS, and Quebec markets are currently about \$15, \$6, \$25, and \$15, respectively.[83] Put another way, RPS programs appear to be achieve a small fraction of the CO2 reductions per dollar of cost, relative to cap-and-trade markets.

There are several caveats and implications of these results that bear noting. First, the analysis is reduced form so we cannot assign precise shares of the RPS programs full costs to differences in generation costs, intermittency, transmission, and stranded assets. Further, it seems reasonable to assume that these shares vary over time and in ways that further complicate trying to infer their contributions. For example, it seems plausible that any stranded asset costs are declining at the same time that intermittency costs are increasing, because the net requirements grow over time.

Second, there are two reasons that the cost per metric ton calculations may understate the full social costs of RPS programs. This is because the price effects only reflect the portion borne by ratepayers. However, it seems reasonable to presume that at least some of the costs will be borne by owners of capital (e.g., generators or transmission), particularly in states with restructured electricity markets. Further, it is possible that some of the costs are shared by all the participants in wholesale electricity markets, which in several cases includes states with and without RPS programs. If the costs are partially reflected in retail prices in non-adopting states, then the difference in differences approach would understate the full costs borne by ratepayers because it would miss the portion borne by ratepayers in non-adopting states and understate the effect in adopting states.

Third, more broadly, a randomized control trial is unavailable here, so there will always be a form of unobserved heterogeneity that could explain the results without RPS programs playing a causal role. For example, our measures of other state programs that influence retail electricity prices are limited in their detail, only measuring the years in which states adopted three of these types of programs. So while our estimates are adjusted for the presence of three of these types of programs, this may fail to capture their full impact on electricity prices and that could cause

[83]Because there are mandates inside these cap-and-trade programs, the permit price may not be reflective of marginal abatement costs across the entire covered sectors.

us to understate or overstate the impacts of RPS programs on retail electricity prices, depending on their correlation with RPS programs.

Fourth, it is often claimed that renewable policies provide an external benefit by reducing the costs of future generation that is generic and cannot be fully appropriated by the firm that is expanding its operations. If there are such spillovers or positive externalities, then our estimates of the costs per metric ton of abatement will be systematically too high because they will not account for the benefits received by customers outside of the RPS states jurisdiction. In principle, these benefits could be global and thus quite substantial. The coincidence of the proliferation of policies that support renewable energy and the decline in solar prices over the last decade are consistent with the possibility of such spillovers. However, research that isolates the magnitude of any such spillovers from other factors is probably best described as emerging, making this is a rich area for future research (Gillingham and Stock, 2018).

8 Conclusion

This paper has provided the first comprehensive evaluation on the impacts of RPS programs, which are perhaps the most popular and prevalent carbon policy in the United States. First, these programs mandated increases in renewable generation that are often smaller than is advertised. Seven years after passage, RPS programs require a 1.8 percentage point increase in renewables share of generation, and 12 years after it is 4.2 percentage points. Second, RPS program passage leads to substantial increases in electricity prices that mirror the programs increasing stringency over time. Seven years after passage, we estimate that average retail prices are 1.3 cents per kWh or 11% higher than they otherwise would be. The corresponding effect twelve years later is 2.0 cents per kWh or 17% higher. Third, the estimates indicate that passage of RPS programs lead to reductions in the generating mixs carbon intensity, although these estimates can be noisier and more sensitive to specification than is ideal. Putting the results together, the cost per metric ton of CO2 abated exceeds $115 in all specifications and ranges up to $530, making it at least several times bigger than conventional estimates of the social cost of carbon.

A particularly striking finding is that the indirect costs of RPS programs, which have not been possible to comprehensively measure to date, appear to account for the majority of RPS program costs. A recent study suggests that the direct costs of RPS increase retail electricity prices by 2% (Barbose, 2018)), which is substantially smaller than our estimates that prices are about 11% higher 7 years after passage. Although there are several differences between these two studies, it seems likely that the indirect costs, including intermittency, transmission, and stranded asset payments, account for a substantial fraction of RPS program costs. This finding suggests caution in extrapolating declines in the direct generation cost of renewable energy to its overall impact on electricity prices, and suggests that reducing indirect costs associated with grid integration could represent the more important barrier to substantially increasing renewable energy's share of generation and meaningfully decreasing carbon dioxide emissions.

Overall, the paper's results underscore the importance of research on policy and technology mechanisms to reduce the costs of renewable energy, and imply that mechanisms to facilitate the integration of intermittent sources onto the grid, such as advanced storage technologies or time-of-use pricing, could be especially beneficial. While the potential damages from global climate change have been widely documented, it is almost self-evident that failing to cost-effectively reduce emissions will ultimately limit the magnitude of these cost reductions. Further, policies that substantially increase the price of electricity tend to have a regressive impact that hits low-income consumers hardest, and therefore may be especially unattractive in developing countries that account for a large and growing share of global emissions. The most effective climate policy in technologically advanced and innovative nations such as the United States will reduce emissions domestically, but also involves developing low-carbon energy systems that are cost-effective enough to promote adoption in the rest of the world.

Chapter 3: A Global View of Creative Destruction

with Chang-Tai Hsieh and Pete Klenow

Abstract

In the wake of the U.S.-Canada Free Trade Agreement, both the U.S. and Canada experienced a sustained increase in job reallocation, including firms moving into exporting. The change was concentrated in industries with steeper tariff cuts, and involved big firms as much as small firms. To mimic these patterns, we formulate a model of innovation by both domestic and foreign firms. In the model, trade liberalization quickens the pace of creative destruction, thereby speeding the flow of technology across countries. The resulting dynamic gains from trade are several times larger than the gains in a standard static model.[84]

[84]We thank Beiling Yan of Statistics Canada for her assistance with the Canadian manufacturing data, Eric English and Erxiao Mo for excellent research assistance, and Sam Kortum for very helpful comments. Hsieh acknowledges support from Chicago Booth's Polsky Center and Klenow from the Stanford Institute for Economic Policy Research. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed by U.S. Census Bureau and Statistics Canada to ensure that no confidential information is disclosed.

1 Introduction

Landmark studies by Bernard and Jensen (1999), Eaton and Kortum (2002), Melitz (2003), and others placed heterogeneous firms at the center of research on international trade. The first wave of follow-up research has mostly focused on models in which trade liberalization leads to a burst of job reallocation and growth, but no medium or long run effect on either.

A growing literature seeks to assess the dynamic costs (such as time consuming job reallocation) and the dynamic benefits of trade (such as faster growth). Empirical studies on the cost side include Autor, Dorn and Hanson (2013b) and Dix-Carneiro and Kovak (2017), and on the benefit side include Bloom, Draca and Van Reenen (2016) and Aghion, Bergeaud, Lequien and Melitz (2018). Efforts at modeling the growth effects of trade build on the foundational models of Rivera-Batiz and Romer (1991) and Grossman and Helpman (1993). These include Alvarez, Buera and Lucas (2013), Perla, Tonetti and Waugh (2015), Buera and Oberfield (2017), and Akcigit, Ates and Impullitti (2018).

In this paper we present facts and a model focusing on the role of creative destruction in trade. We document the magnitude of job reallocation tied to exports in U.S. and Canadian manufacturing firms before vs. after the 1988 U.S.-Canada Free Trade Agreement. Industries which saw larger import tariff cuts experienced elevated job reallocation rates for decades after the agreement. Exit and job destruction rates rose for big firms as much as for small firms, a result in line with the findings of Holmes and Stevens (2014) for the U.S. in the wake of the China shock.

In our model, ideas flow across two countries through trade. Innovators – both entrants and incumbents – draw from a Pareto distribution building on the technology of the firm selling in the domestic market. When innovators take over the local market for an existing product (creative destruction), job reallocation takes place. Domestic firms can also take over foreign markets for a product, as can foreign firms the domestic market.

The first version of the model features exogenous arrival rates of innovation as in Garcia-Macia, Hsieh and Klenow (2018). It is a two-economy version of the influential Klette and Kortum (2004) model of creative destruction, only with exogenous arrival rates. Our second version of the model endogenizes the arrival rates. We build in diminishing returns to the stock of ideas *a la* Jones (1995) and Bloom, Jones, Van Reenen and Webb (2018), so that growth is semi-endogenous and linked to growth in the number of researchers. In both models, the two countries grow at the same rate in the long run.

We calibrate the model to fit moments in U.S. manufacturing vs. manufacturing in the rest of the OECD. We match TFP growth, growth in research investment, exports relative to manufacturing shipments, and the share of entrants in total employment. To pin down the Pareto shape parameter we fit the gap in revenue per worker for exporters vs. non-exporters in U.S. manufacturing plants. We also target manufacturing value added per worker and employment of the U.S. vs the rest of the OECD. We infer higher innovation rates in the U.S. than in the rest of the OECD to match the lower GDP per worker in the latter.

Once calibrated, we analyze steady states and transition dynamics in response to tariff changes. In the exogenous arrival rate version of the model, lower tariffs boost the growth rate in both the U.S. and the rest of the OECD. Because the U.S. is more innovative, the rest of the OECD benefits more and its real consumption wage rises relative to the U.S. version. Lower tariffs lead to more job destruction. In the model, there is a spike immediately after tariffs are lowered, but job destruction remains higher in the new steady state.[85]

In the endogenous arrival rate version of the model, lower tariffs boost growth only temporarily. This is because of diminishing returns in idea production. Ideas do spread faster with lower tariffs, so that each country ascends to a higher TFP path than before the liberalization. The rest of the OECD benefits more because they receive more U.S. ideas than they send to the U.S. Welfare gains from trade, in consumption-equivalent terms, are much higher than in a model with no changes in technology: 2.5 times bigger in the U.S., and 5 times bigger in the rest of the OECD.

We also do a calibration to fit U.S. and Canadian manufacturing. We use it to carry out a

[85]Our model features full employment, but the evidence cited above makes clear that job reallocation is not so seamless in the real world.

counterfactual to fit the change in trade flows after the U.S.-Canada Free Trade Agreement. We then compare the model's quantitative predictions to the changes in job reallocation that we observed in the data, and find them broadly similar. The model predicts persistently higher job reallocation rates after trade liberalization, much of it at large and newly exporting firms.

To dissect our dynamic gains from trade, we entertain alternative assumptions about idea flows across countries. When we assume countries learn partially from *domestic producers* rather than from sellers into the domestic market, the gains from trade shrink toward the static gains. Thus idea flows are critical to our large dynamic gains from trade. When we assume countries specialize in innovating on products they produce rather than import, however, the dynamic gains from trade are even larger for the U.S. Given its innovativeness, the U.S. gains a lot from specializing its draws on a subset of products. Due to limited idea flows across countries, the rest of the OECD benefits less from trade when there is research specialization. We did not try disconnecting idea flows from trade entirely, because then trade liberalization would not lead to sustained increases in job reallocation as seen in Canada after its 1988 Free Trade Agreement with the U.S.

Our effort is most closely related to three recent papers. Perla, Tonetti and Waugh (2015) study the impact of trade on exit, entry, domestic technology diffusion, and growth in a model of symmetric countries. Like us, they find large dynamic gains from trade. They derive analytical solutions in a model of many countries, whereas we simulate a two-country model calibrated to evidence on trade and job flows. Our focus is innovation, idea flows across countries, and creative destruction, whereas their focus is on the interaction of trade with domestic technology diffusion.

We follow Buera and Oberfield (2017) in studying international technology diffusion in a model with Bertrand competition. They arrive at Frechet distributions of productivity within countries, allowing them to characterize multilateral trade flows as in Bernard, Eaton, Jensen and Kortum (2003). Our focus is more empirical, as we try to match evidence on job reallocation associated with creative destruction from trade. They stress that the dynamic effects of trade could be negative depending on whether firms learn from domestic producers or from sellers into the domestic market.

Akcigit, Ates and Impullitti (2018) are similar to us in characterizing the impact of tariffs on growth in a two-country model with technology spillovers. Theirs is a step-by-step innovation model, with escape-from-competition effects that are crucial for how trade can induce more innovation. They analyze transition dynamics and optimal R&D subsidies. Their knowledge spillovers take the form of followers catching up to leaders in one big jump if they fall too far behind. They emphasize the convergence of patenting in other advanced countries toward the U.S. in recent decades. In our model and empirics, in contrast, we focus on how trade affects job reallocation.

The rest of the paper is organized as follows. Section 2 lays out nine facts from U.S. and Canadian manufacturing that we attempt to explain. In Section 3 we present a two-country model of creative destruction and growth with exogenous innovation rates. Section 4 endogenizes the innovation rates. In Section 5 we carry out additional exercises (U.S.-Canada trade liberalization, alternative assumptions about idea flows). Section 6 concludes.

2 Facts from Canadian and U.S. Manufacturing

We use data from the U.S. Longitudinal Business Database (LBD) and Canada's Annual Survey of Manufactures. The U.S. LBD is based on administrative employment records of every nonfarm private establishment in the U.S. We have this data every year from 1977 to 2013, and restrict the sample to establishments owned by firms that own at least one manufacturing establishment in the given year. We include the non-manufacturing establishments of such firms to account for the relocation of jobs from establishments classified as manufacturing to establishments of the same firm that are classified as non-manufacturing. The Canadian data covers all but the smallest manufacturing establishments every year from 1973 to 2012.[86] The variables we use from the U.S. LBD and Canadian manufacturing data are the plant and firm identifiers, employment and industry (four digit SIC or six digit NAICS). The Canadian data has information on exports every five years from 1974 to 1989, 1993 and every year from 1996 to 2012. The U.S.

[86]The survey threshold is currently annual sales of 30 thousand Canadian dollars.

LBD does not measure exports but this information is available in the micro-data of the U.S. manufacturing census every five years starting in 1987. We merge the establishments in the manufacturing census with the LBD to measure exports in our LBD sample.[87] We aggregate establishment data in the U.S. and Canada to the firm level. We highlight nine facts:

1. **Large Job Flows**. Table 20 (rows 1 and 2) presents the job creation and destruction rates over five years for Canadian (from 1973 to 2012) and U.S. manufacturing (from 1987 to 2012).[88] As in the classic work by Davis, Haltiwanger and Schuh (1996), job flows are large. The average job creation and destruction rate over five years is about 30% in Canada. The average job creation rate in U.S. manufacturing from 1973 to 2012 is also about 30%. The U.S. job destruction rate is about 5 percentage points higher.

2. **Job destruction due to "large" firms**. Row 3 in Table 20 presents the job destruction rate due to exit or employment declines only among *large* firms. Large is defined as above-average employment in the initial period. Large firms account for 84% of all job destruction in the U.S. and 48% of all job destruction in Canadian manufacturing.

3. **Job creation due to exports.** We impute employment due to exports as the product of a firm's employment and the ratio of its exports to total shipments. Job creation from exports is the sum of imputed employment in year $t + 5$ of new exporters (firms that enter into exporting between year t and $t + 5$) and the change in imputed employment from firms where exports increased between the two years. We divide this measure of job creation from exports by the average of aggregate employment in years t and $t + 5$. The resulting

[87]The LBD and the manufacturing census use the same plant identifiers.

[88]The job creation rate between year t and $t+5$ is defined as the ratio of the sum of employment of new firms (established between the two years) in year $t + 5$ and the change in employment among *expanding* firms between the two years to average total employment (in year t and $t + 5$). The job destruction rate between years t and $t + 5$ is the sum of employment in year t of firms that exit in the next five years and the change in employment between years t and $t + 5$ among *contracting* firms divided by average total employment (in the beginning and ending years). Job flows for the U.S. are calculated for every five year period from 1987 to 2012. Job creation, destruction, and job destruction from large firms for Canada are calculated every five years from 1973 to 2008. For 2008 to 2012, we multiply by 5/4 to impute the flow over five years. Job creation from exports in Canada is calculated from 1974–1979, 1979–1984, 1984–1989, 1989–1993, 1993–1998, 1998–2003, 2003–2008, and 2008–2012, where we multiply the rate from 1989–1993 and 2008–2012 by 5/4 to impute the flow over five years.

Table 20: Job Flows in the U.S. and Canada

	U.S.	Canada
Job Creation Rate	31.4%	32.4%
Job Destruction Rate	36.6%	31.6%
Job Destruction from Large Firms	30.7%	15.3%
Job Creation from Exports	2.0%	23.3%

Note: Job creation and destruction rate calculated over successive five year periods from 1987 to 2012 for the U.S. and 1973 to 2012 for Canada. Jobs from exports imputed as the product of firm employment and ratio of exports to total shipments. "Large" refers to above-mean employment in the initial year of each five year period.

number, in row 4 in Table 20, shows that the job creation rate due to exports is 23% in Canada. The job creation rate due to exports in the U.S. is much smaller at 2%.[89]

We next document how job flows changed after the Canada-U.S. Free Trade Agreement (CUSFTA). This agreement was signed on January 2, 1988, and mandated annual reductions in tariffs and other trade barriers over a ten-year period starting on January 1, 1989.[90] For Canada, we focus on the difference between 1973–1988 ("Pre-CUSFTA") and 1988–2003 ("Post-CUSFTA"). We highlight three facts in the Canadian data:

4. **Job flows increased after trade liberalization**. Table 21 shows that job creation and destruction rates increased in Canada after the trade agreement with the U.S. Figure 22 plots the change in the job destruction rate in two digit Canadian industries from 1973–1988 to 1988–2003 against the change in the tariff rate on Canadian imports in those same industries over the same period due to CUSFTA.[91] Job destruction rates increased more

[89]Lincoln, McCallum and Siemer (2017) estimate that 29% of U.S. exports in 2002 were by firms that had been exporting for fewer than 5 years.

[90]The average tariff on manufacturing imports among the CUSFTA partners fell from over 8% to below 2% in Canada and from 4% to below 1% in the U.S.

[91]We use the tariff cuts constructed by Trefler (2004), which give changes in bilateral tariffs between Canada and the U.S. following CUSFTA net of the changes in the respective most-favored-nation tariffs.

Table 21: Job Flows in Canada

	Pre-CUSFTA	Post-CUSFTA
Job Creation Rate	30.0%	31.3%
Job Destruction Rate	25.5%	32.7%
Job Destruction from Large Firms	22.1%	24.0%
Job Creation from Exports	9.0%	17.7%

Note: Pre-CUSFTA is 1978 to 1988. Post-CUSFTA is 1988 to 2003. Job creation and destruction calculated over five year periods.

in industries where tariffs declined the most, thought the relationship is not particularly tight.[92]

5. **Large firms increased job destruction after trade liberalization.** Holmes and Stevens (2014) show that large U.S. manufacturing firms were adversely affected by the surge in imports from China. Row 3 in Table 21 documents a similar fact in Canada. The job destruction rate among large (above-mean employment) firms increased by 2 percentage points after CUSFTA, out of an overall increase in job destruction of 7 percentage points.

6. **Job creation from exports increased after trade liberalization.** The last row in Table 21 shows that job creation from exports increased by almost 9 percentage points in Canada after the trade agreement. Figure 23 shows that a similar fact holds across two-digit Canadian industries. Job creation from exports increased more in sectors where U.S. tariffs declined the most, though again this relationship is not a tight one.[93]

Table 22 presents the change in job flows in U.S. manufacturing after CUSFTA. The timing of the U.S. data does not align as well with the trade agreement so here we focus on the 1972–1987 as the "pre-CUSFTA" period and 1992–2012 as the "post-CUSFTA" period. As documented by a

[92]The coefficient of the OLS regression in Figure 22 is $-.096$ with a standard error of $.031$.
[93]The coefficient of the OLS regression in Figure 23 is $-.189$ with a standard error of $.096$.

Figure 22: Δ Job Destruction in Canada vs. Δ Canadian Tariffs

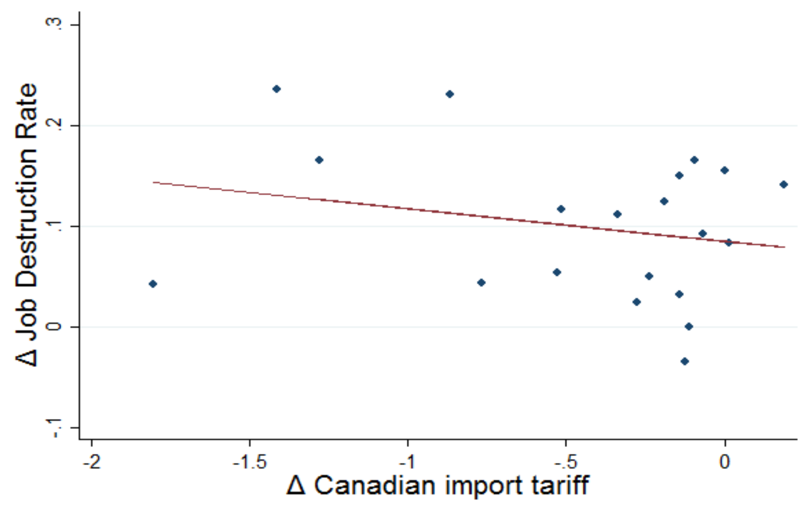

Note: Each observation is a two digit Canadian industry. Δ job destruction is the difference between the average job destruction rate (calculated over five years) in 1988 to 2003 and 1978 to 1988.

Figure 23: Δ Job Creation in Canada from Exports vs. Δ U.S. Tariffs

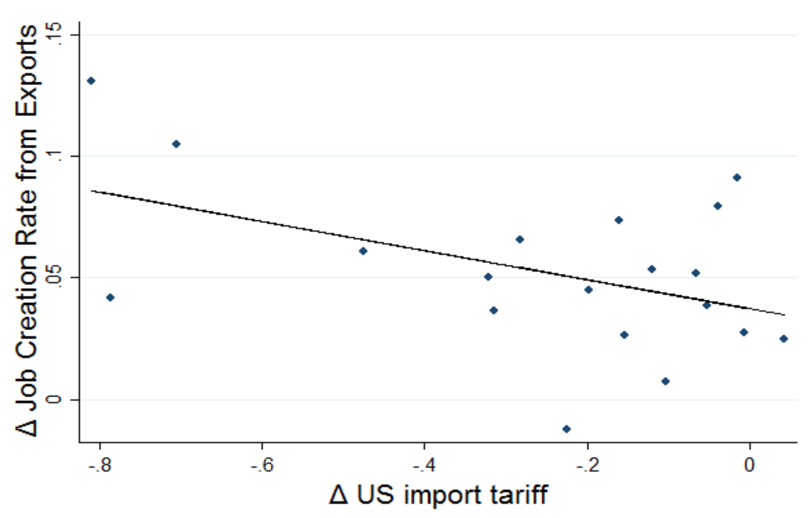

Note: Each observation is a two digit Canadian industry. Δ job creation from exports is the difference between the average job creation rate from exports from 1988 to 2003 and 1978 to 1988, both calculated every five years.

Table 22: Job Flows in the U.S.

	1977–1987	1987–1992	1992–2012
Job Creation Rate	33.7%	33.0%	31.1%
Job Destruction Rate	31.6%	31.0%	37.8%
Job Destruction from Large Firms	25.8%	24.9%	32.1%
Job Creation from Exports	–	2.0%	2.0%

Note: Calculated from U.S. manufacturing census micro-data. Job creation and destruction calculated over five year periods. "Large" firms are above average employment firms in the initial year.

large literature, there was also a surge of imports from China in the 1992–2012 period, so one should not interpret the changes in Table 22 as coming only from CUSFTA.

Three facts are clear from Table 22. First, job destruction increased markedly after 1987, by about 6 percentage points (row 2).[94] Second, the increase in job destruction was entirely driven by large firms. Third, there is no U.S. export data prior to 1987, but job creation from exports was stable at 2% between 1987–1992 and 1992–2012. If we only consider manufacturing establishments, job creation from exports increased modestly from 2.7% to 3.1%.

We now look at differences between exporting and non-exporting firms. Figure 24 plots the distribution of employment (in the left panel) and labor productivity (revenue per worker, in the right panel) from the U.S. manufacturing census in 2012. This figure reveals two additional facts:

7. **Average labor productivity and employment is higher for exporters than for non-exporters.** This can easily be seen in Figure 24.

8. **Overlap of labor productivity and employment between exporters and non-exporters.** Figure 24 makes clear that the distributions overlap substantially. Many exporting firms are smaller than non-exporters, and many exporters have lower labor productivity than

[94] This may seem surprising given the evidence on declining dynamism in Decker, Haltiwanger, Jarmin and Miranda (2014). For U.S. manufacturing firms, at least, this decline was concentrated in job creation and took place well after CUSFTA.

Figure 24: Distribution of Employment and Labor Productivity

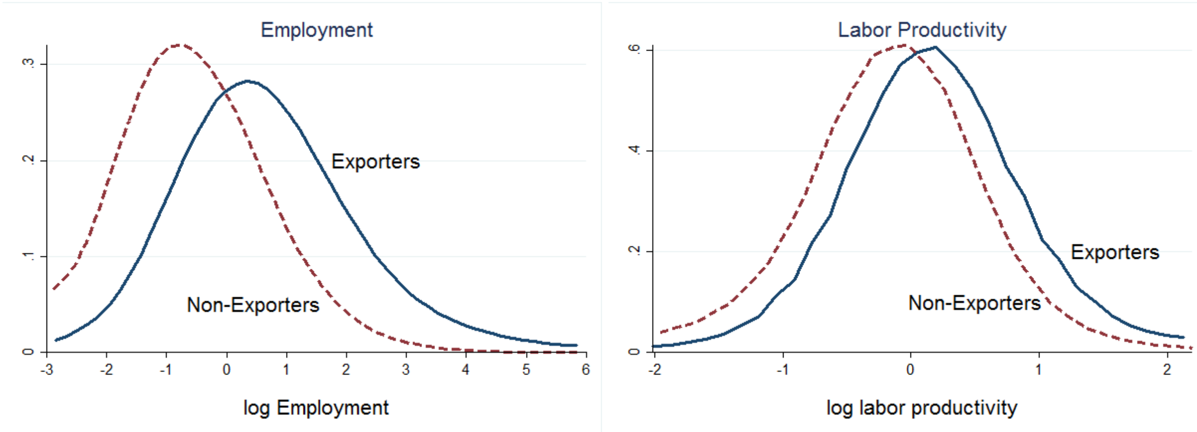

Note: The figures plot the distribution of labor productivity (value-added per worker) and employment of exporting and non-exporting firms in the 2012 U.S. Census of Manufacturing.

many non-exporters.

Our last fact is the rapid and continuous turnover in a country's export products highlighted by Hanson, Lind and Muendler (2018).

9. **Churning of export products.** Table 23 presents two measures of churn of an exported product.[95] The first two rows show the probability that a product exported by the US (column 1) or other OECD countries (column 2) in a given year is no longer exported by the country the following year. The first row shows this statistic for all exported products and the second row for products in the bottom 50% of the export sales distribution (based on sales per product). About 8% of all products exported by the U.S. in a given year is no longer exported by a U.S. firm the following year. And about 15% of the bottom 50% of U.S. exports in year t is not exported by the US in year $t+1$.[96] Our second measure of export churn is based on the country's *top* exported product in each year. Row 3 shows the share of the top export in total exports in year t, and Row 4 shows the ratio of export share of the

[95] A product in Table 23 is one of the 540 4-digit SITC (revision 2) manufacturing industries in Feenstra et al. (2005)'s World Trade Database.
[96] Rows 1 and 2 in Table 23 are the average of one-year panels from 1982-1983 to 2002-2003.

Table 23: Export Product Churn

	U.S.	Rest of OECD
Annual Exit Rate		
All Exported Products	7.8%	8.4%
Bottom 50% in Export Sales	15.1%	15.7%
Top Export Product		
Share of total exports year t	5.7%	12.8%
Share of total exports year $t-5$ / year t	66.5%	86.2%

Note: A product is one of 540 4-digit manufacturing industries in Feenstra, Lipsey, Deng, Ma and Mo (2005)'s World Trade Flows. Rows 1 and 2 show the probability an exported product in year t is no longer exported by the country in year t+1 for all exported products and products in bottom half of export sales, respectively. Row 3 shows the share of the top exported product in total exports in year 4. Row 4 shows the ratio of the share in total exports of the top export five years before to the export share of the top product in a given year. Entries are average of one year panels from 1982–1983 to 2002-2003 (rows 1-2) or five year panels from 1982–1987 to 2002–2007 (rows 3-4).

same product in year $t-5$ relative to the share in year t. This ratio averages 66.5% for the US and 86.2% for the other OECD countries.[97]

3 Exogenous Innovation

This section presents a model of growth driven by creative destruction, where innovation can come from domestic or foreign firms. The goal is to examine the dynamic gains from trade liberalization, and to see whether this model can mimic the nine facts described in section 2.

3.1 Static Equilibrium

Utility of the home-country representative consumer is given by consumption of a continuum of varieties C_j with measure 1:

$$U = \int_0^1 \ln C_j \, dj. \tag{34}$$

[97] These numbers are the average of five-year panels from 1982–1987 to 2002–2007.

Table 24: Markups

	Traded Produced in Home	Non-Traded	Traded Produced in Foreign
Home	$\frac{A_i}{\max\left[A_i', \frac{\omega A_i^*}{\tau}\right]}$		$\frac{A_i^*/\tau}{\max\left[\frac{A_i^{*\prime}}{\tau}, \frac{A_i}{\omega}\right]}$
Foreign	$\frac{A_i/\tau}{\max\left[\frac{A_i'}{\tau}, \omega A_i^*\right]}$		$\frac{A_i^*}{\max\left[A_i^{*\prime}, \frac{A_i}{\omega\tau}\right]}$

This utility function implies that consumers spend the same share of their income on each variety.[98]

Output of each variety is the product of labor and the quality of the blueprint for the product. We assume that there is a choice of blueprints for each product. We denote A_j as the "best" blueprint among domestic firms. A_j^* is the corresponding best blueprint among foreign firms. Furthermore, suppose the product index j is decreasing in A_j/A_j^*. Then products $j \in [0, x]$ are traded and produced at home, $j \in [x, x^*]$ are non-traded, and $j \in [x^*, 1]$ are traded and produced abroad. The cutoff products x and x^* are defined by

$$\frac{A_x}{\tau} = \omega A_x^* \tag{35}$$

$$A_{x^*} = \frac{\omega A_{x^*}^*}{\tau} \tag{36}$$

where ω denotes the relative wage (domestic relative to foreign) and $\tau \geq 1$ is the symmetric gross trade cost. When $\tau = 1$, $x = x^*$ and all products are traded.

The owner of the best blueprint sets their quality-adjusted price to push their closest competitor out of the market (Bertrand competition), so the gross markup is the gap between the incumbent firm's marginal cost and the cost of its closest competitor — domestic or foreign. Table 24 summarizes the markup of domestic firms μ_i and foreign firms μ_i^*. Here A_i' and $A_i^{*\prime}$ denote

[98]We suppress the equations for the foreign country when they are the same as that of the home country. For example, utility of the foreign consumer is given by $U^* = \int_0^1 \ln C_j^* \, dj$.

the productivity of the second best producer in the domestic and foreign markets, respectively. These potential competitors do not produce in equilibrium but affect markups.

The relative wage is pinned down by balanced trade:

$$I^* \cdot x = I \cdot (1 - x^*) \tag{37}$$

where I and I^* denote nominal GDP at home and abroad, respectively. The left hand side of equation (37) is the home country's exports and the right hand side is the home country's imports. Nominal GDP in each country is given by

$$I = \frac{\overline{\mu}\, wL}{1 - \frac{1-\tau}{\tau} \cdot (1 - x^*)} \qquad \text{and} \qquad I^* = \frac{\overline{\mu}^* w^* L^*}{1 - \frac{1-\tau}{\tau} \cdot x}$$

where $\overline{\mu}^*$ and $\overline{\mu}$ denote the average gross markup of foreign and domestic firms, w and w^* are the home and foreign wage, and L and L^* are labor supply at home and abroad.[99] More exactly, the average price-cost markup in the U.S. satisfies

$$\frac{1}{\overline{\mu}} \equiv \frac{\int_0^{x^*} \frac{1}{\mu_j} dj + \frac{1}{\tau} \cdot \int_0^{x} \frac{1}{\mu_j^f} dj}{x^* + x/\tau}$$

where μ_j^f denotes the markup of domestic firms on their exported products. The expression for the foreign firms' average markup is analogous.

We can now express the real (consumption) wage as a function of the distribution of the best blueprints, markups, the cutoff indexes, the relative wage, and the trade cost. The real wage at home W and in the foreign country W^* are given by

$$W = \int_0^{x^*} \ln\left(\frac{A_j}{\mu_j}\right) dj + \int_{x^*}^{1} \ln\left(\frac{A_j^*}{\mu_j^*} \cdot \frac{\omega}{\tau}\right) dj \tag{38}$$

$$W^* = \int_0^{x} \ln\left(\frac{A_j}{\mu_j} \cdot \frac{1}{\omega\,\tau}\right) dj + \int_{x}^{1} \ln\left(\frac{A_j^*}{\mu_j^*}\right) dj. \tag{39}$$

[99] The expression for nominal income comes from equating nominal income to the revenue of local firms plus tariff revenue: $I = \overline{\mu}\, wL + (\tau - 1)\frac{I}{\tau}(1 - x^*)$ and $I^* = \overline{\mu}^* w^* L^* + (\tau - 1)\frac{I^*}{\tau} \cdot x$.

Remember the home country buys products $j \in [x^*, 1]$ from the foreign country so the real domestic real wage is increasing in the productivity of foreign firms of these products. Likewise, the foreign country purchases products $j \in [0, x]$ from the home country so the foreign real wage depends on the productivity of domestic firms of these products.

Our static model is simply the Bernard, Eaton, Jensen and Kortum (2003) model, which is essentially the Dornbusch, Fischer and Samuelson (1977) model with heterogeneity in markups. Here equations (35), (36), (37) and Table 24 jointly determine the real wage in the two countries, the relative wage ω, the cutoff products x_1 and x_2, and the markup for each product.

3.2 Innovation

We now introduce dynamics to the model. First, following Klette and Kortum (2004) we assume a firm is a portfolio of products, where a firm produces a product if it owns the best technology for that product. Second, we assume all growth comes from creative destruction – i.e., when a firm (another incumbent or an entrant) improves a product's technology and steals it from the incumbent producer. Third, we assume creative destruction can also come from a firm located in another country. Our goal is to show that the trade friction τ is a key parameter that determines whether innovation in one country results in creative destruction in the other country. We make the following assumptions about innovation. First, we assume a constant exogenous arrival rate for each type of innovation. (We will endogenize the arrival rate in the next section.) Second, we assume that arrivals are in proportion to the number of products owned by a firm. For example, a firm with two products is twice as likely to creatively destroy another firm's variety compared to a firm with one product. Third, we assume that innovation builds on the existing quality level of the product ***consumed domestically***.[100] Specifically, the quality drawn by an innovation follows a Pareto distribution with shape parameter θ and scale parameter equal to the existing quality level. The average percent improvement in quality of an

[100] If innovation was endogenous, there would be a positive externality to research unless all research was done by firms on their own products. Such knowledge externalities are routinely assumed in the quality ladder literature, such as Grossman and Helpman (1991), Aghion and Howitt (1992), Kortum (1997), and Acemoglu, Akcigit, Alp, Bloom and Kerr (2018). See Coe et al. (1997, 2009) for evidence consistent with learning by importing.

Table 25: Channels of Innovation

	Domestic Firm	Foreign Firm
Innovation by incumbents	λ	λ^*
Innovation by entrants	η	η^*

Note: The average improvement in quality is $\frac{1}{\theta-1}$.

existing variety (conditional on innovation) is thus $\frac{1}{\theta-1} > 0$.[101] Finally, we assume that entrants have one product, while incumbent firms potentially produce many varieties.

The notation for innovation probabilities is given in Table 25. The probability a product is improved upon by an incumbent domestic firm is λ. Conditional on not being improved by a domestic incumbent, η is the probability the product is improved by a new domestic firm. Conditional on not being improved by *any* domestic firm, λ^* is the probability the product will be improved by a foreign incumbent firm. Finally, conditional on the product not being improved upon by either a domestic firm or by a foreign incumbent, η^* is the probability a foreign entrant innovates on the best blueprint.

In short, a given product can be improved upon by a domestic incumbent firm, a domestic entrant, a foreign incumbent firm, or a foreign entrant. The probability a product will be improved upon by a domestic incumbent is λ. The *unconditional* probability of innovation by domestic entrant is $\tilde{\eta} \equiv \eta(1-\lambda)$. The *unconditional* probability the product will be improved by foreign incumbent is $\tilde{\lambda^*} \equiv \lambda^*(1-\lambda)(1-\tilde{\eta})$. The unconditional probability of innovation by a foreign entrant is $\tilde{\eta}^* \equiv \eta^*(1-\lambda)(1-\tilde{\eta})(1-\tilde{\lambda}^*)$. Finally, $\mu \equiv \lambda + \tilde{\eta}$ denote the unconditional probability a domestic firm (entrant or incumbent) improves a product, and $\mu^* \equiv \lambda^* + \tilde{\eta}^*$ is the unconditional innovation rate of a foreign firm.

Table 26 summarizes the probability of creative destruction in the domestic market (rows 1-

[101]In the simulations we run in the next section, we add a reflecting barrier whereby the bottom ψ percent of products by quality redraw from the top $1-\psi$ percent of products. This will maintain a stationary distribution of qualities and allow us to match empirical estimates of the elasticity of trade flows with respect to trade barriers.

Table 26: Probability of Creative Destruction

Market	Product Type	Domestic Firm	Foreign Firm
Home	Exported	μ	$\mu^* \cdot \min\left[(\frac{\omega}{\tau})^\theta, 1\right]$
	Non-Traded	μ	$\mu^* \cdot \min\left[\left(\frac{\omega A_j^*}{\tau A_j}\right)^\theta, 1\right]$
	Imported	$\mu \cdot \min\left[\left(\frac{\tau}{\omega}\right)^\theta, 1\right]$	μ^*
Foreign	Exported	μ	$\mu^* \cdot \min\left[(\omega\tau)^\theta, 1\right]$
	Non-Traded	$\mu \cdot \min\left[\left(\frac{A_j}{\omega\tau A_j^*}\right)^\theta, 1\right]$	$\mu*$
	Imported	$\mu \cdot \min\left[(\frac{1}{\omega\tau})^\theta, 1\right]$	μ^*

3) and the foreign market (rows 4-6) due to innovation by domestic firms (column 1) and foreign firms (column 2). The first row shows the arrival rate of new ideas in the domestic market for a product that is also exported to the foreign market. The probability this product is improved upon by another domestic firm is μ, and a domestic innovator will always replace the incumbent firm in this market. A foreign firm also improves upon the same product with probability μ^* but a successful foreign innovator does not necessarily replace the domestic incumbent. Since quality improvement follows a Pareto distribution, the probability that the quality improvement of the foreign innovator is large enough to replace the domestic incumbent is min $\left[\left(\frac{\omega}{\tau}\right)^\theta, 1\right]$. For a given innovation rate by foreign firms, higher relative wages ω and lower trade costs τ increases the probability that innovation by a foreign firm benefits domestic consumers. Intuitively, higher domestic wages increase the probability a foreign firm that innovates will be competitive enough to replace the incumbent in the domestic market. Higher trade costs make the foreign innovator less competitive compared to the domestic incumbent. Effectively, trade costs insulate domestic firms from foreign competition in the domestic market.

The expected growth rate of the real consumption wage in the domestic market is the product of the weighted average of the rate of creative destruction in rows 1-3 in Table 26 and the weighted average of the increase in product quality (conditional on the product being replaced). Likewise, the corresponding growth rate of the foreign real consumption wage is the product of the weighted average of the arrival rates in rows 4-6 in Table 26 and the corresponding step size improvement in quality. Real growth rates in the two countries depend on the arrival rates of innovation μ and μ^*, the relative wage ω and the share of each type of product (x and x^*). As discussed in the previous section, the relative wage and the share of products made by each country are pinned down by balanced trade and the distribution of relative technologies $\frac{A_i}{A_i^*}$. And in the model with innovation, the distribution of $\frac{A_i}{A_i^*}$ changes endogenously when foreign and domestic firms innovate on the best blueprints.

To understand the equilibrium in the model with innovation, it is useful to consider the case of completely free trade ($\tau = 1$). In this case, all products are traded so the relevant arrival rates in Table 26 are rows 1 and 3 (for the domestic market) and 4 and 6 (for the foreign market). The probability a domestic firm creatively destroys another firm is thus given by:

$$\underline{\text{Domestic}}\text{ creative destruction rate} = \mu \cdot x^* + \mu \min\left[\omega^{-\theta}, 1\right] \cdot (1 - x^*)\,.$$

The corresponding rate of creative destruction by a foreign firm is:

$$\underline{\text{Foreign}}\text{ creative destruction rate} = \mu^* \cdot (1 - x^*) + \mu^* \min\left[\omega^{\theta}, 1\right] \cdot x^*.$$

Ceteris paribus, higher ω lowers the rate of creative destruction of domestic firms and raises that of foreign firms. In the steady state of the model with innovation, the equilibrium relative wage equates the rate of creative destruction by domestic firms to that of foreign firms. So if domestic firms are more innovative, domestic wages are higher but the rate of creative destruction of domestic firms is the same as that of less innovative foreign firms. The steady state in the dynamic model is thus pinned down by the condition that the rate of creative destruction is the same in the two countries and the conditions that determine the static equilibrium discussed in the previous section (balanced trade and the definition of the cutoff products).

Figure 25: Adjustment to Increase in U.S. Innovation Rate

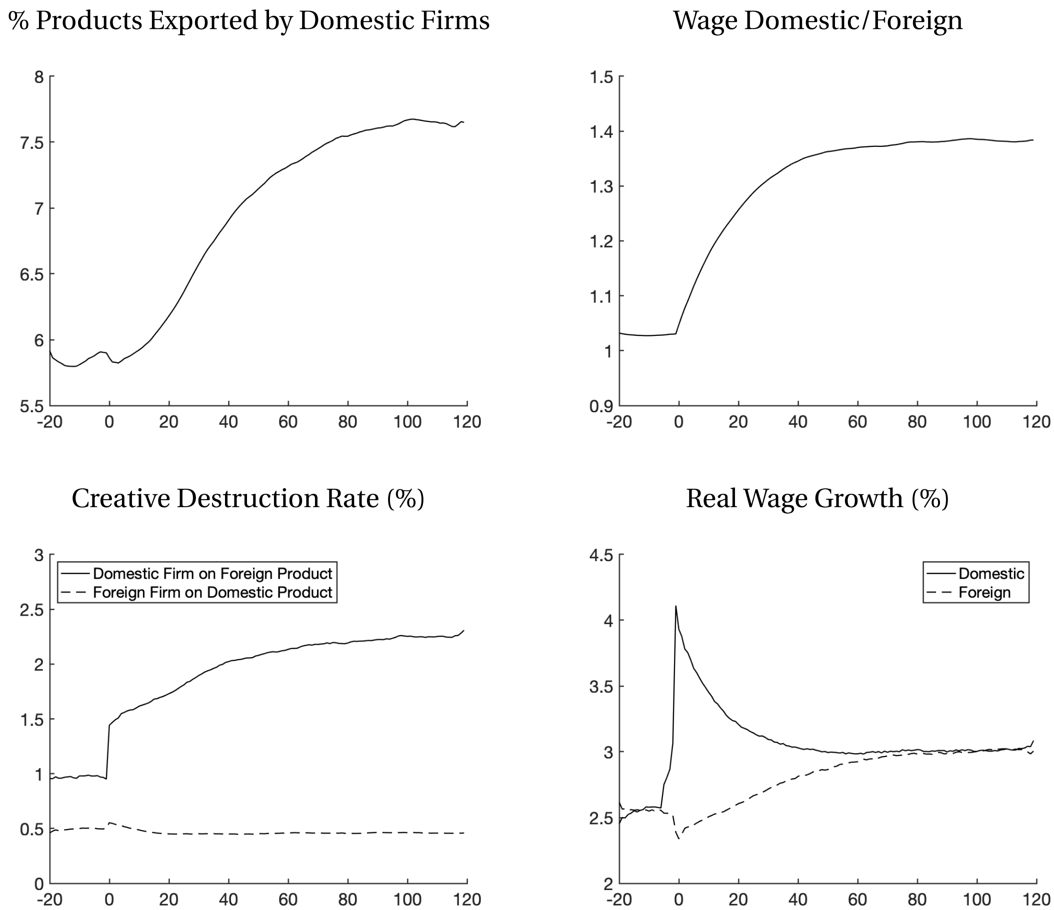

Note: The figure simulates the effect of an increase in λ from 6.75% to 13.45% in year 0 and holds all other parameters constant. The values of the other parameters are $\widetilde{\eta} = 2.55\%$, $\widetilde{\lambda}^* = 9.64\%$, $\widetilde{\eta}^* = 2.55\%$, $\theta = 10.94$, and $\tau = 1.491$. We explain in the next subsection where these numbers come from, although for the purposes of illustrating the effect of a change in λ any set of parameter values will do.

Figure 25 illustrates the determinants of the steady state equilibrium by simulating the adjustment to an increase in the *domestic* innovation rate.[102] Higher domestic innovation increases the share of products exported by domestic firms and the domestic wage relative to the foreign wage. This is illustrated in the two panels in the first row in Figure 25. The second row in Figure 25 shows that, after the increase in the domestic innovation rate at time 0, the rate of creative destruction rate rises for domestic firms. The share of products exported by domestic firms rises and ω rises as long as the creative destruction rate by domestic firms exceeds that of foreign firms. In equilibrium, the relative wage increases by enough such that the rates of creative destruction are the same in the two countries. When this is the case, the relative wage and share of products made by each country is constant. The growth rate of the real wage is higher in the new steady state and the same in the two countries. In sum, in a steady state, countries with higher innovation rates have higher relative wages but the growth rate of real wage is the same rate in all countries.[103]

We close this subsection by contrasting autarky and free trade when the two countries are symmetric in size and in their innovation arrival rates. In this special case the relative wage $\omega = 1$ and the expressions become simply:

$$\underline{\text{Autarky}}\text{ growth rate} = (\lambda + \widetilde{\eta})\,\frac{1}{\theta - 1}$$

$$\underline{\text{Frictionless}}\text{ growth rate} = \left(\lambda + \widetilde{\eta} + \widetilde{\lambda}^{*} + \widetilde{\eta}^{*}\right)\frac{1}{\theta - 1}.$$

Under autarky, each country benefits only from its domestic arrivals. Under frictionless grade, however, each country benefits fully from both domestic and foreign arrival rates. This underscores the scale effect at the heart of the dynamic gains from trade in this model.

[102]The simulation assumes the economy was in a steady state with a constant relative wage prior to year 0. See notes to Figure 25 for the parameters used for the figure.

[103]This result is reminiscent of Acemoglu and Ventura (2002).

Table 27: Data Moments used for Calibration

Data Moment	Source	Value
Revenue per worker exp./non-exp.	U.S. mfg	1.066
TFP growth rate	U.S. mfg	3.01%
Value added per worker home/foreign	U.S. and OECD mfg	1.29
Employment share of entrants	U.S. mfg	16.9%
Export share of revenues (home)	U.S. mfg	10.2%
Employment home/foreign	U.S. and OECD mfg	0.389
Employment growth rate	U.S. mfg	–1.1%
Trade elasticity from halving τ	Head and Mayer (2014)	–5

3.3 Calibration

The model is summarized by two innovation rates (for incumbents and entrants) in each country, the shape parameter of the Pareto distribution of the innovation draws, and a trade cost. In this section, we infer the value of these parameters from simple moments of the data.

For consistency with the model, we assume the world consists of the manufacturing sectors in the U.S. and the rest of the OECD ("foreign"). The shape parameter of the Pareto distribution of innovation draws (θ), relative employment (L/L^*), innovation rates in Table 25, and the trade cost (τ) jointly determine the growth rate, the trade share and the relative wage. For a given value of θ, we can back out the trade cost and innovation rates from data on total employment, the growth rate, the trade share, and the relative wage.[104] We use the employment share of new firms to pin down the innovation by entrants vs. incumbents.

The data moments we use are displayed in Table 27. We fit a steady state where TFP in the two countries grows at 3% per year, employment shrinks at 1.3% per year, output per worker in the home country is 29% higher than in the foreign country, and the home trade share is 10%.[105] The

[104] We describe in the next section how we back out θ from the gap in labor productivity (revenue per worker) between exporters and non-exporters.

[105] Overall OECD employment was not contracting during our sample, but employment in manufactur-

Table 28: Estimates of Model Parameters

Variable	Description	Value
θ	Shape parameter of innovation draws	10.94
λ	Home innovation rate incumbent	13.45%
η	Home innovation rate entrants	2.95%
μ^*	Foreign innovation rate incumbents + entrants	14.81%
τ	Trade cost	1.491
ψ	Reflecting barrier for product quality	0.011

innovation rates and trade cost needed to fit these facts are shown in Table 28.[106] The innovation rates in Table 28 are *conditional*. The *unconditional* innovation rate is $\lambda + \tilde{\eta} = 0.16$ for domestic firms and $\tilde{\lambda}^* + \tilde{\eta}^* = 0.122$ for foreign firms. The innovation rate of domestic firms has to be higher than that of foreign firms to explain the 29% higher real wage in the home country than in the foreign country. Conditional on the innovation rates and the relative size of the two economies, the trade share pins down the trade cost, which is roughly equivalent to a 50% tariff.[107]

3.4 Firm Dynamics

We now show that the model can generate the nine facts described in Section 2. The driver of growth here is the creation of better quality products and the resulting destruction of product lines of incumbent producers. In our open economy setting, the destruction of product lines can come from firms located in other countries. This has four implications:

First, growth is associated with job creation by innovating firms and job destruction by firms

ing was. We target contracting employment to fit net job creation in the U.S. Negative employment growth poses no problem for obtaining steady state TFP growth and relative wages with exogenous innovation. But we will revisit this for endogenous innovation.

[106]We simulate the model for 5000 varieties in each country. Each variety receives an innovation draw and is assigned to an existing incumbent or a new entrant with random probability governed by the innovation rates. The relative wage is selected to balance trade between the two countries. We simulate this process for several hundred years until it settles down to a steady-state, at which point we calculate the firm-level and aggregate moments produced by the model. We utilize a simulated annealing procedure to search for the parameter values that match the moments in the data.

[107]Eaton and Kortum (2002) and others infer high trade costs to explain bilateral trade flows.

whose products are innovated upon. Table 29 (rows 1 and 2, column 2) shows the job creation and destruction rates in the steady state of the model parameterized to fit the moments in Table 27. The job creation rate (over five years) is 32%, and the job destruction rate is 6% higher at about 38%. For comparison, the first column in Table 29 replicates the U.S. data. The job flows predicted by the model with the parameters in Table 28 are roughly of the same magnitude as in the data (fact 1).

Second, firms shrink in this model when their products are replaced by other firms. A firm exits (and all its jobs are destroyed) when it loses *all* of its products. But job destruction is not only due to exit by small firms; jobs are destroyed at large (multiproduct) firms when a subset of their products are innovated upon. The third row in Table 29 shows that, consistent with the evidence from U.S. and Canadian manufacturing (fact 2), job destruction in the model is mostly driven by large (above-mean employment) firms. The job destruction rate by large firms is 22% in the model, two-thirds of the overall job destruction rate.

Third, some of the job destruction is the result of innovation from firms located in other countries, and some of the job creation is the result of domestic firms replacing producers in the foreign market. The fourth row in Table 29 shows that the job creation rate from exports in the model is 5.7%. In the U.S. data, the number is 2%. We cannot measure empirically when jobs are destroyed because domestic producers are replaced by imports, but we can calculate this moment in the model. The last row shows that the job destruction rate due to imports is 7.4%, so about a fifth of the overall job destruction in the model comes from creative destruction by foreign firms.

Fourth, another implication of the fact that foreign firms replace domestic firms is some of the country's exported products are replaced by a foreign producer. If an exported product is replaced by a another domestic firm, the country will continue to export the given product. In our baseline simulation, the group of products with the highest level of exports tends to account for a lower share of total exports when re-examined five years later. The top exported product's share of total exports is 77.5% as high five years later in the simulation, compared with 66.5% in the data.

Table 29: Firm Dynamics, Data vs. Simulations

Moment	U.S. Data	Simulations
Job Creation Rate	31.4%	31.6%
Job Destruction Rate	36.6%	38.2%
Job Destruction from Large Firms	30.7%	22.3%
Job Creation from Exports	2.0%	5.7%
Job Destruction from Imports	–	7.4%
Top Export Product Turnover	66.5%	77.5%
Probability of Losing an Export	7.8%	12.8%

The U.S. data is the average from 1987 to 2012. The second column show the simulated moments in the steady state of the model with the parameter values from Table 28.

Table 29 shows that the model can replicate, at least qualitatively, facts 1-3 and 9 that we presented in Section 2 (large job flows, lots of job destruction at large firms, some job creation due to exports, and high rates of export turnover). The model can also speak to facts 4-6 (job flows, job destruction at large firms, and job creation due to exports all increase after trade liberalization). We simulate the new steady state of the model with lower trade costs, holding constant all the other parameters. A key assumption in this exercise is that domestic and foreign innovation rates do not change when trade costs change.[108]

We first show the consequence of reducing trade costs for TFP growth in the two countries. Lower trade costs make it more likely that innovation by a domestic firm on a foreign product will replace the foreign producer in the foreign market, and thus raise the quality foreign consumers can consume. Similarly, lower trade costs increase the probability that a foreign firm who innovates on a domestically-produced variety will replace the domestic producer and thus improve the quality that domestic consumers enjoy. The effect of lower trade costs, shown in Figure 26 (left panel), is to boost the growth rate of TFP in the two countries in tandem. Holding

[108]We relax this assumption in the next section, when we endogenize innovation rates.

Figure 26: Simulated TFP Growth Rate and Relative Wage vs. Trade Costs

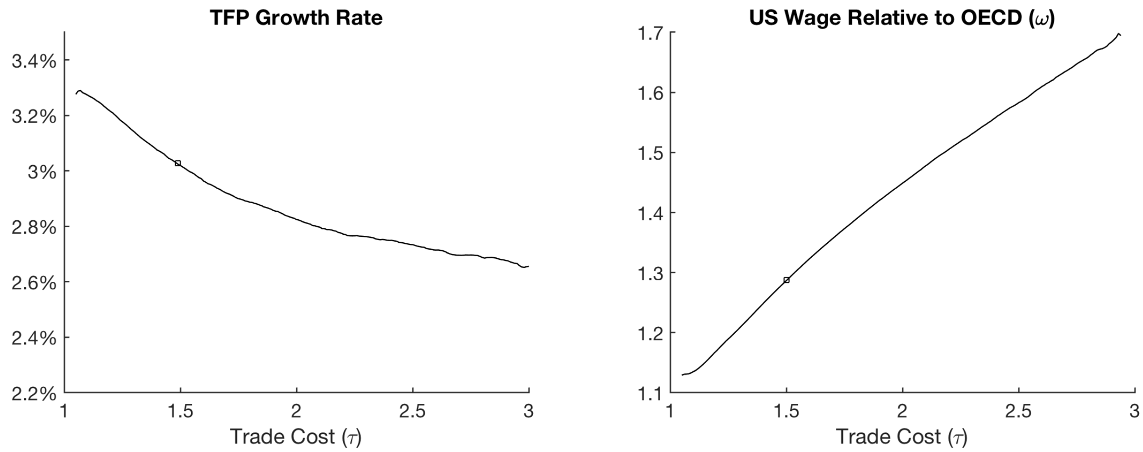

Note: The left panel shows simulated growth rate of the real wage and the right panel the relative wage (home/foreign) in the steady state for different values of τ. All other parameters of model are kept fixed at the values in Table 28.

constant innovation rates, moving to a world with frictionless trade increases the steady-state TFP growth rate from 3% to 3.3%.

Lower trade costs also raise the foreign wage relative to the domestic wage. This is shown in the right panel in Figure 26. The intuition is that a country that innovates less benefits more from trade liberalization since it is now easier for the country to "import" ideas. The home/foreign wage is 1.15 with frictionless trade, versus 1.29 in the baseline with a roughly 50% tariff rate.

Figure 27 shows the effect of the higher TFP growth rate on job creation and destruction. The model predicts that, relative to our baseline steady state ($\tau = 1.491$), free trade would increase the job creation and destruction rates by about 10 percentage points. So a 0.3 percentage point increase in the growth rate is associated with a 10 percentage point increase in the job creation and destruction rate. The impact on job reallocation is more proportionally bigger than the impact on growth, because lower trade costs facilitate creative destruction from trade with smaller step sizes. Consistent with the evidence from Canada and the U.S. (fact 4), the model predicts that job flows rise when trade costs fall.

Figure 28 shows the counterfactual job destruction rate from large firms (those with above-

Figure 27: Simulated Job Creation and Destruction vs. Trade Costs

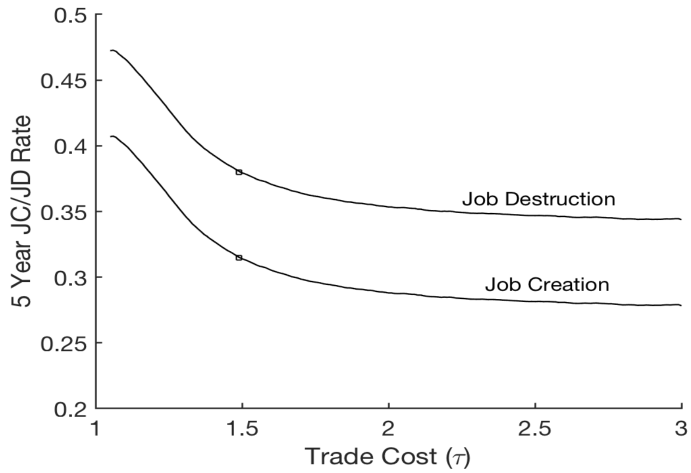

Note: The figure simulates the steady state job creation and destruction rates when we vary τ but keeping constant all other parameters at the values in Table 28.

average employment) for different values of τ. Consistent with the evidence from the U.S. and Canada (fact 5), in the model job destruction from large firms rises when trade costs fall.

Next we show the model predicts an increase in job creation *from exports* (fact 6) in the aftermath of trade liberalization. In Figure 29 (left panel), the job creation rate from exports increases by 10 percentage points when trade costs fall from our benchmark value ($\tau = 1.491$) to completely free trade ($\tau = 1$). The right panel in Figure 29, meanwhile, plots the predicted job destruction rate from domestic firms who are replaced by imports. The model predicts that moving to frictionless trade would raise the overall job destruction rate by 10 percentage points. Unfortunately we do not have an empirical counterpart for this model statistic.

The left panel in Figure 30 plots the simulated distribution of employment for exporters and non-exporters in the steady state, using our baseline parameter values from Table 28. These are the model analogues to facts 7 and 8. Firm size is determined by the number of products the firm controls, and whether the quality of the product is high enough to overcome the trade friction. A firm that exports has at least one product whose quality is high enough to overcome the trade

Figure 28: Simulated Job Destruction from *Large* Firms vs. Trade Costs

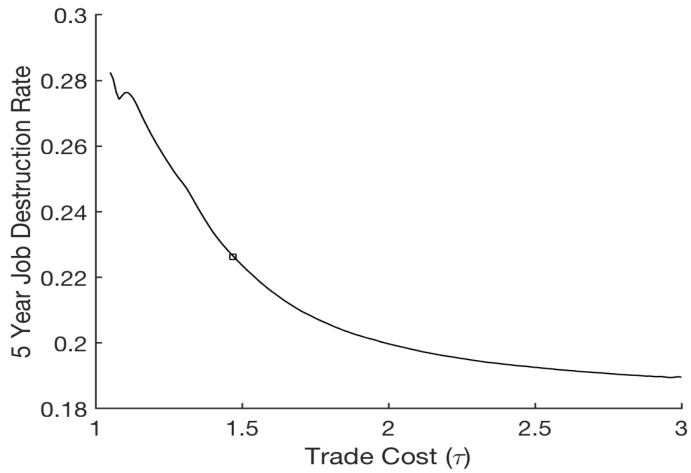

Note: The figure simulates the steady state job destruction rate from large (above mean employment) firms when we vary τ but keeping constant all other parameters at the values in Table 28.

Figure 29: Simulated Job Flows from Trade vs. Trade Costs

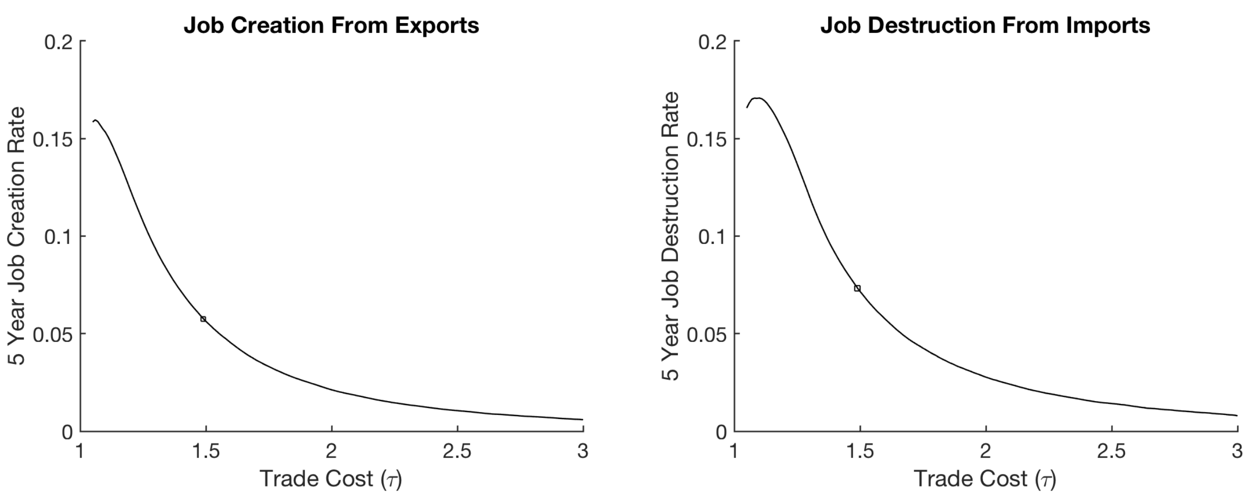

Note: The left panel shows the job creation rate due to exports and the right panel the job destruction rate due to imports in the steady state for different values of τ. All other parameters of model are kept fixed at the values in Table 28.

Figure 30: Simulated Distribution of Employment and Labor Productivity

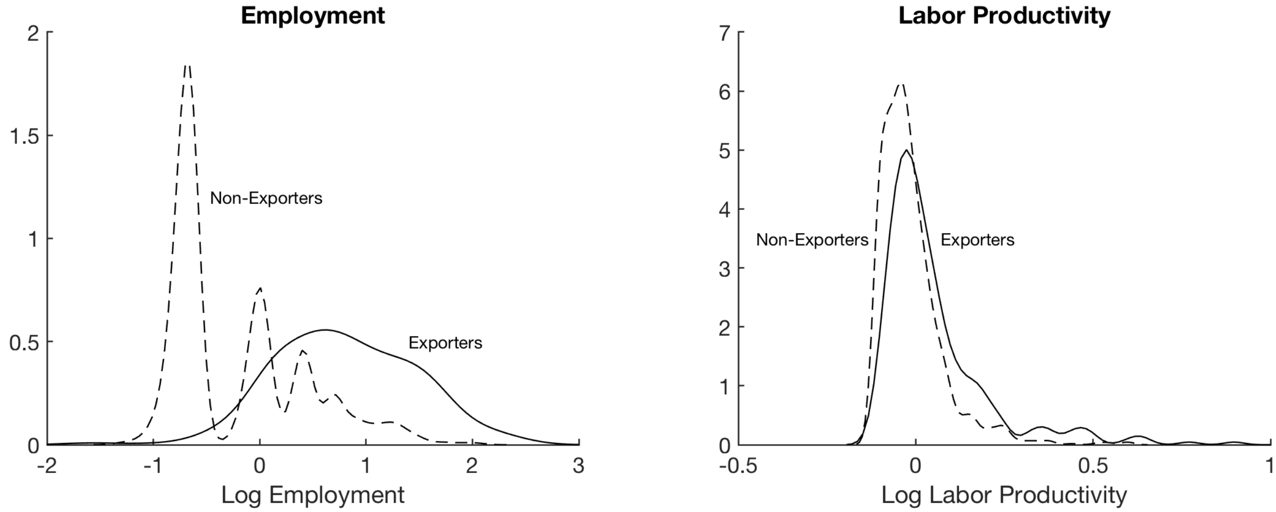

Note: The distributions of employment (left panel) and labor productivity (right panel) in the steady state of the model with the parameters given in Table 28.

cost. In the model, this probability is increasing in the firm's number of products. Larger firms own more products, and firms with more products are more likely to have at least one product with sufficient quality to export. The gap in average size between exporters and non-exporters is not due to any fixed cost of exporting, but rather the difference in the number and quality of products between the two groups of firms. Consistent with the empirical distribution of employment in Figure 24, the model predicts substantial overlap in the distribution of firm size of exporters and non-exporters. Although firms with fewer products are less likely to export, some of these products are high enough quality to overcome the trade cost.

The model also predicts that labor productivity is higher, on average, among exporters compared to non-exporters. This can be seen in the right panel in Figure 30, which plots the distribution of labor productivity (revenue per worker) of exporters and non-exporters. In the model, dispersion in labor productivity is entirely driven by markup heterogeneity. Since the markup is given by the quality gap between the best and the second best blueprint, this gap is increasing in the quality of the best blueprint. Because a firm with high quality varieties is also more likely to export, such firms are also more likely to charge higher markups. The gap in average labor productivity between exporters and non-exporters reflects the gap in average quality between the two groups of firms. This is similar to Bernard et al. (2003).

A key parameter that governs the gap in average quality between exporters and non-exporters (and quality dispersion more generally) is the shape parameter of the Pareto distribution of innovation draws θ. We therefore calibrate this parameter to make the simulated model match the average gap in labor productivity between exporters and non-exporters in the U.S. data.[109]

Figure 24 makes clear that, while average labor productivity (and thus quality when viewed through the lens of our model) of exporters is higher than for non-exporters on average, there is substantial overlap in their distributions. As depicted in Figure 30, our model generates such an overlap because many non-exporters. This reflects overlap in markups between the two groups of firms.

To be fair, the empirical *dispersion* of employment and labor productivity (Figure 24) is substantially larger than in the simulated data (Figure 30). Our assumption that preferences over varieties is Cobb-Douglas implies that product quality only matters for employment when higher quality helps the firm overcome the trade barrier. Conditional on selling in a given market at a given price, product quality has no effect on employment. We could make product quality matter more for firm employment, and thus get more employment dispersion, if we relaxed the Cobb-Douglas assumption. As for the dispersion of labor productivity, our model abstracts from differences in factor costs, overhead costs, adjustment costs and measurement error, all of which are likely present in the data and behind some of the dispersion in labor productivity in the data.[110] We leave these useful extensions for future work.

4 Endogenous Innovation

A key assumption we have made so far is that the innovation rates are exogenous parameters. We now consider the effect of a reduction in trade costs when innovation rates are endogenously

[109] As stated in Table 27, this gap is 6.6%.

[110] Bartelsman, Haltiwanger and Scarpetta (2013) emphasis the role of overhead costs, Asker, Collard-Wexler and De Loecker (2014) the importance of adjustment costs, and Bils, Klenow and Ruane (2018) the contribution of measurement error.

determined.

Suppose the innovation rate (per variety owned) of a domestic incumbent satisfies

$$\lambda = \left(\frac{R_i}{\gamma \chi_i \bar{A}^{(1-\phi)/\gamma}} \right)^{\gamma}, \tag{40}$$

where R_i denotes labor used for research (per variety owned), $\bar{A}$ is the geometric average quality of products *sold* in the domestic market, χ_i is an efficiency parameter, $\gamma < 1$ captures the returns to research effort, and ϕ captures the external returns to the stock of ideas. As in Klette and Kortum (2004), underlying (40) is the assumption of constant returns at the firm level to research effort and the number of varieties the firm owns (i.e., elasticities of γ and $1-\gamma$, respectively). When $\phi < 1$ we have diminishing returns to the stock of ideas so growth is semi-endogenous and linked to the population growth rate as in Jones (1995).

Similarly, suppose the unconditional innovation rate of domestic entrants is

$$\widetilde{\eta} = \left(\frac{R_e}{\gamma \chi_e \bar{A}^{(1-\phi)/\gamma}} \right)^{\gamma}, \tag{41}$$

where R_e is labor used for research (per variety in the economy) by potential entrants and χ_e is an efficiency parameter.[111]

The return to innovation is the product of the probability of grabbing a variety from another firm and the expected value of that variety. The new product can either be sold in both markets or only in the domestic market, and the value of this new product depends on whether it is traded or non-traded. So the return to innovation is the sum of the expected value of a traded product and the expected value of a non-traded product (multiplied by the probability of getting each type of product).

It will be convenient to normalize the value of a product by $\bar{A}^{(1+\gamma-\phi)/\gamma}$. We define v_x and v_n as the *expected* normalized value of a traded and non-traded product. The following arbitrage

[111] The innovation rates for foreign firms are given by equations analogous to (40) and (41) with R_i and χ_i replaced by R_i^* and χ_i^* in (40), R_e and χ_e replaced by R_e^* and χ_e^* in (41), and average quality sold in the foreign market instead of in the home market.

equation pins down v_x at time t:

$$
\begin{aligned}
v_{x,t} = \pi_{x,t} - \gamma\chi_i\lambda_t^{\frac{1}{\gamma}} \\
+ \frac{(1+g_t)^{\phi}}{1+r_t}\left[\lambda_t\left(\beta_{x,t}v_{x,t+1} + \beta_{n,t}v_{n,t+1}\right)\right] \\
+ \frac{(1+g_t)^{\phi}}{1+r_t}\left[\left(1-\delta_{x,t}\right)v_{x,t+1} - \delta'_{x,t}\left(v_{x,t+1} - v_{n,t+1}\right)\right].
\end{aligned}
\tag{42}
$$

Here g denotes the growth rate of the real wage, r the interest rate, π_x expected profits (normalized by $\bar{A}^{(1+\gamma-\phi)/\gamma}$), β_x and β_n the probability conditional on innovating of grabbing a traded and non-traded product, δ_x the probability of losing a traded variety in both markets, and δ'_x the probability of losing a traded product only in the foreign market.[112] The first term in equation (42) is profit net of research expenses (normalized by $\bar{A}^{(1+\gamma-\phi)/\gamma}$), the second term is the expected value of grabbing a new product, and the last term is the expected value of an exported variety in the next period minus the expected value of losing the product to a competitor in the foreign market.

The corresponding arbitrage equation for the non-traded product v_n is

$$
\begin{aligned}
v_{n,t} = \pi_{n,t} - \gamma\chi_i\lambda_t^{\frac{1}{\gamma}} \\
+ \frac{(1+g_t)^{\phi}}{1+r_t}\left[\lambda_t\left(\beta_{x,t}v_{x,t+1} + \beta_{n,t}v_{n,t+1}\right)\right] \\
+ \frac{(1+g_t)^{\phi}}{1+r_t}\left[\left(1-\delta_{n,t}\right)v_{n,t+1}\right].
\end{aligned}
\tag{43}
$$

where π_n denotes expected profits (normalized by $\bar{A}^{(1+\gamma-\phi)/\gamma}$) of a non-traded product and δ_n is the probability a non-traded product is creatively destroyed by another firm.

The privately optimal innovation rates are given by equating the marginal revenue from innovation to the marginal cost of innovation, which yields:

$$
\lambda_t = \left(\frac{\beta_{x,t}v_{x,t} + \beta_{n,t}v_{n,t}}{\chi_i}\right)^{\frac{\gamma}{1-\gamma}} \tag{44}
$$

[112] δ_x and δ'_x are the creative destruction rates in Table 26. Since innovation is not targeted, β_x and β_n depend on the share of each type of product and the creative destruction rates.

$$\widetilde{\eta}_t = \left(\frac{\beta_{x,t} v_{x,t} + \beta_{n,t} v_{x,t}}{\chi_e} \right)^{\frac{\gamma}{1-\gamma}} \tag{45}$$

An increase in v_x and v_n raises the innovation rate with an elasticity that depends on γ. As in the model where innovation is exogenous, the equilibrium is determined by equations (35), (36), (37), and the markup formulas in Table 24, except that the innovation rates are now pinned down by (42) through (45).[113]

In steady-state, β_x, β_n, v_x, and v_n are constant so the innovation rates λ and η are constant as well. Differences in innovation rates between countries now reflect differences in the innovation cost parameter χ, but it is still the case that differences in innovation rates show up as differences in the relative wage. What is new in the endogenous innovation model is that the growth rate of the real wage in steady state is ultimately given by the product of the population growth rate and $\gamma/(1-\phi)$ with $\phi < 1$.

We set ϕ and γ to match the empirical growth rate of real wages (TFP growth), the growth rate of the "population" (growth of investments in intellectual property products), and the share of "labor" devoted to research (share of value added invested in intellectual property products).[114] This yields $\phi = 0.165$ and $\gamma = 0.61$. The calibrated values of these parameters are given in Table 30. The innovation cost parameter χ is lower in the U.S. to generate higher U.S. innovation rates and match the higher wages in the U.S. (relative to the foreign country).

We can now re-examine the effect of reducing trade costs, this time with endogenous arrival rates of innovation. Here, trade liberalization has no permanent effect on the long run growth rate, which is pinned down by population growth. The initial increase in the growth rate due to trade liberalization increases the level of TFP, but with $\phi < 1$ this raises the cost of innovation. Thus, in the new steady state with lower trade costs, innovation rates by each country are actually *lower* but the growth rate of TFP is the same.

[113]Plus the corresponding optimal innovation rates for foreign firms. We assume linear utility so that the consumption Euler equation implies $r = \rho$. We set ρ = 0.05.

[114]TFP growth averaged 3.01% per year from 1995–2008 according to the U.S. Bureau of Labor Statistics KLEMS data. Real intellectual property investments grew 4.12% per year from 1995–2008 according to the U.S. Bureau of Economic Analysis, and such investments averaged 7.03% of value added in U.S. manufacturing from 1997–2008.

Table 30: Estimates of Model Parameters, Endogenous Innovation

Variable	Description	Value
ϕ	Return to stock of ideas	0.165
γ	Return to research intensity	0.61
χ_e/χ_i	Home entrant/incumbent research cost	2.89
χ_i^*/χ_i	Foreign/home incumbent research cost	7.26
χ_e^*/χ_i	Foreign entrant/home incumbent research cost	16.98

Figure 31 plots the effect of a permanent, unanticipated reduction of trade costs, from $\tau = 1.491$ to $\tau = 1.245$, on innovation rates. It shows that innovation rates initially rise in the aftermath of a reduction in trade costs. But as the level of TFP rises, arrival rates fall due to the rising difficulty of innovating. This is due to diminishing returns to the stock of ideas $\phi < 1$. In the new steady state, innovation rates are *lower* compared to the initial steady state, though TFP is on a higher path (parallel to its initial path). TFP is higher despite lower arrival rates within each country because ideas flow across countries more with lower trade costs.

Figure 32 plots the effect of a permanent reduction in trade costs on the share of labor devoted to research in each country. Like the arrival rates, the shares spike on impact. Unlike the arrival rates, the share of labor doing research ends up higher in the long run. The bigger market for each successful innovation makes higher research effort worthwhile despite the endogenously greater difficulty in coming up with ideas. This result contrasts with Eaton and Kortum (2001), wherein these two forces exactly offset each other and leave research effort unchanged.

Figure 33 simulates the effect of permanently lower trade costs on the level of real consumption. Real consumption is closely tied to TFP in each country, but also accounts for tariff revenues and the share of labor diverted from production to R&D. The figure expresses variables relative to the path of U.S. consumption in the absence of trade liberalization. Real consumption

Figure 31: Simulated arrival rates after trade liberalization

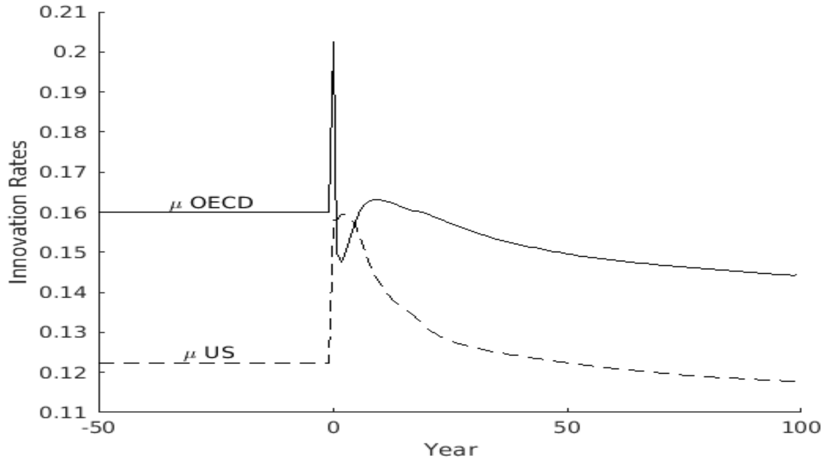

Note: The figure simulates the path of arrival rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 30. Here $\mu = \lambda + \eta$, where λ and η are the endogenous arrival rates of innovation by incumbents and entrants, respectively.

Figure 32: Research labor shares after trade liberalization

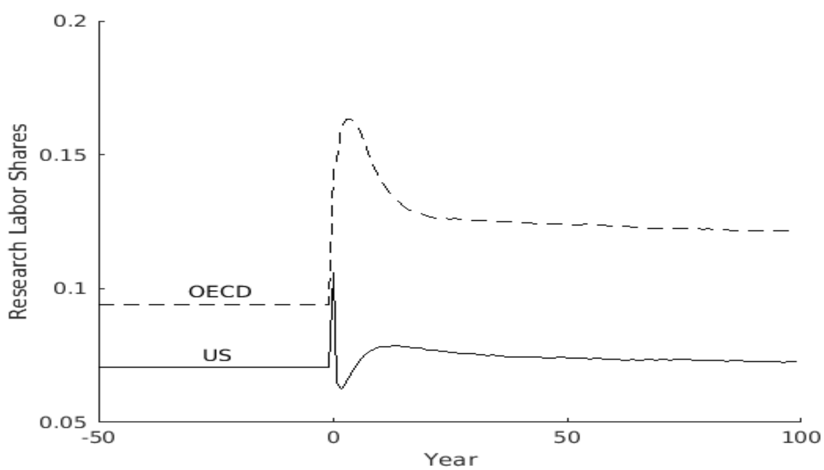

Note: The figure simulates the share of research labor in total employment in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 30.

Figure 33: Simulated real consumption after trade liberalization

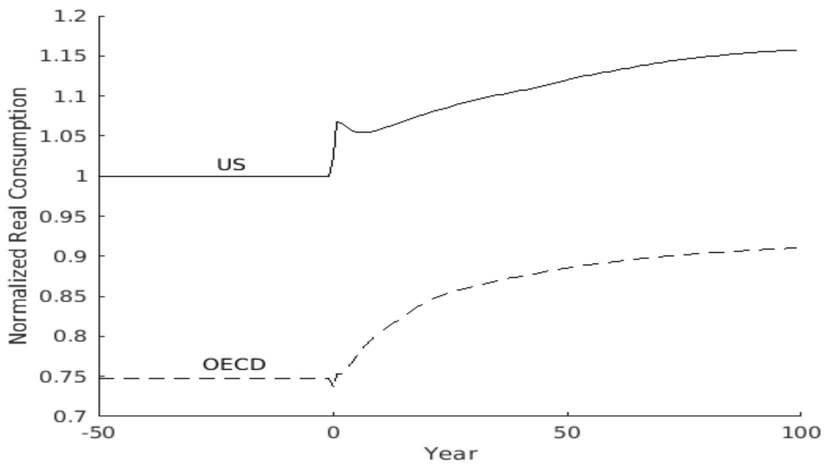

Note: The figure simulates the path of real consumption in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 30.

rises upon impact of liberalization in the US but initially grows at a slower rate in the OECD due to the temporary large increase in the share of workers engaged in innovation rather than production. Following the initial shock, the growth rate of consumption increases temporarily for a few decades in response to a decline in trade costs, but eventually slows down as innovation becomes more difficult with higher TFP. In the new steady state, real consumption and TFP are permanently higher (compared to the initial steady state path), but the growth rate is the same as in the initial equilibrium. As in the exogenous innovation case, the rest of the OECD gains more than the U.S. because the U.S. is more innovative.

Figure 34 maps the impact of a host of different tariff levels on the level of real consumption in the long run. Real consumption is almost always higher as a result of lower tariffs. The real wage actually falls, however, as frictionless trade is approached. At very low tariff levels, the high rate of creative destruction from imports discourages research effort so much that it outweighs the quicker spread of ideas.

Figure 35 illustrates how freer trade affects job flows. Consistent with the evidence from the

Figure 34: Simulated steady-state real consumption vs. various trade costs

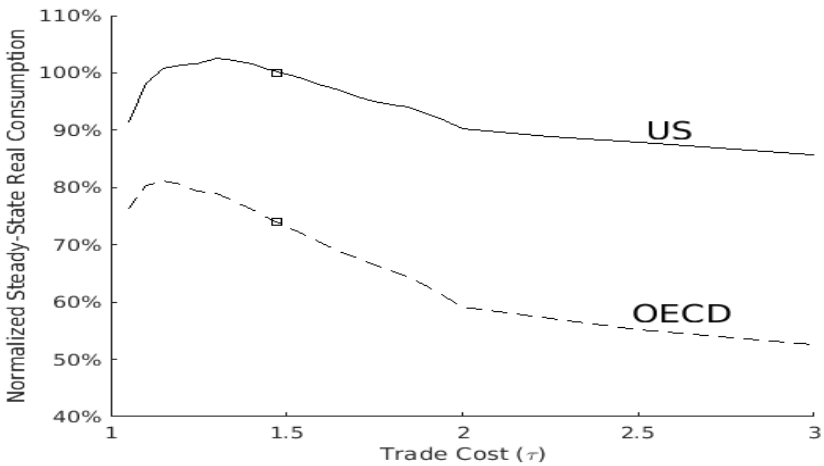

Note: The figure plots simulated equilibrium real consumption in steady-states with differing trade costs, keeping constant all other parameters at the values in Table 30.

U.S. and Canada after CUSFTA, job flows surge in the aftermath of a reduction in trade costs. The pace of job flows remains elevated for decades after a tariff reduction — certainly within the 15-year window we examine after the 1988 U.S.-Canada Free Trade Agreement. This pattern drives home that there can be dynamics costs (job destruction) as well as dynamic benefits (knowledge flows) to trade liberalization.[115]

In Table 32 we present the welfare gains from lower trade barriers. Here we measure welfare as the present discounted value of consumption, consistent with linear utility and a constant real interest rate, using a discount rate of 5%. In the first counterfactual, we reduce tariffs from 4, a high tariff at which the aggregate trade share is about 1%, to our estimated value of 1.491. For comparison, we first calculate the static gains — the jump in real consumption on impact. The static gains are 23.9% in the U.S. and 21.5% in the rest of the OECD. The U.S. is the smaller economy (due to smaller population) with the larger trade share, so it gains more.[116]

[115] Bernard, Redding and Schott (2007) analyze a multi-sector Melitz model which also features effects of trade liberalization on steady state job flows.

[116] Our static gains from trade are larger than implied by the formula in Arkolakis, Costinot and

Figure 35: Simulated job flows from trade liberalization

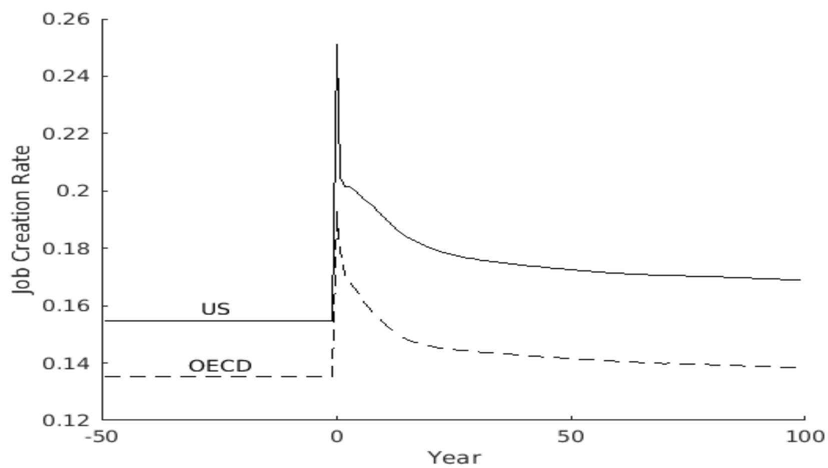

Note: The figure simulates job creation and destruction rates in response to an unanticipated reduction in trade costs, keeping constant all other parameters at the values in Table 30.

Next we present the full gains from freer trade, including effects on innovation. With exogenous arrival rates of innovation, the total gains are quite large at 37% in the U.S. The rest of the OECD gains even more (104%) because it gets more ideas than it gives. The dynamic gains are markedly smaller with endogenous innovation rates, but still sizable at about 30% in the U.S. and 45% in the rest of the OECD. The gains are much more modest with endogenous research effort for two reasons. First and foremost, we built in diminishing returns to the stock of ideas ($\phi < 1$) and a congestion externality ($\gamma < 1$). These temper the cumulative TFP gains from endogenously rising research effort and a faster flow of ideas across countries. Second, the higher research effort comes at the cost of less labor devoted to production and consumption. Even in the endogenous innovation case, however, the full gains are an order or magnitude larger than the static gains.

The second counterfactual in Table 32 reduces tariffs from our estimated value of 1.491 to

Rodriguez-Clare (2012). When the trade share goes from 1% to 10.2%, their formula predicts gains of 1.9% (9.2%/5), if one uses our trade elasticity of 5. This could be because our Cobb-Douglas preferences result in less substitutability across products than in the class of models ACR examine.

Table 31: Gains From Trade

	Real Consumption	
Method	U.S.	OECD
Static Gains	23.9%	21.5%
Dynamic Gains - Exogenous Innovation	37.0%	104.0%
Dynamic Gains - Endogenous Innovation	30.1%	44.5%

Entries give the percentage increase in the present discounted value of the real (consumption) wage and real consumption as a result of reducing tariffs from 4 to 1.491. The aggregate trade share at $\tau = 4$ is about 0.1%. We use a discount rate of 5%.

Table 32: Gains From Trade Liberalization

	Real Consumption	
Method	U.S.	OECD
Static Gains	5.5%	3.6%
Dynamic Gains - Exogenous Innovation	9.6%	15.5%
Dynamic Gains - Endogenous Innovation	9.7%	13.8%

Entries are the percentage increase in the present discounted value of the real wage and real consumption as a result of reducing tariffs from 1.491 to 1.245. We use a discount rate of 5%.

1.245. The dynamic gains are smaller here than when going from near-autarky to current trade flows. The ideas flowing after further trade liberalization involve endogenously smaller step sizes. The dynamic gains are still two to five times bigger than the static gains.

5 Additional Exercises

5.1 U.S.-Canada Simulations

The U.S.-Canada Free Trade Agreement provided three of the nine motivating facts we offered in Section 2. We emphasized that job flows, job destruction at large firms, and job creation associated with exports all increased noticeably in Canada after the agreement. Here we simulate a simplified version of CUSFTA in our model with endogenous arrival rates.[117] Specifically, we analyze the effect of lowering tariffs from 1.39 to 1.25 to match Canada's trade shares of 25% before CUSFTA (1978–1988) and 37% after CUSFTA (1988–2003). We re-estimate other parameter values so that the model matches the relative wage and population between the U.S. and Canada.

Table 33 presents the changes in job flow moments in Canada after trade liberalization.[118] The first column has (changes in) empirical moments, and the second column the changes in simulated moments. The model qualitatively matches the patterns in the data: rising job creation, job destruction, job destruction from large firms, and job creation from exports. The model somewhat overstates overstates job reallocation and job creation from exports, but somewhat understates job destruction from large firms. Note that no parameter values were chosen to target any of the data moments in Table 33.

5.2 Alternative Assumptions about Idea Flows

All of our simulations so far assume that innovators build on the productivity level of *sellers* into the domestic market. To gauge the importance of this assumption for the dynamic gains from trade, we now consider two alternative assumptions about idea flows across countries. For simplicity, these simulations involve exogenous arrival rates of innovation. And they involve going from near-autarky to current levels of trade.

[117]Until now, we have carried out simulations with the rest of the OECD because its larger size helped drive home the importance of scale for the dynamic gains from trade. Because the rest of the OECD is so large, trade flows with them are larger and, presumably, idea flows as well.

[118]We focus on changes in Canada rather than the U.S. because CUSFTA was a much bigger shock to the Canadian economy given its large trade share with the U.S.

Table 33: Job Flows in Canada — Post-CUSFTA versus Pre-CUSFTA

	Data	Simulations
Job Creation Rate	+1.3%	+7.9%
Job Destruction Rate	+7.2%	+7.9%
Job Destruction from Large Firms	+9.1%	+5.0%
Job Creation from Exports	+8.7%	+10.1%

Note: Pre-CUSFTA is 1978 to 1988. Post-CUSFTA is 1988 to 2003. Job creation and destruction calculated over five year periods. Simulations use a version of the model estimated to match the relative wage and population between the U.S. and Canada, with a trade liberalization event that matches the twelve percentage point increase in Canadian export share from the pre-CUSFTA to post-CUSFTA period.

We first assume that innovators probabilistically build on sellers (with probability κ) or on domestic *producers* (with probability $1 - \kappa$).[119] We next assume that innovators build on all products (with probability ν) or on the subset of products that are domestically produced (with probability $1 - \nu$). Such research specialization ($\nu < 1$) allows countries to experience more frequent innovations on the subset of products they produce, and more so the higher the share of products imported.

Table 34 presents our findings. Column one reproduces our baseline specification of $\kappa = 1$ and $\nu = 1$ to facilitate comparison. Column two reduces the frequency with which innovators learn from sellers to 30% ($\kappa = 0.3$), but does not allow research specialization ($\nu = 1$). Doing so dampens the increase in job reallocation from trade in the U.S. and OECD and substantially reduces the steady-state growth kick from trade from 47 basis points per year in the baseline to 13 basis points per year with limited idea flows. The consumption-equivalent gains from trade fall from 37% to 27% in the U.S., and from 104% to 48% in the rest of the OECD.[120]

[119] If a product is imported, innovators build on the quality of the last domestic producer of the product. There is sufficient turnover of imported products to domestic producers so that there is always a previous (fallow) domestic producer.

[120] Recall that the rest of the OECD gains more from idea flows than the U.S. because it is less innovative in our calibration. It therefore suffers more from limiting idea flows.

Column three of Table 34 allows for research specialization along with limited idea flows ($\kappa = 0.3$ and $\nu = 0.3$). The effect of trade on steady-state growth and job reallocation here is similar to that of the baseline specification, though the consumption-equivalent gains from trade are smaller in the OECD because the growth kick from research specialization in the less innovative country occurs gradually relative to the gains it receives from technology spillovers. The U.S. actually gains slightly more in consumption-equivalent terms in this case, 40% versus 38% in the baseline. Thus, even if ideas flow in a more limited way, trade may create large dynamic gains if it facilitates research specialization.

How can we assess the realism of these assumptions? All three cases feature increases in job reallocation after trade liberalization, as seen in Canada after CUSFTA.[121] But some insight can be gleaned from product turnover. Recall that we reported in Table 23 that the SITC with the highest level of exports in a given year tend to account for a lower share of total exports five years earlier (66.5%). In our baseline simulation, this fraction is similar at 77.5%.[122] In simulations with $\kappa = 0.3$ and $\nu = 1$, the fraction rises to 87.9%. It will rise further if we move to even lower κ values. Specialized research ($\nu < 1$) lowers product turnover more dramatically. In the case where $\kappa = 0.3$ and $\nu = 0.3$, the export share of the top product five years earlier was 95.4% of its current share. Thus deviating from our baseline would move us further away from the data on product churn.

When we consider even lower values of κ and ν we find that growth rates diverge between the U.S. and Canada or the rest of the OECD in our simulations. Thus the similarity of observed growth rates across these countries (Klenow and Rodriguez-Clare, 2005) is another piece of evidence in favor of tying them together through idea flows.

[121]We do not entertain a case where trade has no impact at all on growth and job reallocation, as such a model cannot explain the elevated job reallocation observed after CUSFTA.

[122]We divide our 5,000 products randomly into 125 categories, identify the top export "product" (category) in each year, and calculate its average relative share 5 years prior.

Table 34: Counterfactuals with Alternative Assumptions on Idea Flows

Near-Autarky to Estimated τ	$\kappa = 1, \nu = 1$	$\kappa = 0.3, \nu = 1$	$\kappa = 0.3, \nu = 0.3$
Δ Job Reallocation U.S.	3.9%	1.9%	3.5%
Δ Job Reallocation OECD	1.7%	0.7%	2.2%
Δ Growth	0.47%	0.13%	0.48%
Gains from Trade U.S.	37.0%	27.3%	40.3%
Gains from Trade OECD	104.0%	47.8%	33.1%
Top Export Product Turnover U.S.	77.5%	87.9%	95.4%
Top Export Product Turnover OECD	78.4%	87.2%	96.4%
Prob. losing exports U.S.	12.8%	5.9%	1.1%
Prob. losing exports OECD	16.7%	7.8%	1.1%

Note: Column one is the baseline specification. Column two reduces the frequency with which innovators learn from foreign producers to 30% ($\kappa = 0.3$); the other 70% of draws build on domestic producers. Column three allows for research specialization, so that 70% of innovation ($\nu = 0.3$) draws are in products a country produces; the other 30% draw from all products. In each case, innovation parameters are re-estimated to match the targeted moments while θ and ψ are held fixed, and counterfactuals are shown for the exogenous arrival rate near-autarky and estimated steady-states. Estimated τ is 1.491 for column one, 1.78 for column two, and 2.49 for column three. Near-autarky is defined as the τ that achieves a trade share of approximately 0.5%. The gains from trade are in terms of consumption-equivalent welfare.

6 Conclusion

We documented facts about trade and job reallocation in U.S. and Canadian manufacturing in recent decades. After the U.S.-Canada Free Trade Agreement in 1988, job destruction rates spiked and remained elevated through 2012 (our latest year of data). The increase in job destruction rates and exit from exporting occurred equally at big and small firms in Canada.

We used these facts as motivation to construct a two-country model of creative destruction and trade. In the model, lower tariffs increase the probability that foreign firms take over domestically produced products. As in the data, this occurs in an initial burst but remains higher for decades. Foreign and domestic firms take over each other's markets more frequently when trade barriers are lower. This dynamic disruption is a byproduct of faster growth. Growth remains high in the version of the model with exogenous innovation rates.

When we endogenize innovation rates and build in diminishing returns, lower tariffs boost growth only temporarily. Still, trade liberalization raises levels of productivity permanently. Compared to (near) autarky, such dynamic gains are an order of magnitude larger than the static gains from trade in our model. Trade, like technological change more generally, brings dynamic benefits but also dynamic costs.

We see several potential directions for future research. One is to explicitly incorporate frictions to reallocating workers in response to trade-induced creative destruction. Another route is to study events such as China joining the WTO. A third avenue would be to obtain more direct empirical discipline on the form and magnitude of knowledge spillovers (e.g. the frequency of imitation of rich country producers by developing country producers, or of learning from domestic producers vs. foreign sellers in the local market). As we stressed, we think idea flows will need to be tied to trade flows somehow to explain why trade liberalization ushers in more rapid job reallocation in a sustained way.

We end with a conjecture about optimal innovation policy in our setting. Because of domestic knowledge spillovers, national governments may find it optimal to subsidize domestic R&D. But they might not internalize knowledge spillovers to foreign producers who build on domestic innovations. The world might need “Global Technical Change” accords to internalize these positive externalities, just as we need Global Climate Change agreements to internalize negative pollution externalities.

References

Acemoglu, Daron and Jaume Ventura, "The World Income Distribution," *Quarterly Journal of Economics*, 2002, *117* (2), 659–694.

__ , Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, 2018, *108* (11), 3450–91.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), pp. 323–351.

__ , Antonin Bergeaud, Matthieu Lequien, and Marc Melitz, "The Impact of Exports on Innovation: Theory and Evidence," 2018.

Akcigit, Ufuk, Sina T. Ates, and Giammario Impullitti, "Innovation and Trade Policy in a Globalized World," 2018.

Alfaro, Laura and Maggie X Chen, "Selection and market reallocation: Productivity gains from multinational production," *American Economic Journal: Economic Policy*, 2018, *10* (2), 1–38.

Alvarez, Fernando E, Francisco J. Buera, and Robert E. Lucas, "Idea Flows, Economic Growth, and Trade," 2013.

Arkolakis, Costas, Arnaud Costinot, and Andres Rodriguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Asker, John, Allan Collard-Wexler, and Jan De Loecker, "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 2014, *122* (5), 1013–1063.

Autor, David, David Dorn, and Gordon Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Baqaee, David Rezza and Emmanuel Farhi, "The macroeconomic impact of microeconomic shocks: beyond Hulten's Theorem," 2017.

Barbose, Galen, "U.S. renewable portfolio standards: 2018 annual status report," *Lawrence Berkeley National Lab (LBNL), Berkeley, CA*, 2018.

Barnes, Justin, "DSIRE dataset on renewable portfolio standards," *Database of State incentives for Renewables & Efficiency (DSIRE), North Carolina State University, Raleigh, NC*, 2014.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 2013, *103* (1), 305–334.

Baumol, William J and William G Bowen, "Performing Arts: The Economic Dilemma (New York: The Twentieth Century Fund)," *BaumolThe Performing Arts The Economic Dilemma1966*, 1966.

Bernard, Andrew B. and J. Bradford Jensen, "Exceptional Exporter Performance: Cause, Effect, or Both?," *Journal of International Economics*, 1999, *47* (1), 1–25.

—, Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, "Plants and Productivity in International Trade," *American economic review*, 2003, *93* (4), 1268–1290.

Bernard, Andrew B, Stephen J Redding, and Peter K. Schott, "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 2007, *74* (1), 31–66.

Bils, Mark, Peter J. Klenow, and Cian Ruane, "Misallocation or Mismeasurement?," 2018.

Bjørner, Thomas, Mikael Togeby, and Henrik Jensen, "Industrial Companies' Demand for Electricity: Evidence From a Micropanel," *Energy Economics*, 2001, *23* (5), 595–617.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," 2018.

—, Mirko Draca, and John Van Reenen, "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity," *The Review of Economic Studies*, 2016, *83* (1), 87–117.

Borenstein, Severin, "The Market Value and Cost of Solar Photovoltaic Electricity Production," 2008. Working Paper.

—, "To what electricity price do consumers respond? Residential demand elasticity under increasing-block pricing," 2009.

Buera, Francisco J. and Ezra Oberfield, "The Global Diffusion of Ideas," 2017.

Burnside, Craig and Martin Eichenbaum, "mFactor $HoardingandthePropagationofBusiness$ Cycle Shocks. m The American Economic Review," 1996.

—, —, and Sergio Rebelo, "Labor hoarding and the business cycle," *Journal of Political Economy*, 1993, *101* (2), 245–273.

Bushnell, James, Michaela Flagg, and Erin Mansur, "Capacity markets at a crossroads," *Report to the Department of Energy, Office of Energy Policy and Systems Analysis, Washington DC*, 2017.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural productivity and structural transformation: Evidence from Brazil," *American Economic Review*, 2016, *106* (6), 1320–65.

Carleton, Tamma, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Amir Jina, Robert E Kopp, Kelly McCusker, Ishan Nath et al., "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," 2018.

Chen, Cliff, Ryan Wiser, and Mark Bolinger, “Weighing the Costs and Benefits of State Renewables Portfolio Standards: A Comparative Analysis of State-Level Policy Impact Projections,” Technical Report 61580, LBNL 2007.

Chen, Wei, Xilu Chen, Chang-Tai Hsieh, and Zheng Michael Song, “A Forensic Examination of Chinas National Accounts,” 2019.

Cline, William R, *Global warming and agriculture: End-of-century estimates by country*, Peterson Institute, 2007.

Coe, David T., Elhanan Helpman, and Alexander W. Hoffmaister, “North-South R&D Spillovers,” *Economic Journal*, 1997, *107* (440), 134–149.

— , — , and — , “International R&D Spillovers and Institutions,” *European Economic Review*, 2009, *53* (7), 723–741.

Comin, Diego A, Danial Lashkari, and Martí Mestieri, “Structural change with long-run income and price effects,” 2015.

Costinot, Arnaud, Dave Donaldson, and Cory Smith, “Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world,” *Journal of Political Economy*, 2016, *124* (1), 205–248.

Cuaresma, Jesús Crespo, “Income projections for climate change research: A framework based on human capital dynamics,” *Global Environmental Change*, 2017, *42*, 226–236.

Cullen, Joseph, “Measuring the environmental benefits of wind-generated electricity,” *American Economic Journal: Economic Policy*, 2013, *5* (4), 107–133.

Davis, Steven J, John C Haltiwanger, and Scott Schuh, “Job Creation and Destruction,” *MIT Press Books*, 1996.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, “The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism,” *Journal of Economic Perspectives*, 2014, *28* (3), 3–24.

Denholm, Paul and Robert M Margolis, “Evaluating the limits of solar photovoltaics (PV) in traditional electric power systems,” *Energy policy*, 2007, *35* (5), 2852–2861.

Deryugina, Tatyana and Solomon M Hsiang, “Does the environment still matter? Daily temperature and income in the United States,” 2014.

Deschenes, Olivier and Michael Greenstone, “The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather,” *American Economic Review*, 2007, *97* (1), 354–385.

Deschênes, Olivier and Michael Greenstone, “Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US,” *American Economic Journal: Applied Economics*, 2011, *3* (4), 152–85.

Desmet, Klaus and Esteban Rossi-Hansberg, "On the spatial economic impact of global warming," *Journal of Urban Economics*, 2015, *88*, 16–37.

Dingel, Jonathan I, Kyle C Meng, and Solomon M Hsiang, "Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate," 2019.

Dix-Carneiro, Rafael and Brian K. Kovak, "Trade Liberalization and Regional Dynamics," *American Economic Review*, 2017, *107* (10), 2908–46.

Dornbusch, Rdiger, Stanley Fischer, and Paul Samuelson, "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, 1977, *67* (5), 823–39.

Driscoll, John C and Aart C Kraay, "Consistent covariance matrix estimation with spatially dependent panel data," *Review of economics and statistics*, 1998, *80* (4), 549–560.

Eaton, Jonathan and Samuel Kortum, "Technology, Trade, and Growth: A Unified Framework," *European Economic Review*, 2001, *45* (4-6), 742–755.

_ and _ , "Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

EEI, "Transmission projects: At a glance," *Edison Electric Institute (EEI)*, 2011.

EIA, "Annual Energy Outlook 2019," *U.S. Energy Information Administration (EIA)*, 2019.

EPA, "Social cost of carbon," *Environmental Protection Agency (EPA): Washington, DC, USA*, 2016.

Fabrizio, Kira, Nancy Rose, and Catherine Wolfram, "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency," *American Economic Review*, 2007, *97* (4), 1250–1277.

Feenstra, Robert C., Robert E. Lipsey, Haiyan Deng, Alyson C. Ma, and Hengyong Mo, "World Trade Flows: 1962-2000," *NBER Working Paper 11040*, 2005.

Fischer, Carolyn, "Combining Policies for Renewable Energy: Is the Whole Less Than the Sum of Its Parts?," *International Review of Environmental and Resource Economics*, jun 2010, *4* (1), 51–92.

GAO, "Four Regions Use Capacity Markets to Help Ensure Adequate Resources, but FERC Has Not Fully Assessed Their Performance," *U.S. Government Accountability Office (GAO) Report to Congressional Committees*, 2017.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow, "How Destructive is Innovation?," 2018.

Gillingham, Kenneth and James Stock, "The cost of reducing greenhouse gas emissions," *Journal of Economic Perspectives*, 2018, *32* (4), 53–72.

Gollin, Douglas, Casper Worm Hansen, and Asger Wingender, "Two blades of grass: The impact of the green revolution," 2018.

—, Stephen L Parente, and Richard Rogerson, "The food problem and the evolution of international income levels," *Journal of Monetary Economics*, 2007, *54* (4), 1230–1255.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez, "Capital allocation and productivity in South Europe," *The Quarterly Journal of Economics*, 2017, *132* (4), 1915–1967.

Gowrisankaran, Gautam, Stanley S Reynolds, and Mario Samano, "Intermittency and the value of renewable energy," *Journal of Political Economy*, 2016, *124* (4), 1187–1234.

Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton, "Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation," *Review of Environmental Economics and Policy*, 2013, *7* (1), 23–46.

Grossman, Gene M. and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 1991, *58* (1), 43–61.

— and — , *Innovation and Growth in the Global Economy*, MIT press, 1993.

Hanson, Gordon, Nelson Lind, and Marc-Andreas Muendler, "The Dynamics of Comparative Advantage," 2018.

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse, Toolkit, and Cookbook," *Handbook of International Economics*, 2014, *4*, 131–195.

Heeter, Jenny, Galen Barbose, Lori Bird, S Weaver, F Flores-Espino, K Kuskova-Burns, and R Wiser, "A Survey of State-Level Cost and Benefit Estimates of Renewable Portfolio Standards," Technical Report 6A20-61042, NREL/TP 2014.

Holmes, Thomas J. and John J. Stevens, "An Alternative Theory of the Plant Size Distribution, with Geography and Intra-and International trade," *Journal of Political Economy*, 2014, *122* (2), 369–421.

Hulten, Charles R, "Growth accounting with intermediate inputs," *The Review of Economic Studies*, 1978, *45* (3), 511–518.

Ito, Koichiro, "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," *American Economic Review*, 2014, *104* (2), 537–563.

Jacobson, Mark Z, Mark A Delucchi, Mary A Cameron, and Bethany A Frew, "Low-cost solution to the grid reliability problem with 100% penetration of intermittent wind, water, and solar for all purposes," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15060–15065.

Jeffrey, Stephen, Leon Rotstayn, MA Collier, Stacey Dravitzki, Carlo Hamalainen, Chris Moeseneder, KK Wong, and Jozef Syktus, "Australias CMIP5 submission using the CSIRO-Mk3. 6 model," *Aust. Meteor. Oceanogr. J*, 2013, *63*, 1–13.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, *103* (4), 759–784.

Jones, Charles I and Peter J Klenow, "Beyond GDP? Welfare across countries and time," *American Economic Review*, 2016, *106* (9), 2426–57.

Joskow, Paul L, "Comparing the costs of intermittent and dispatchable electricity generating technologies," *The American Economic Review*, 2011, *101* (3), 238–241.

Klenow, Peter J. and Andres Rodriguez-Clare, "Externalities and Growth," *Handbook of Economic Growth*, 2005, *1*, 817–861.

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, *65* (6), 1389–1419.

Lagakos, David and Michael E Waugh, "Selection, agriculture, and cross-country productivity differences," *American Economic Review*, 2013, *103* (2), 948–80.

Lamont, Alan D, "Assessing the long-term system value of intermittent electric generation technologies," *Energy Economics*, 2008, *30* (3), 1208–1231.

Lincoln, William, Andrew McCallum, and Michael Siemer, "The Great Recession and a Missing Generation of Exporters," 2017.

Lyu, Changjiang, Kemin Wang, Frank Zhang, and Xin Zhang, "GDP management to meet or beat growth targets," *Journal of Accounting and Economics*, 2018, *66* (1), 318–338.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw, "The impact of global warming on agriculture: a Ricardian analysis," *The American economic review*, 1994, pp. 753–771.

Milligan, Michael, Erik Ela, Bri-Mathias Hodge, Brendan Kirby, Debra Lew, Charlton Clark, Jennifer DeCesaro, and Kevin Lynn, "Integration of variable generation, cost-causation, and integration costs," *The Electricity Journal*, 2011, *24* (9), 51–63.

Mills, Andrew, Ryan Wiser, and Kevin Porter, "The cost of transmission for wind energy: A review of transmission planning studies," *Lawrence Berkeley National Lab (LBNL), Berkeley, CA*, 2009.

Perla, Jesse, Christopher Tonetti, and Michael E. Waugh, "Equilibrium Technology Diffusion, Trade, and Growth," 2015.

Rivera-Batiz, Luis A. and Paul M. Romer, "Economic Integration and Endogenous Growth," *Quarterly Journal of Economics*, 1991, *106* (2), 531–555.

Schlenker, Wolfram and David B Lobell, "Robust negative impacts of climate change on African agriculture," *Environmental Research Letters*, 2010, *5* (1), 014010.

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

Schmalensee, Richard, "Evaluating Policies to Increase Electricity Generation from Renewable Energy," *Review of Environmental Economics and Policy*, 2012, *6* (1), 45–64.

Seppanen, Olli, William J Fisk, and QH Lei, "Effect of temperature on task performance in officeenvironment," 2006.

Shrimali, Gireesh, Steffen Jenner, Felix Groba, Gabriel Chan, and Joe Indvik, "Have State Renewable Portfolio Standards Really Worked? Synthesizing Past Policy Assessments to Build an Integrated Econometric Analysis of RPS Effectiveness in the U.S.," 2012. Working Paper.

Somanathan, E, Rohini Somanathan, Anant Sudarshan, Meenu Tewari et al., "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing," Technical Report, Indian Statistical Institute, New Delhi, India 2015.

Teignier, Marc, "The role of trade in structural transformation," *Journal of Development Economics*, 2018, *130*, 45–65.

Tombe, Trevor, "The missing food problem: Trade, agriculture, and international productivity differences," *American Economic Journal: Macroeconomics*, 2015, *7* (3), 226–58.

Törnqvist, Leo, "The Bank of Finland's consumption price index," 1936.

Trefler, Daniel, "The Long and Short of the Canada-U.S. Free Trade Agreement," *American Economic Review*, September 2004, *94* (4), 870–895.

Tuerck, David., Paul Bachman, and Michael Head, "The Economic Impact of Wisconsin's Renewable Portfolio Standard," 2013, *26* (4).

Upton, Gregory and Brian Snyder, "Funding renewable energy: An analysis of renewable portfolio standards," *Energy Economics*, 2017, *66*, 205–216.

Uy, Timothy, Kei-Mu Yi, and Jing Zhang, "Structural change in an open economy," *Journal of Monetary Economics*, 2013, *60* (6), 667–682.

Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang, "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," *Journal of Environmental Economics and Management*, 2018, *88*, 1–17.

Zivin, Joshua Graff and Matthew Neidell, "Temperature and the allocation of time: Implications for climate change," *Journal of Labor Economics*, 2014, *32* (1), 1–26.

Appendix A: Additional Regression Results

Figure A-1: Predicted Heterogeneous Response of Annual Manufacturing Revenue to Daily Maximum Temperature

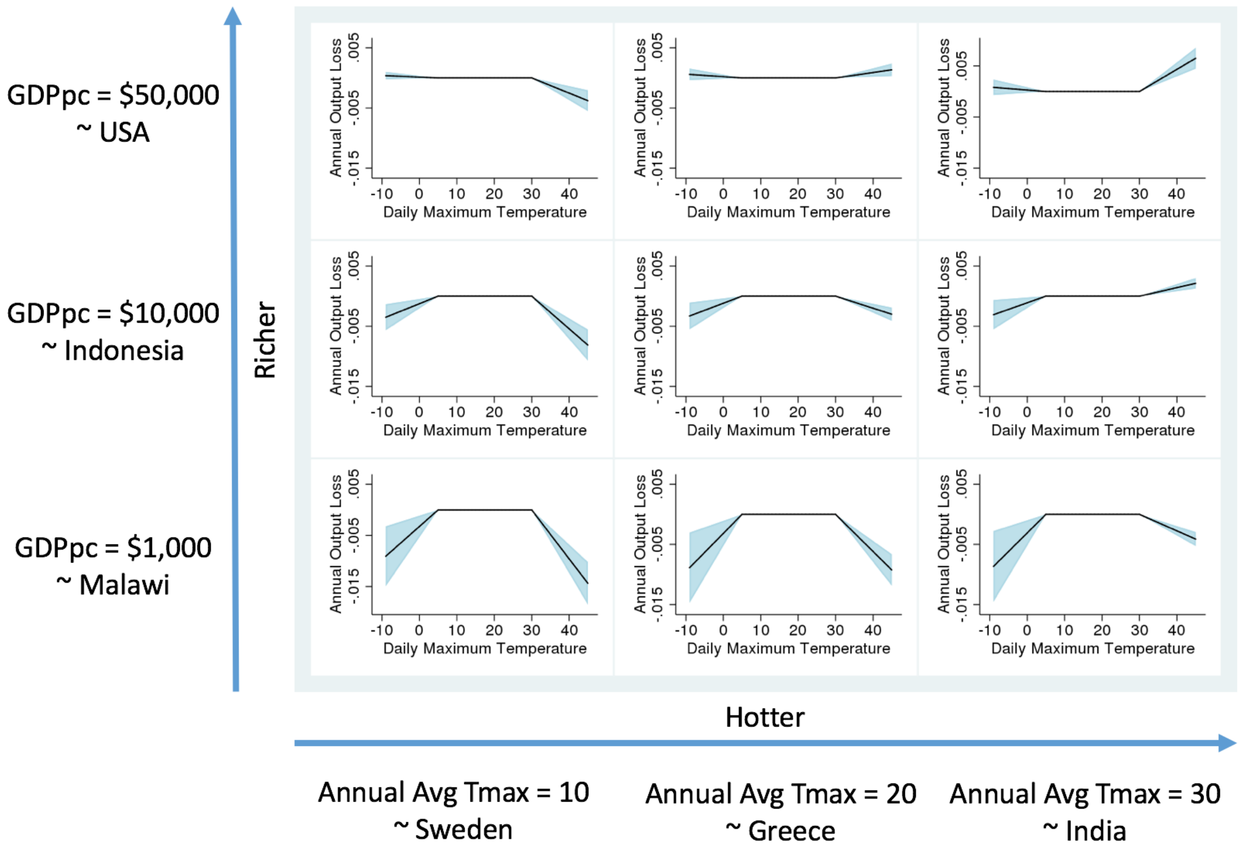

Notes: Figure shows the predicted effect of temperature on the log of manufacturing revenues at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 3 of Table 2. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-2: Predicted Heterogeneous Response of Annual Manufacturing Employment to Daily Maximum Temperature

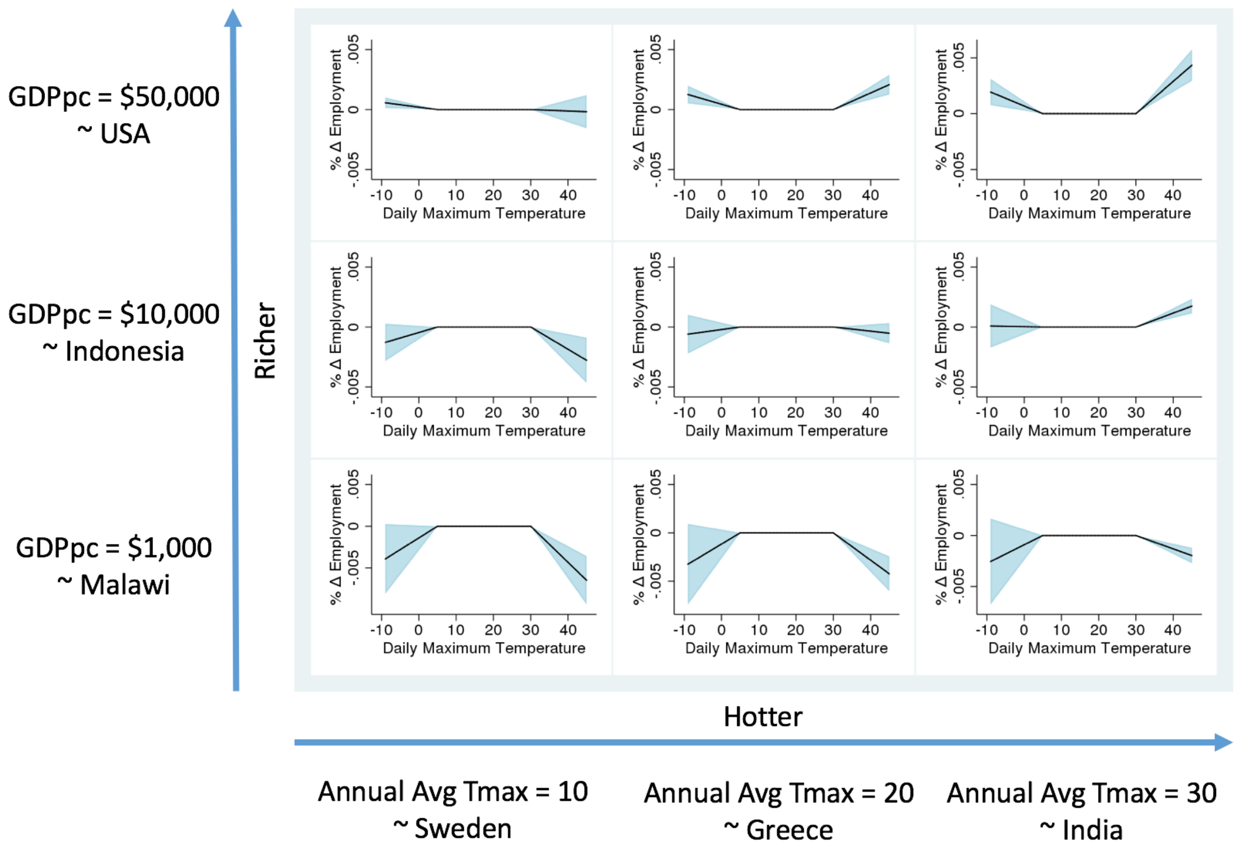

Notes: Figure shows the predicted effect of temperature on the log of manufacturing employment at varying levels of income and long-run average temperature by evaluating the interacted regression from Column 4 of Table 2. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-3: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature

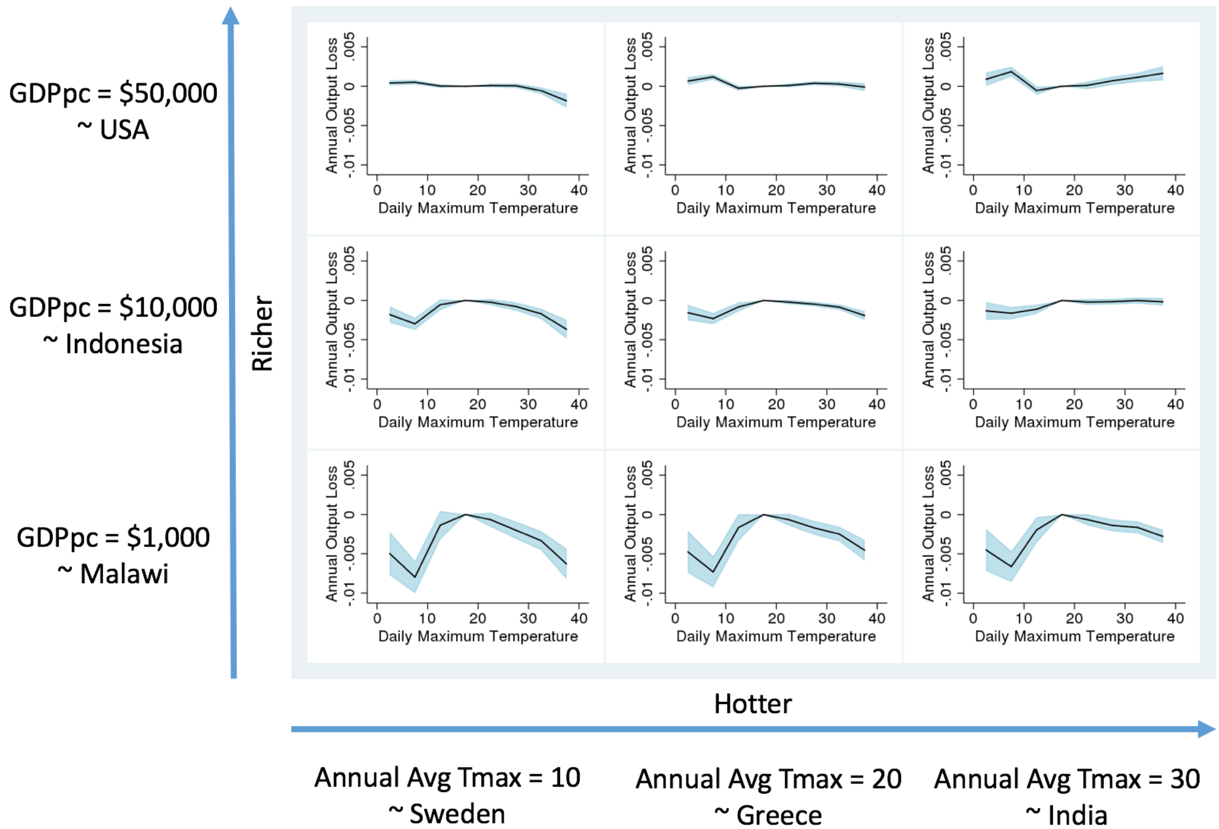

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using bins of daily maximum temperature in the specification from Equation 8.Days are divided into 5°C bins. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-4: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature

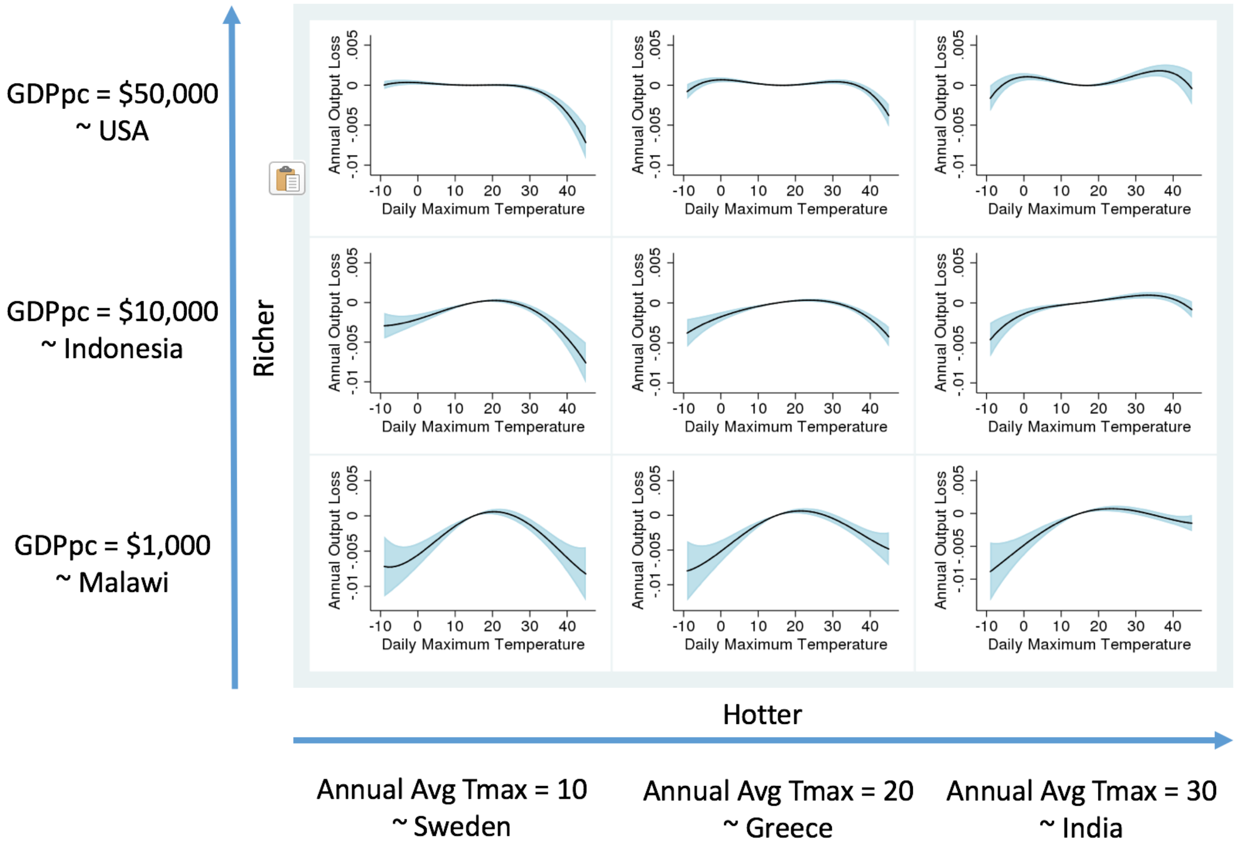

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using a polynomial of degree four in daily average temperature in the specification from Equation 8.Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-5: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE

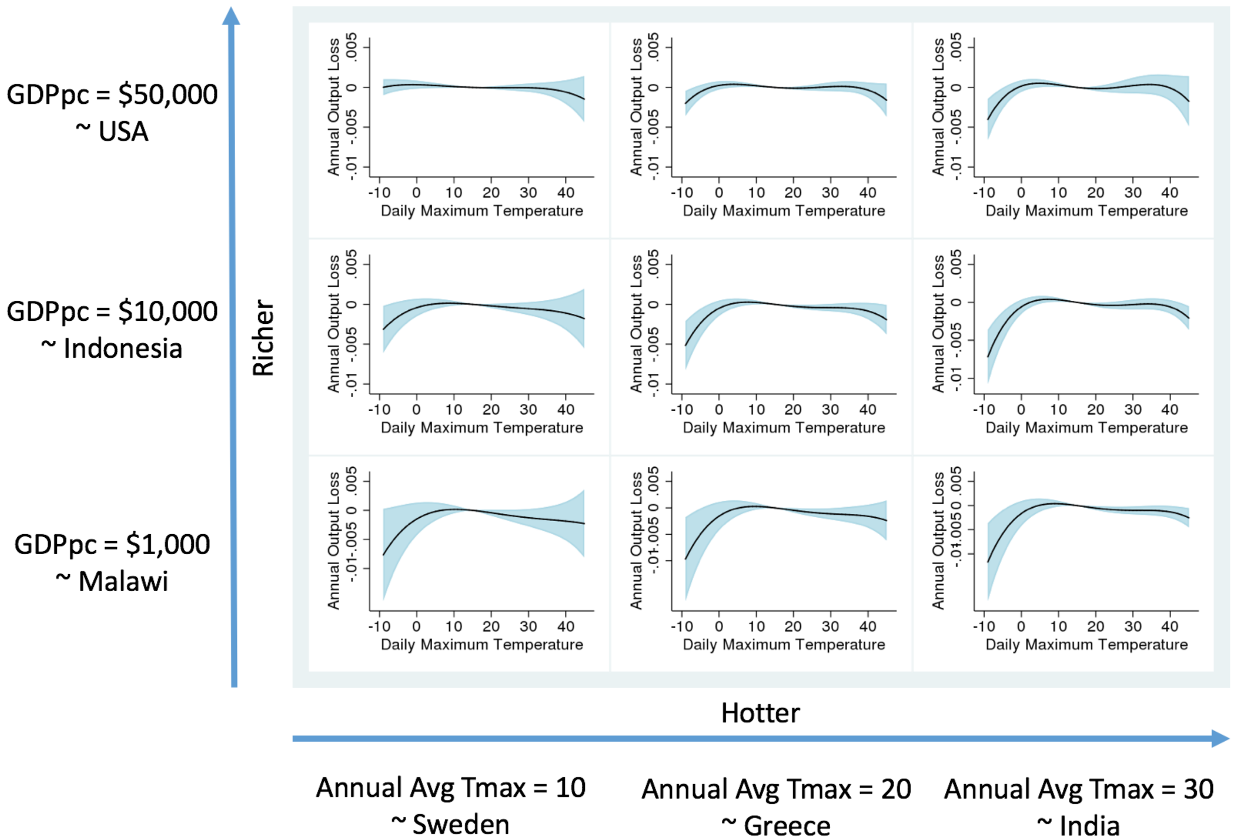

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-6: Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE

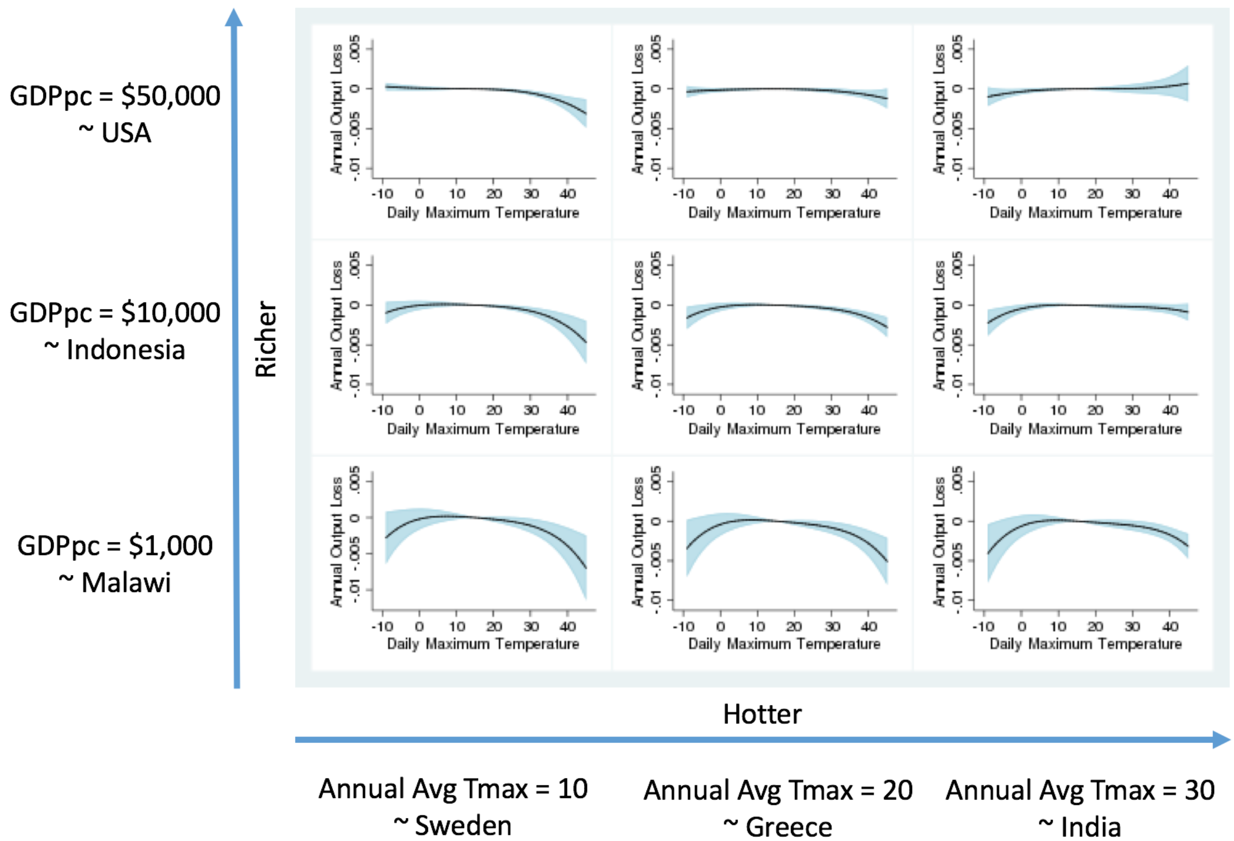

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and a polynomial of degree four in daily maximum temperature. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-7: Predicted Heterogeneous Response of Annual Manufacturing/Services Revenue Per Worker to Daily Maximum Temperature - State-by-Year FE

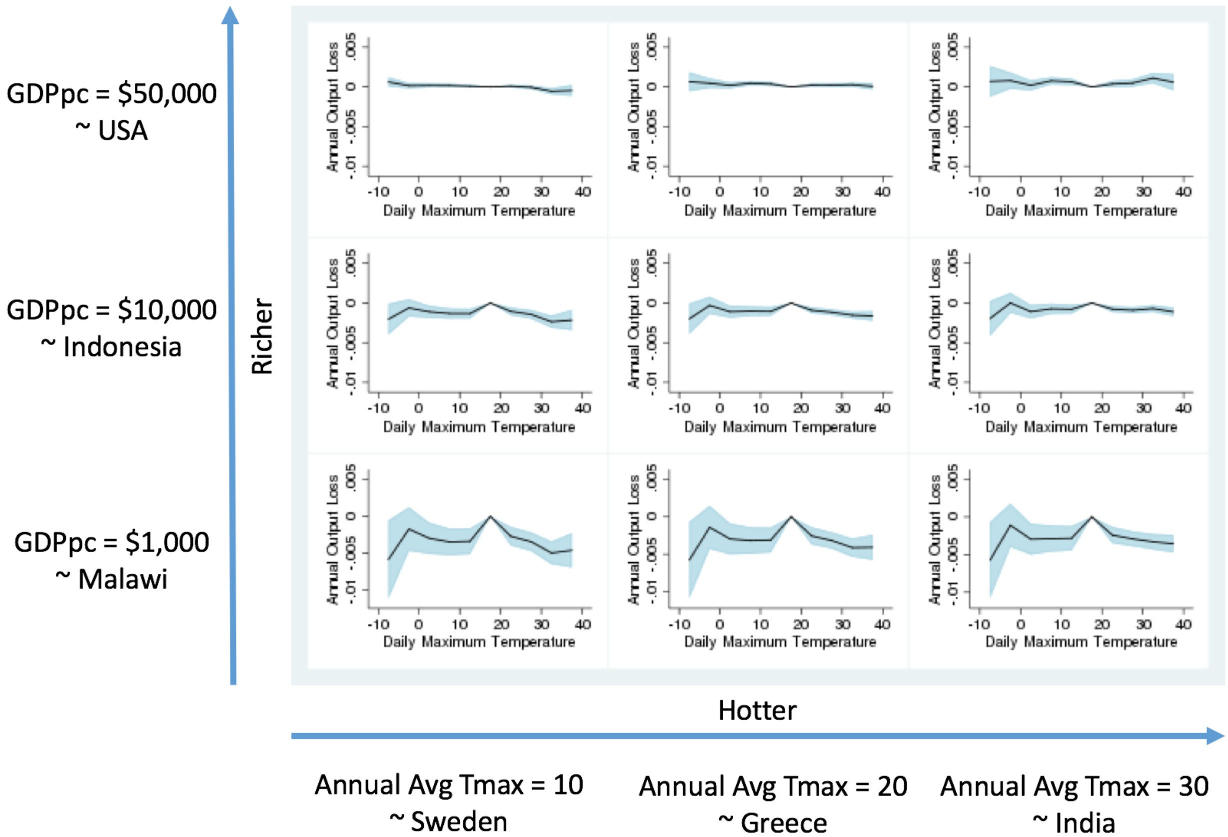

Notes: Figure shows the predicted effect of temperature on revenue per worker at varying levels of income and long-run average temperature for a pooled sample of manufacturing and services firms using the specification from Equation 8 with state-by-year fixed effects and bins of daily maximum temperature. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-8: Predicted Heterogeneous Response of Annual Manufacturing Revenue Per Worker to Daily Maximum Temperature - Controls for Capital

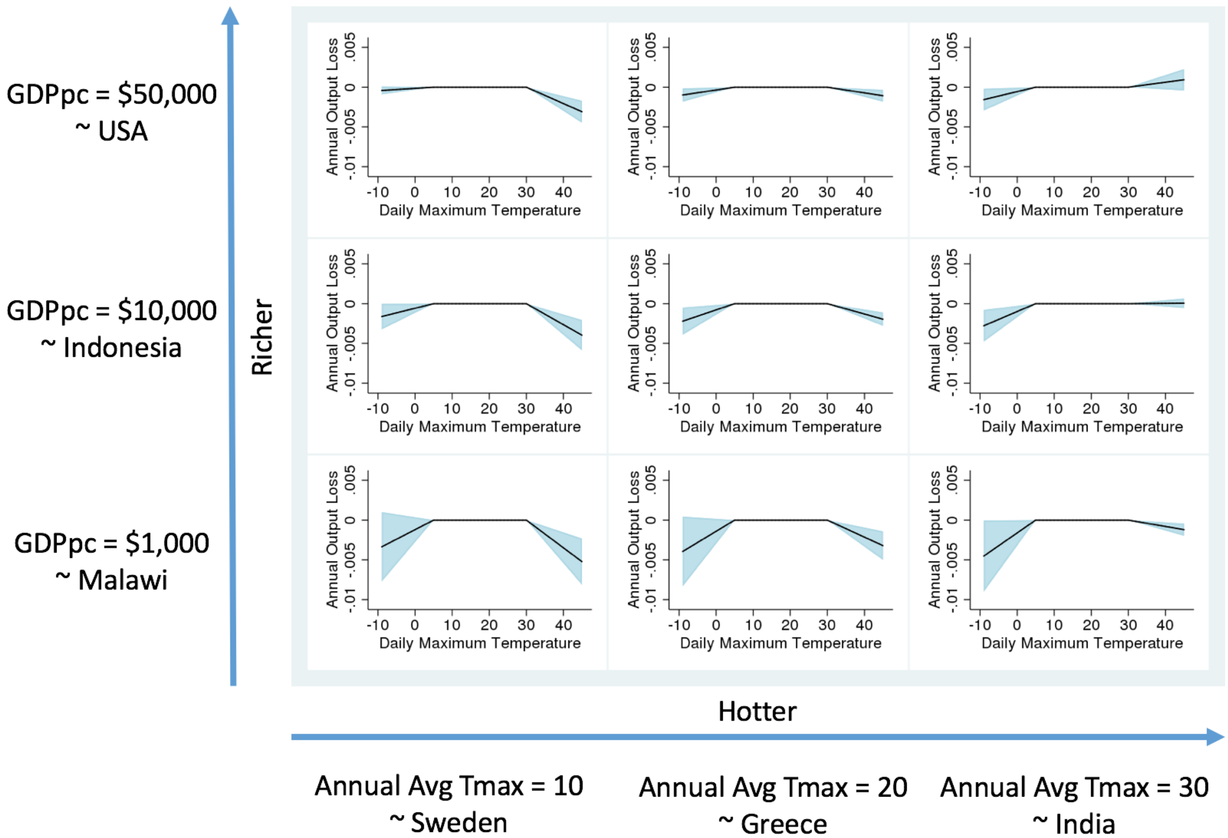

Notes: Figure shows the predicted effect of temperature on manufacturing revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with controls for capital. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-9: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature

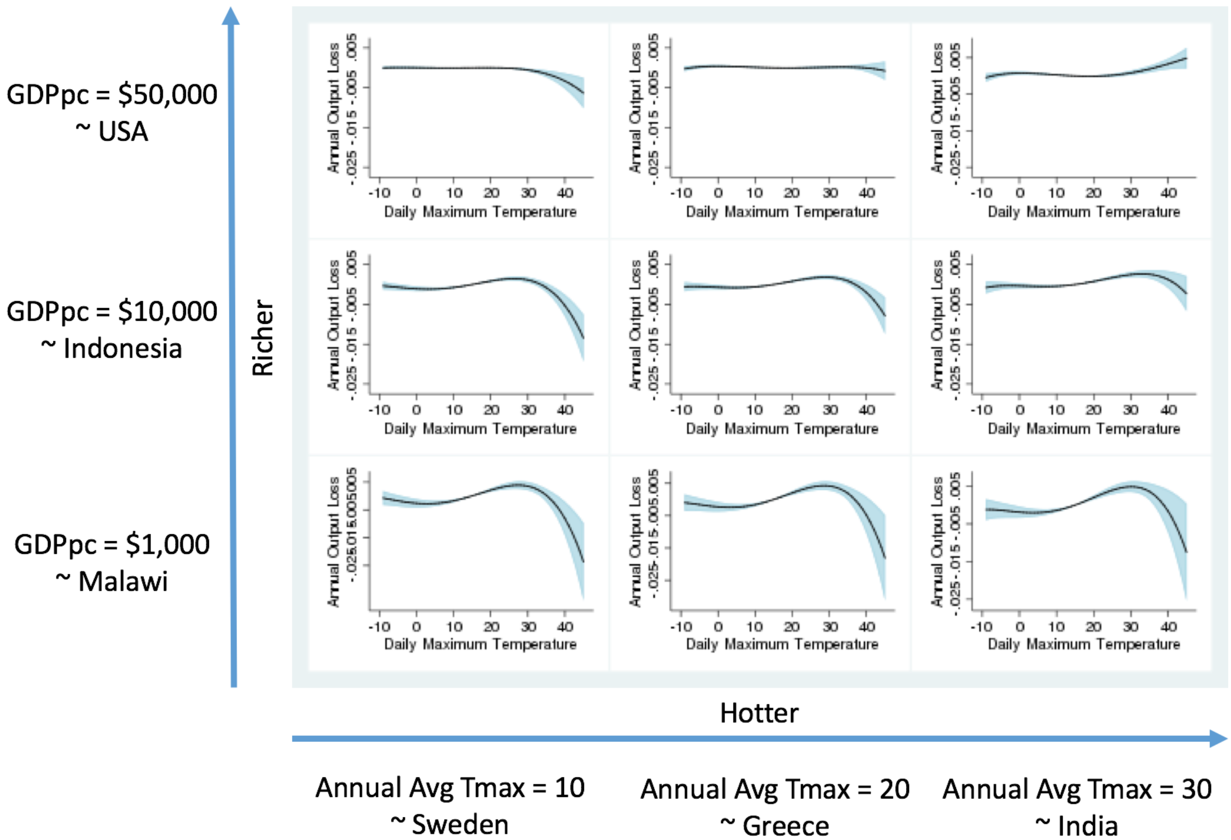

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with a polynomial of degree four in daily maximum temperature. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-10: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature

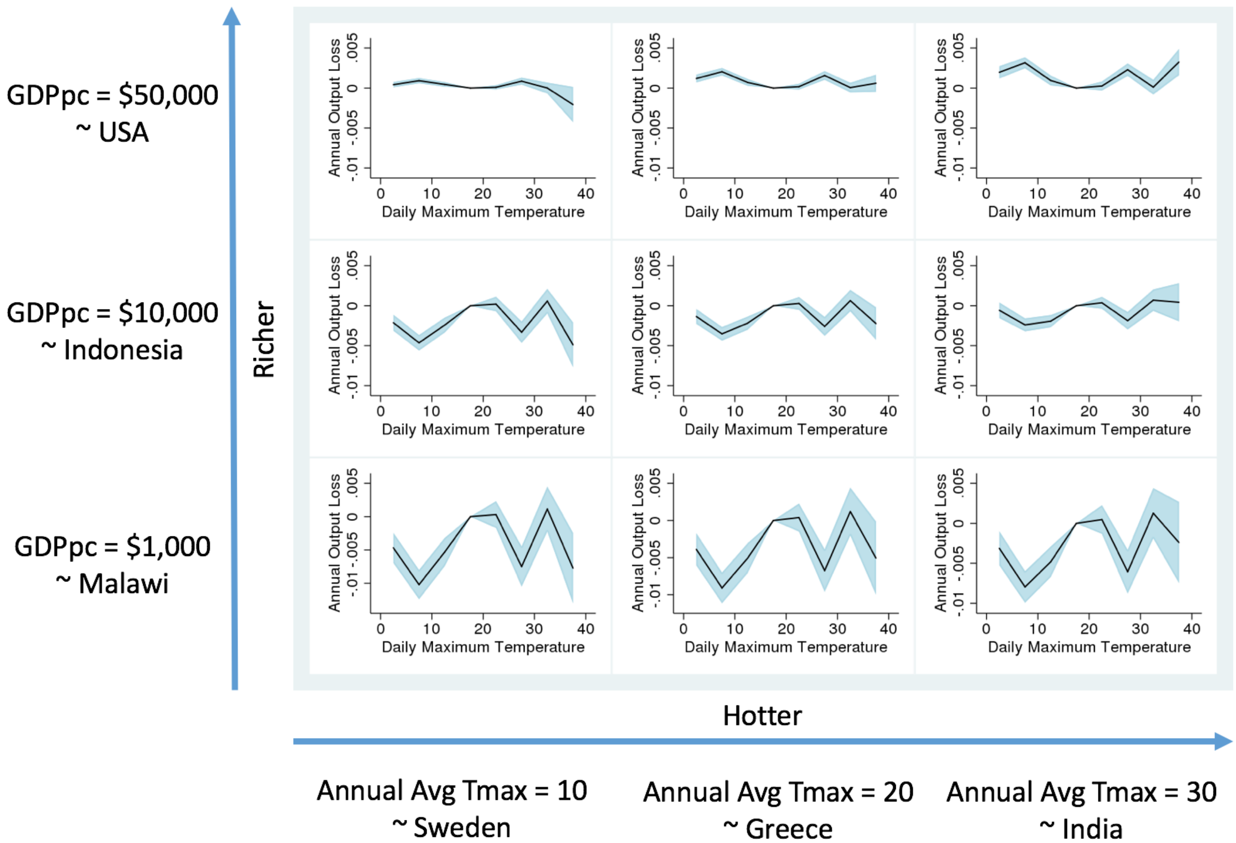

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8 with bins of daily maximum temperature. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Figure A-11: Predicted Heterogeneous Response of Annual Services Revenue Per Worker to Daily Maximum Temperature

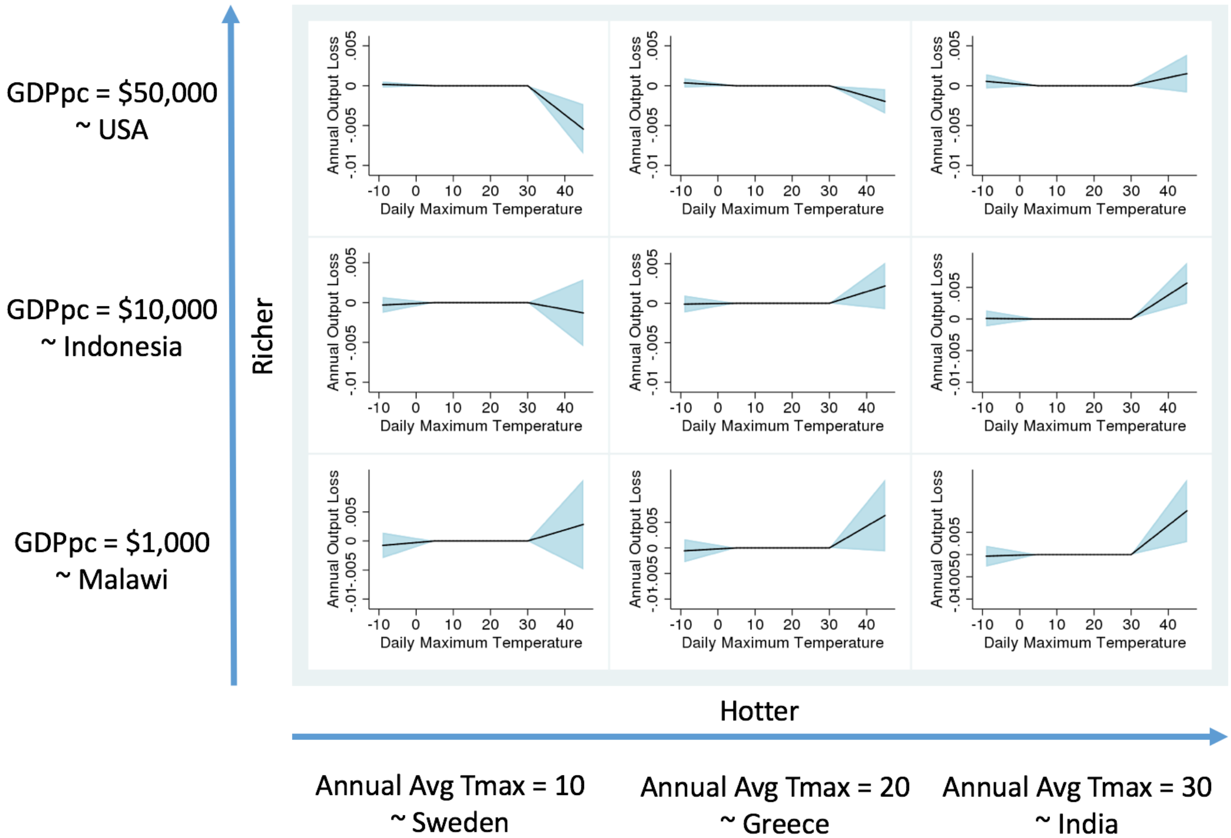

Notes: Figure shows the predicted effect of temperature on services revenue per worker at varying levels of income and long-run average temperature using the specification from Equation 8. Outcome variables come from data sources listed in Table 1 and temperature data is from GMFD.

Appendix B: U.S. Results

Figure A-12: Estimated Response of U.S. Annual Manufacturing TFPR to Daily Maximum Temperature

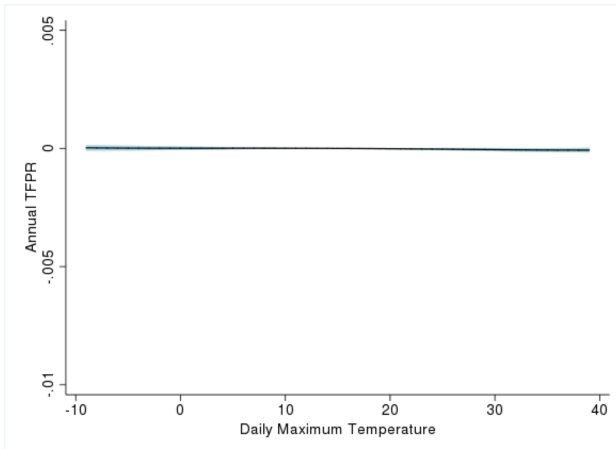

Notes: Figure shows the estimated effect of temperature on manufacturing TFPR using the specification from Equation 7 with a polynomial of degree four in daily maximum temperature. Outcome data comes from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD.

Table A-1: U.S. Results

	Revenue/Worker	Revenue	Employment	TFPR	Revenue/Worker	Revenue/Worker
TMax-30	-0.0000109	0.0000220	0.0000330	0.00000134	-0.0000422	0.0000110
	(-2.21)	(2.01)	(3.49)	(0.33)	(-2.97)	(0.46)
5-TMax	0.0000365	0.0000338	-0.00000269	-0.00000685	-0.0000226	0.000154
	(5.65)	(2.65)	(-0.26)	(-1.30)	(-1.71)	(3.56)
Observations	2852000	2852000 2852000	2852000	2852000	2852000	
Firm FE	X	X	X	X	X	X
Country X Year FE	X	X	X	X	X	X
State X Year FE					X	
Sales Weighting						X

Notes t-statistics in parentheses. Dependent variables all in logs. Standard errors two-way clustered at the firm and county-by-year level. Estimates use the regression model from Equation 7 with outcome variable data from 1976-2014 from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau and temperature data from the Global Meteorological Forcing Dataset.

Figure A-13: Estimated Response of U.S. Annual Manufacturing Plant-Level Energy Expenditures to Daily Maximum Temperature

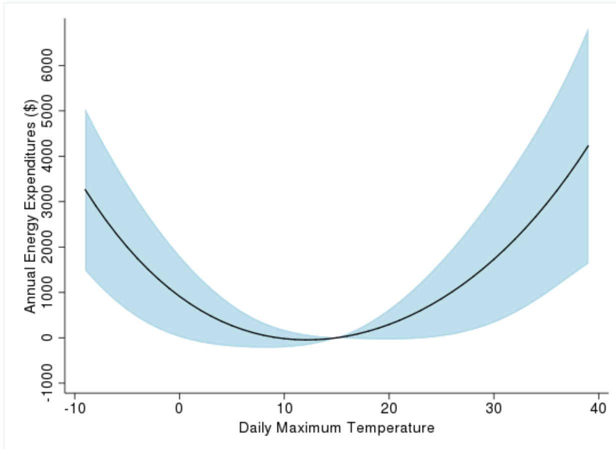

Notes: Figure shows the estimated effect of temperature on manufacturing energy expenditures using the specification from Equation 7 with a polynomial of degree four in daily maximum temperature. Energy expenditures are the sum of cost of fuels and electricity expenditures in the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau. Temperature data is from GMFD.

Table A-2: U.S. Energy Results

	log(Energy Expenditure)	Energy Expenditures	log(Energy Expenditures)	Energy Expenditures
TMax-30	0.0000822	0.0000890	251.1	6056
	(6.03)	(3.24)	(4.45)	(1.32)
5-TMax	0.0000108	0.00000184	490.8	13840
	(0.78)	(0.04)	(3.57)	(1.69)
Observations	2852000	2852000	2852000	2852000
Firm FE	X	X	X	X
Country X Year FE	X	X	X	X
Sales Weighting		X		X

Notes: t-statistics in parentheses. Standard errors two-way clustered at the firm and county-by-year level. Estimates use the regression model from Equation 7 with outcome variable data from 1976-2014 from the Annual Survey of Manufacturers and Census of Manufacturers from the U.S. Census Bureau and temperature data from the Global Meteorological Forcing Dataset. Dependent variable is the sum of electricity expenditures and cost of fuels, in logs or levels.

Appendix C: China Results

This section explains the data quality issues that lead me to estimate the results in Section 4.1 excluding data from China. At a high level, I find evidence consistent with the conclusions of Chen, Chen, Hsieh and Song (2019) that Chinese micro-data after 2007 are unreliable due to systematic manipulation by local officials. The details are as follows.

To start with, Zhang, Deschenes, Meng and Zhang (2018) analyze data from China for the years 1998-2007 and find that both cold and hot temperatures harm output and productivity, consistent with my findings. Using the overlapping subset of years from my data, which goes from 2003-2012, I am able to replicate their findings fairly closely, as shown in Appendix Figure A-14. Notably, I am also able to use my main results from the rest of my global data in Figure 3 to predict the response of output to temperature in China based on their income level and average climate. My prediction and the estimates from Zhang, Deschenes, Meng and Zhang (2018) are shown in Figure A-15. While I slightly overpredict sensitivity to cold and underpredict sensitivity to heat, my results are broadly consistent with their findings, lending external validity to my work.

However, when I estimate the response to temperature in my full sample of Chinese firms from 2003-2012, I produce the highly anomalous results shown in Figure A-16. This estimate using my full sample of Chinese data implies that extreme temperatures sharply and statistically significantly *increase* output, a finding inconsistent with my results from any other country in the world. Notably, this anomalous result begins to appear by including later years starting with 2008 in the regression, the same year Chen, Chen, Hsieh and Song (2019) start to find discrepancies in the data. They state that “local statistics increasingly misrepresent the true numbers after 2008” and “the micro-data of the ASIF [have] overstated aggregate output.”

A somewhat puzzling fact is that my results suggest that this documented manipulation of data in China is systematically correlated with temperature. One plausible hypothesis is that Chinese provincial officials inflate reported manufacturing output to meet GDP targets in response to declines in other sectors more susceptible to temperature, such as agriculture. These targets have historically played a central role in the evaluation and promotion of government officials, and Lyu, Wang, Zhang and Zhang (2018) demonstrate that reported provincial GDP just barely hits target thresholds with implausible frequency. I cannot provide further evidence on the particular sources and methods of manipulation, but given the widespread external documentation of problems with this subset of the Chinese firm data and my very short panel that would remain when excluding these years in China, I exclude this dataset entirely from my main analysis. Still, I view the consistency of both my replication and predictions with the results of Zhang, Deschenes, Meng and Zhang (2018) as validating my central analysis.

Figure A-14: China Replication - Overlapping Years

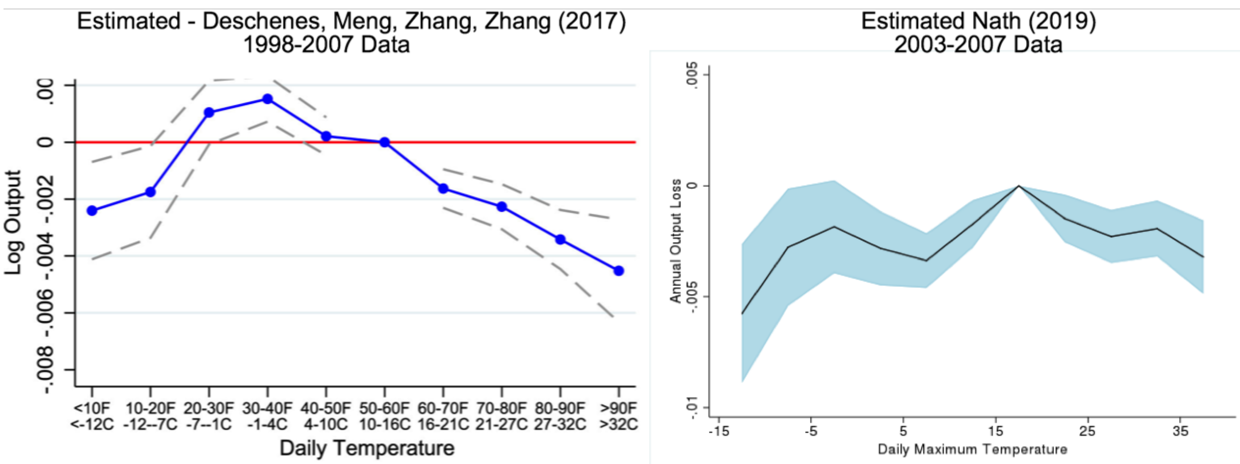

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result using data from the same dataset for 2003-2007 - the overlapping years of my data coverage. Temperature data is from GMFD.

Figure A-15: China Manufacturing Temperature Sensitivity
- Estimated and Predicted

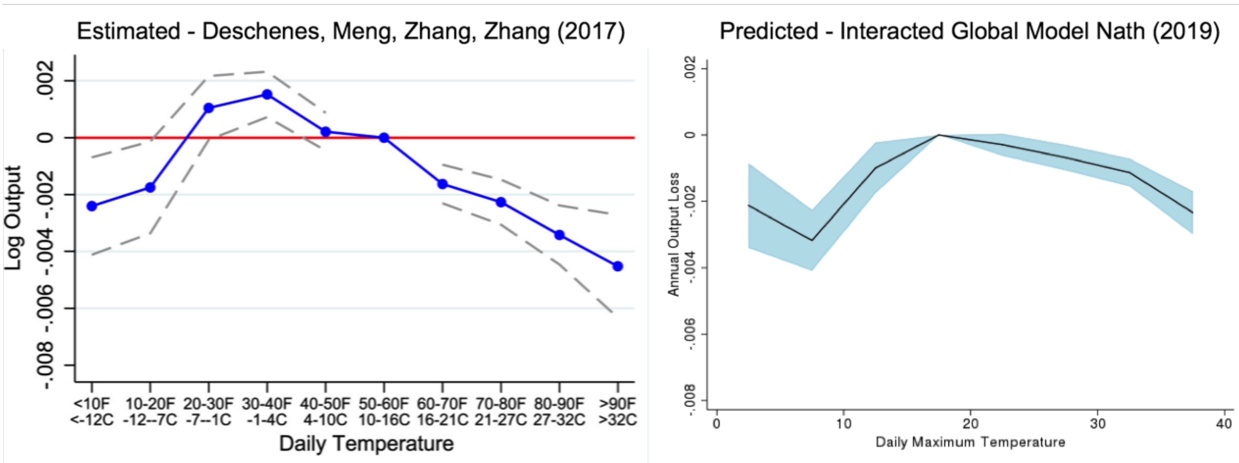

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows the predicted effect of temperature in China from evaluating my global interacted specification from Column 2 of Table 2 at China's income and average long-run temperature from 1998-2007. I do not use any data from China in my estimation or prediction.

Figure A-16: China Replication - Different Years

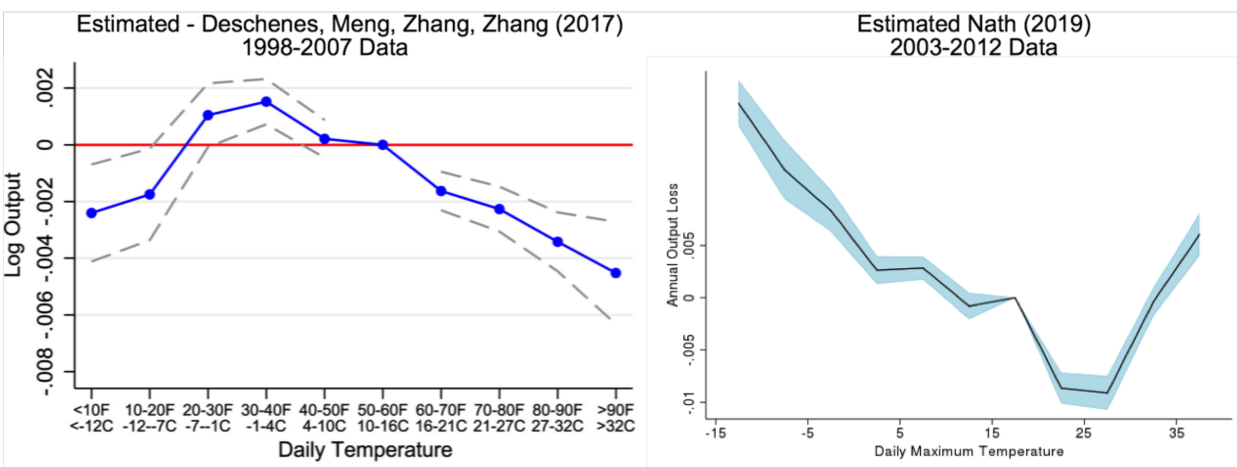

Notes: Left panel of the figure shows the effect of temperature on annual manufacturing output in China estimated by Zhang, Deschenes, Meng and Zhang (2018) using data from the Chinese Industrial Survey of the National Bureau of Statistics from 1998-2007. Right panel shows my replication of their result using data from the same dataset for 2003-2012 - the years of my data coverage. Temperature data is from GMFD.

Appendix D: Adaptation Benefits and Costs

In this section I explain how I use revealed preference methods developed by Carleton et al. (2018) to infer the costs firms incur from reducing the sensitivity of their production to extreme temperatures. To build intuition start by considering a simple example of otherwise identical firms in two cities, Seattle and Houston. Houston is hotter than Seattle, but Seattle heats up over the course of the century such that its exposure to CDD in 2100 is that of Houston in 2020. Let β represent lost annual revenues from exposure to a cooling degree day, a function of the adaptation investments the firm chooses to make. The annual costs of extreme heat to a firm in Seattle are given by $CDD_{Seattle} * \beta_{Seattle}$. Since Seattle suffers little exposure to extreme heat, its firms choose a lower (more negative) β than firms in Houston, as I find in my empirical estimates. If Seattle firms had chosen the Houston β associated with greater expected exposure to heat, the marginal benefits they would obtain are as follows:

$$MB = CDD_{Seattle} * (\beta_{Houston} - \beta_{Seattle})$$

Given that Seattle firms do not choose $\beta_{Houston}$, we know that the marginal costs of this incremental reduction in temperature sensitivity must exceed the marginal benefits. By repeating this logic for the firm's estimated temperature sensitivity for every year of warming from $Seattle_{2020}$ to $Seattle_{2100}$, we can construct the full marginal cost curve for the Seattle firm's projected change in chosen β from 2020 to 2100:

$$TC = \sum_{t=2020}^{2099} MC_t = \sum_{t=2020}^{2099} CDD_t * (\beta_{t+1} - \beta_t) \tag{46}$$

Note that the continuous version of Equation 46 also follows straight from the firm's first-order condition in the framework in Section 3.1. The firm's lost revenues from extreme heat are $CDD * \beta$ so the marginal benefit the firm receives from a reduction in β is given by CDD. Since the firm's optimal choice of β equates marginal benefit to marginal cost, we have marginal cost $c_\beta = CDD$ for the full range of CDDs.

The total benefits of future adaptation for firms in Seattle are given by the change in damages from choosing their optimal level of adaptation for expected heat exposure in 2100 rather than remaining at the adaptation level they choose in 2020:

$$TB = CDD_{2100} * (\beta_{2100} - \beta_{2020}) \tag{47}$$

Because CDDs are increasing as countries become hotter, the benefits of adaptation in Equation 47 exceed the costs in Equation 46. Figure A-17 shows predicted manufacturing sensitivity to a hot day at end-of-century temperatures, which is substantially muted relative to the sensitivities at current temperatures shown in Figure 5. Figure A-18 show the costs of achieving this reduced sensitivity, as calculated using Equation 46, and Figure A-19 show the net benefits of firms adapting to changes in expected exposure to extreme heat.

Figure A-17: Predicted Effect of a 40°C Day on Annual Manufacturing Revenue per Worker At 2080 Average Temperatures

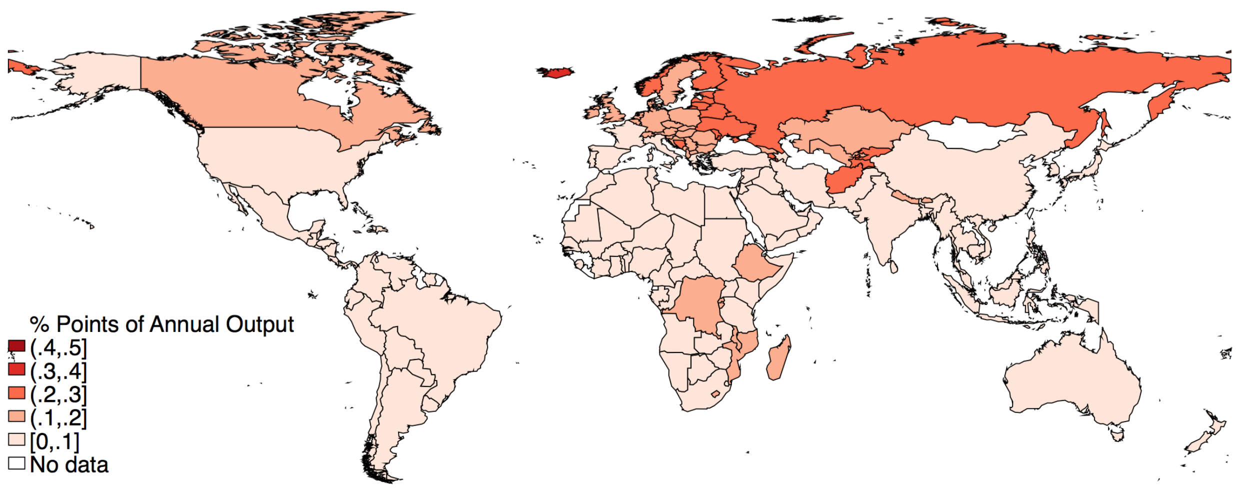

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression in Column 2 of Table 2 at each country's level of income and end-of-century long-run average temperature. Temperature sensitivities are lower in this figure than in Figure 5 because my results predict that firms will adapt to hot temperatures as the world warms.

Figure A-18: Firm-Level Adaptation Costs
(Share of Manufacturing Output)

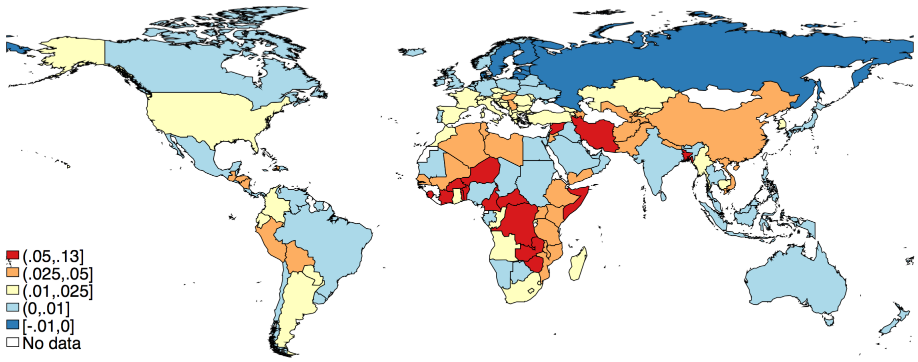

Notes: Map shows my calculations of the costs firms pay to achieve the lower temperature sensitivity shown in Appendix Figure A-17 compared to Figure 5. I infer these costs using a revealed preference approach developed by Carleton et al. (2018) that infers adaptation costs from the foregone benefits firms would have attained by reducing their heat sensitivity. The procedure is detailed in Appendix D.

Figure A-19: Firm-Level Adaptation Net Benefits
(Share of Manufacturing Output)

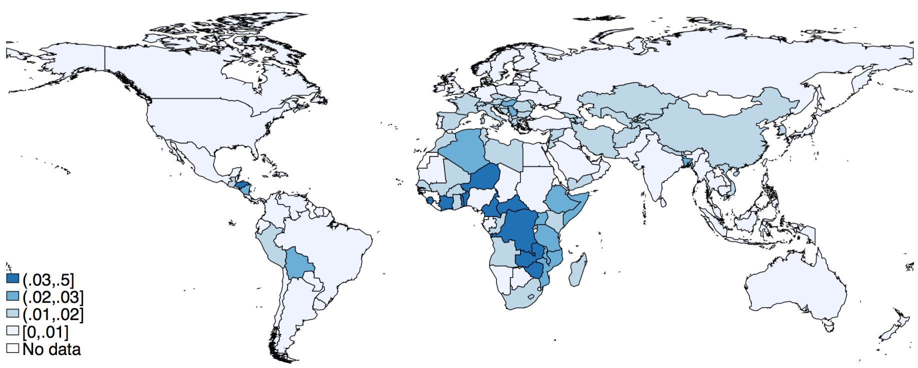

Notes: Map shows my calculations of the net benefits firms receive by investing to reduce their heat sensitivity as the climate warms. The benefits come from reducing heat sensitivity to the level shown in Appendix Figure A-17 compared to the original level in Figure 5. The inferred costs are shown in Appendix Figure A-18. The procedure to calculate these costs and benefits is detailed in Appendix D.

Figure A-20: Predicted Effect of a 40°C Day on Annual Services Revenue per Worker

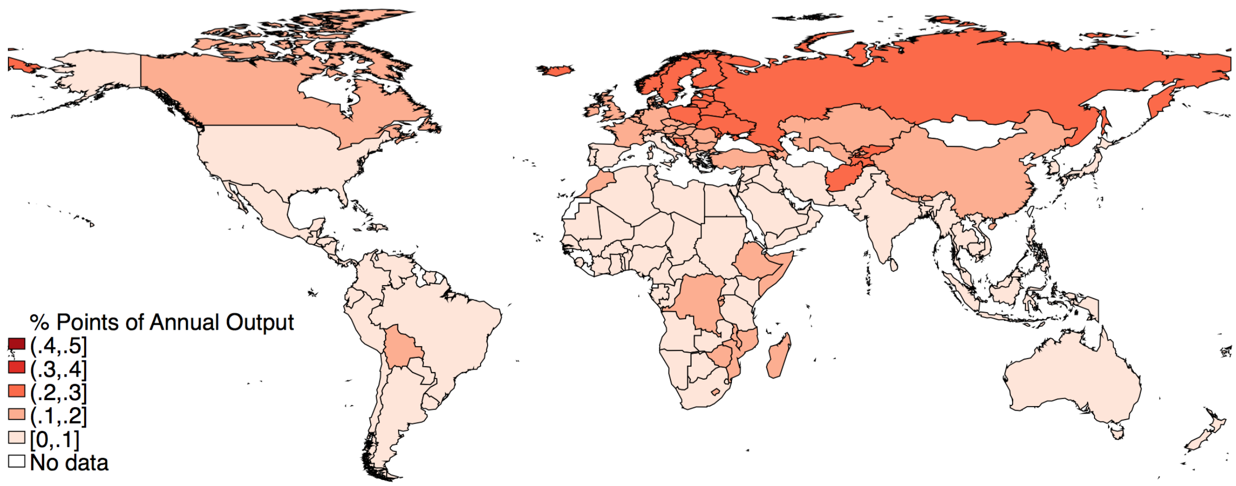

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a 40°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-21: Predicted Effect of a -5°C Day on Annual Services Revenue per Worker

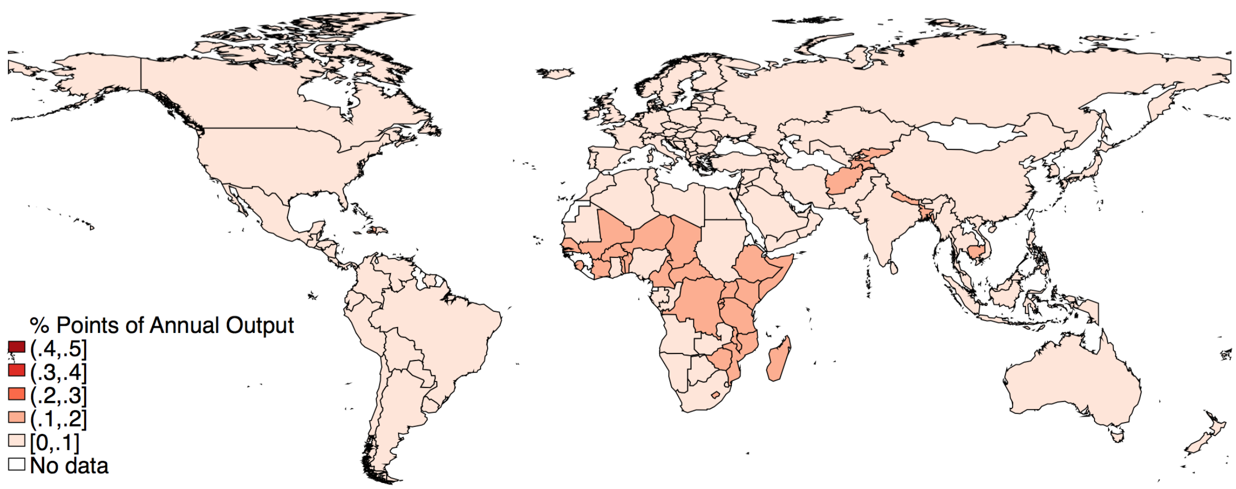

Notes: Map shows the predicted annual percentage point loss in revenue per worker from a -5°C day obtained by evaluating the interaction regression for a pooled sample of manufacturing and services firms in Column 5 of Table 2 at each country's level of income and long-run average temperature.

Figure A-22: Projected Change in Exposure to Extreme Heat

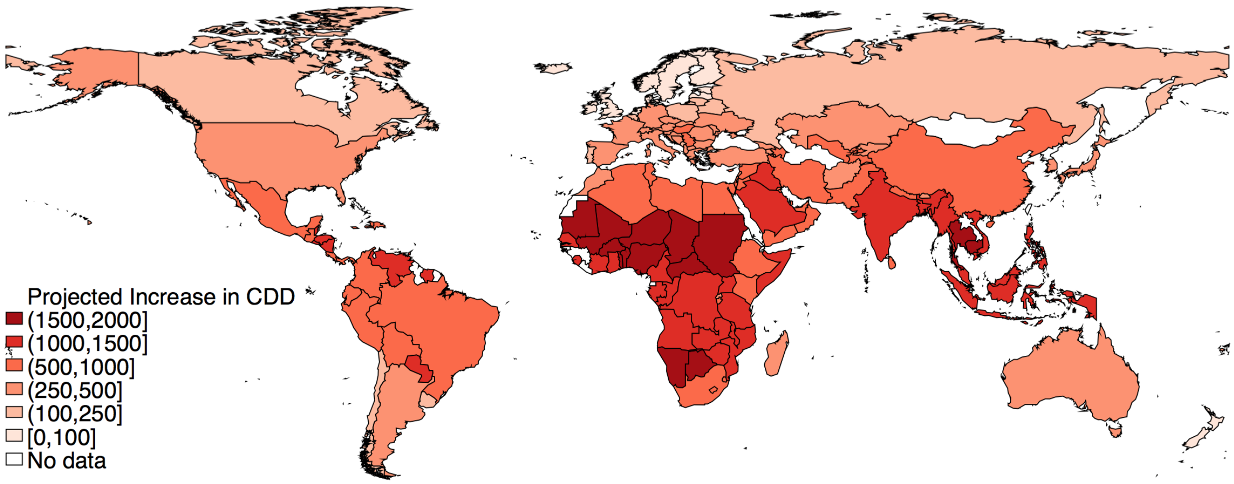

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of changes in exposure to extreme heat as measured by cooling degree days above 30°C between 2015 and the two decade average from 2080 to 2099.

Figure A-23: Projected Change in Exposure to Extreme Cold

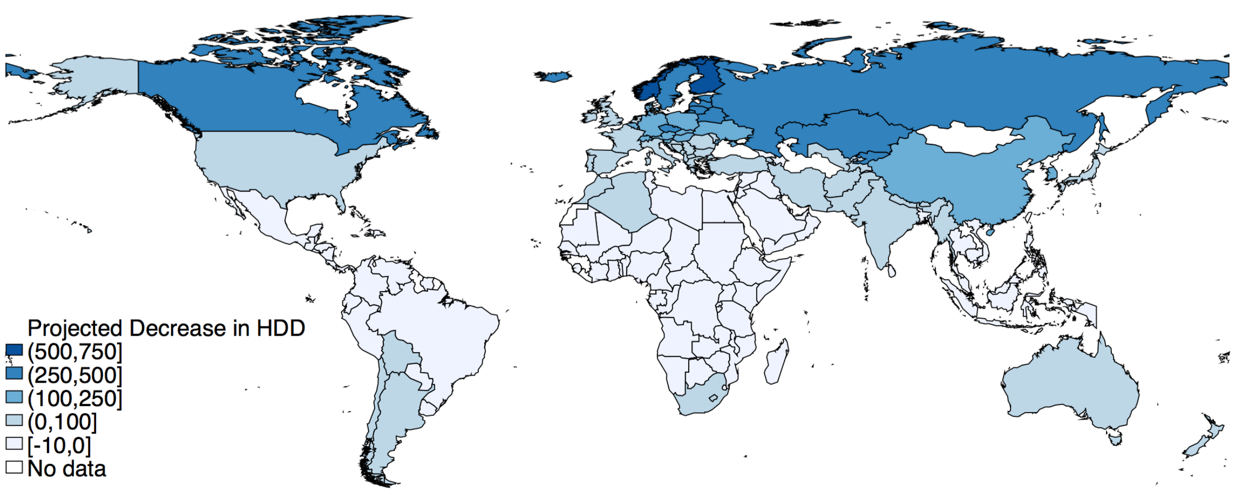

Notes: Map shows projections from the CSIRO-MK-3.6.0 global climate model of changes in exposure to extreme cold as measured by heating degree days below 5°C between 2015 and the two decade average from 2080 to 2099.

Figure A-24: Projected Impact of Climate Change on Services Productivity

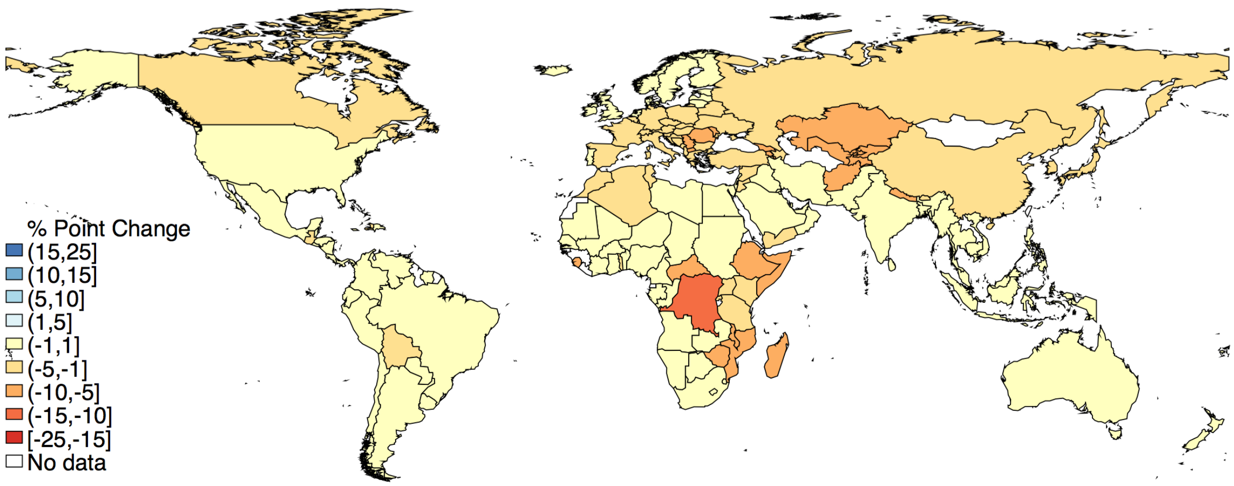

Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 5 of Table 2 at each country's income and end-of-century long-run average temperature.

Figure A-25: Projected Impact of Climate Change on Manufacturing Productivity Accounting for Economic Growth and Adaptation

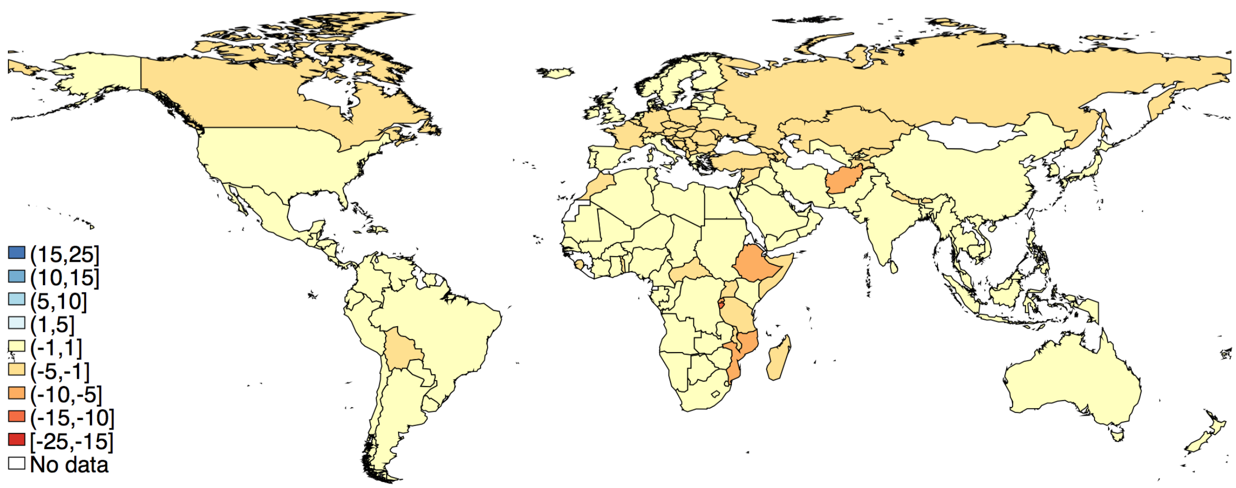

Notes: Map shows the projected impact of climate change on manufacturing productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 2 of Table 2 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017). These estimates that account for economic growth show reduced losses relative to those in Figure 9 because my empirical results suggest that firms in richer countries have reduced exposure to extreme temperatures.

Figure A-26: Projected Impact of Climate Change on Services Productivity Accounting for Economic Growth and Adaptation

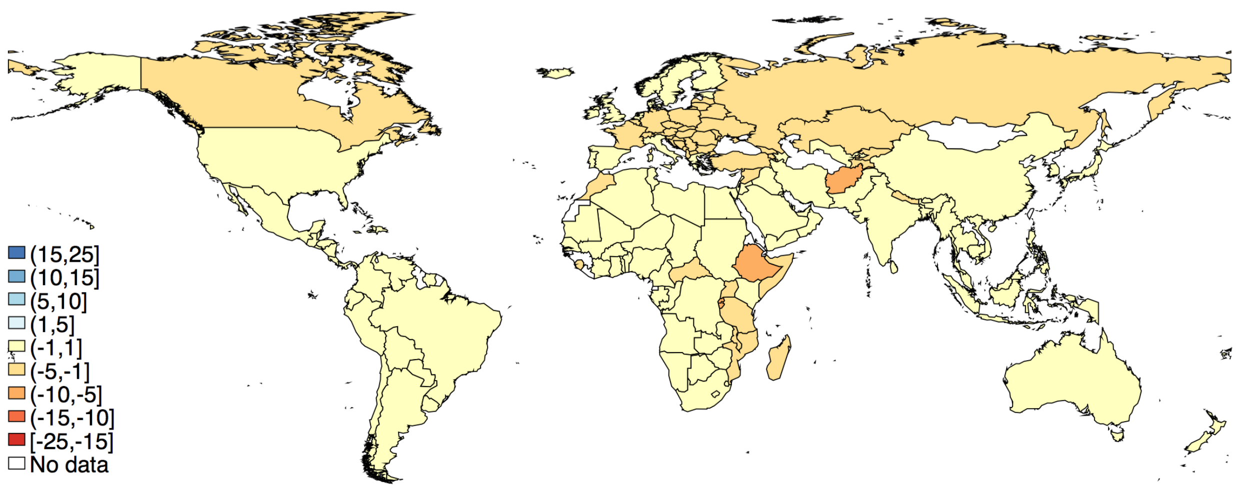

Notes: Map shows the projected impact of climate change on services productivity in 2080-2099 obtained by multiplying predicted temperature sensitivities by CSIRO-MK-3.6.0 global climate model predictions of changes in exposure to extreme heat and cold. Temperature sensitivities are calculated by evaluating the interaction regression from Column 5 of Table 2 at each country's end-of-century long-run average temperature and 2080 income as projected by Cuaresma (2017). These estimates that account for economic growth show reduced losses relative to those in Appendix Figure A-24 because my empirical estimates suggest that firms in richer countries have reduced exposure to extreme temperatures.

Appendix E: Additional Model Fit Figures

Figure A-27: Manufacturing Share of GDP - Data vs. Simulation

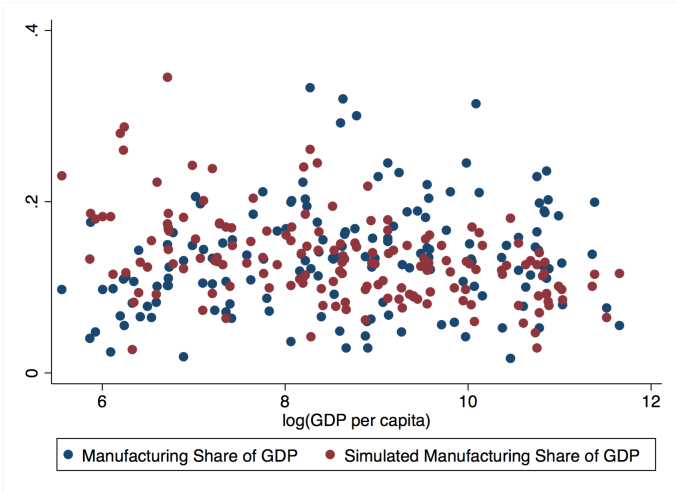

Notes: Graph shows the fit of simulated manufacturing share of GDP in the model to data from the World Bank.

Figure A-28: Services Share of GDP - Data vs. Simulation

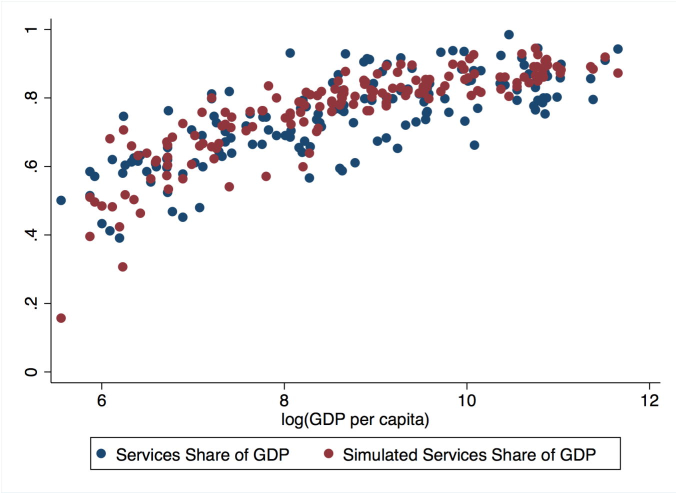

Notes: Graph shows the fit of simulated services share of GDP in the model to data from the World Bank.

Figure A-29: Agriculture Share of GDP - Data vs. Simulation

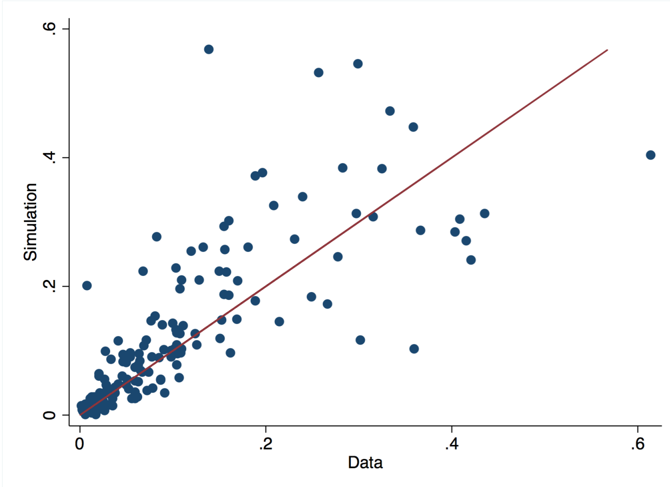

Notes: Graph shows another view of the fit of simulated agriculture share of GDP in the model to data from the World Bank also shown in Figure 7. A perfect fit would have all data points be on the 45°line where the simulated and actual values are equal. The simulation explains over 60% of the variation in the agriculture share of GDP.

Figure A-30: Domestic Production Share of Agriculture Expenditures - Data vs. Simulation

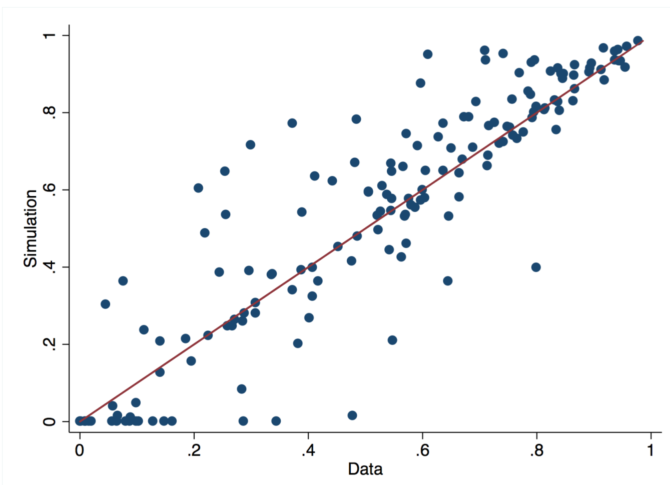

Notes: Graph shows the fit of simulated domestic production share of agricultural consumption in the model to data from Comtrade. As shown in Section 5.2, openness to food imports is a crucial parameter governing the response of labor reallocation to climate change. The simulation explains over 80% of the variation in the data for this moment.

Figure A-31: Manufacturing Domestic Production Share of Expenditures - Data vs. Simulation

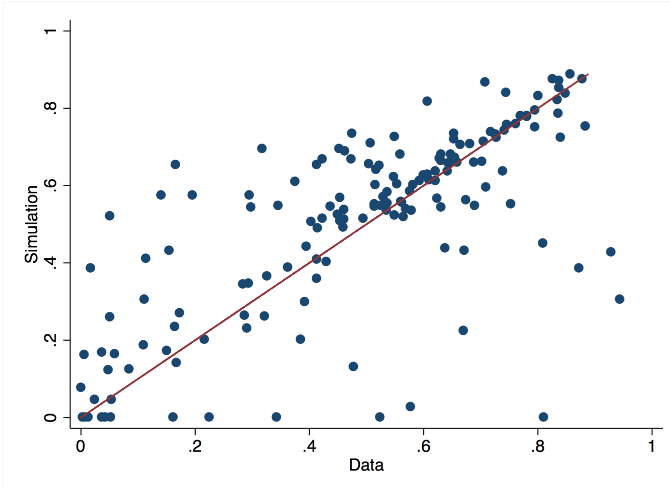

Notes: Graph shows the fit of simulated domestic production share of manufacturing consumption in the model to data from Comtrade.

Figure A-32: Log GDP Per Capita - Data vs. Simulation

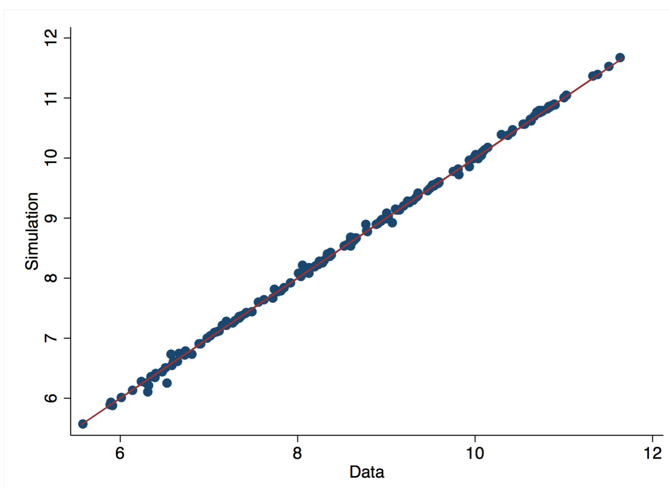

Notes: Graph shows the fit of simulated domestic production share of manufacturing consumption in the model to data from the World Bank.

Figure A-33: Domestic Production Share of Expenditures in Agriculture - Model Simulation

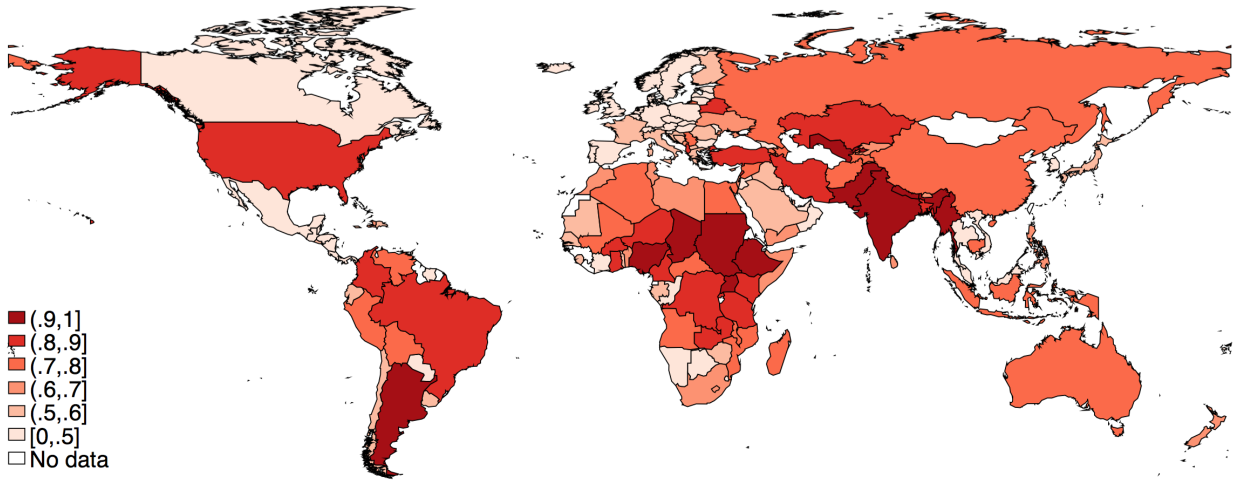

Notes: Figure shows that the share of expenditures on domestically produced goods in agriculture is very high in many developing countries with high barriers to trade. Table 5 shows that these simulated values track closely to the data.

Figure A-34: Agriculture Share of GDP - Data vs. Simulation
Stone-Geary Specification

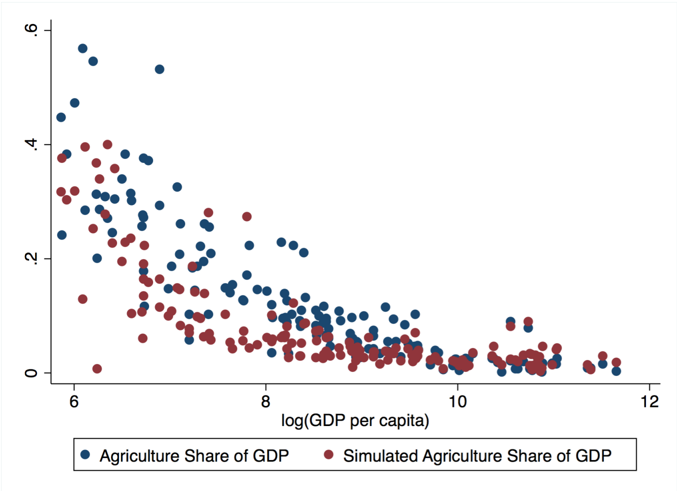

Notes: Graph shows the fit of simulated agriculture share of GDP to data from the World Bank with an alternative model specification using Stone-Geary preferences over sectoral consumption. The best fit with Stone-Geary preferences has an R^2 of only 0.43 and dramatically underpredicts the agriculture share in middle-income countries especially. In contrast, the chosen nonhomothetic CES preferences from Comin, Lashkari and Mestieri (2015) explain over 60% of the variation.

Appendix F: Country-by-Country Model Counterfactual Results

Table A-3: Counterfactual Ag Net Export Share of GDP - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Angola	-.258	0	-.01	-.019
Benin	-.327	0	-.028	-.074
Botswana	-.469	0	-.096	-.176
Burkina Faso	-.243	0	-.003	-.055
Cameroon	-.2	-.005	-.026	-.052
Cape Verde	-.327	-.012	-.107	-.222
Central African Republic	-.601	-.076	-.097	-.385
Chad	-.601	0	.038	-.121
Comoros	-.217	-.075	-.156	-.204
Congo	-.601	0	-.104	-.241
Cote d'Ivoire	-.143	0	.005	0
Democratic Republic of Congo	-.147	-.142	-.045	.002
Ethiopia	-.313	-.102	.052	.026
Gabon	-.601	0	-.056	-.162
Gambia	-.327	0	-.161	-.216
Ghana	-.14	0	.059	.057
Kenya	-.054	-.044	.027	.051
Lesotho	-.469	-.055	-.05	-.172
Madagascar	-.262	-.067	.006	-.034

Table A-4: Counterfactual Ag Net Export Share of GDP - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Malawi	-.313	-.111	-.052	-.041
Mali	-.356	0	-.043	-.106
Mauritania	-.327	0	-.05	-.144
Mauritius	-.262	0	-.002	-.03
Mozambique	-.217	-.104	-.037	-.015
Namibia	-.469	0	-.005	-.097
Niger	-.341	0	-.041	-.174
Nigeria	-.185	0	.006	.005
Rwanda	-.601	-.058	0	-.289
Senegal	-.519	0	-.025	-.182
Seychelles	-.262	0	-.029	-.03
Sierra Leone	-.327	-.071	-.086	-.053
Somalia	-.166	-.125	-.048	-.04
South Africa	-.334	0	.006	-.011
Sudan	-.561	0	.022	-.066
Swaziland	-.469	-.006	-.132	-.258
Tanzania	-.242	-.057	.001	-.03
Togo	-.327	-.042	.13	.075
Uganda	-.168	-.057	.032	.009
Zambia	-.396	0	-.023	-.101
Zimbabwe	-.379	-.099	-.06	-.15

Table A-5: Counterfactual Ag Net Export Share of GDP - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Algeria	-.36	-.019	-.004	-.024
Bahrain	-.219	0	-.015	-.014
Egypt	.113	0	-.002	.024
Iran	-.289	.002	-.001	-.009
Iraq	-.411	0	-.018	-.057
Jordan	-.27	-.019	-.009	-.009
Kuwait	-.219	0	-.022	-.028
Lebanon	-.27	-.012	.061	.047
Libya	-.124	0	.006	-.001
Morocco	-.39	-.038	0	-.054
Oman	-.219	0	-.005	-.012
Qatar	-.219	0	-.003	-.004
Saudi Arabia	-.219	0	-.008	-.013
Syria	-.27	-.049	-.013	-.038
Tunisia	-.36	-.019	-.003	-.038
United Arab Emirates	-.219	0	.003	.005
Yemen	-.282	-.031	-.019	-.037

Table A-6: Counterfactual Ag Net Export Share of GDP - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Afghanistan	-.247	-.086	-.001	-.011
Azerbaijan	-.226	-.044	-.021	-.032
Bangladesh	-.217	-.016	-.03	-.044
Bhutan	-.381	-.034	-.112	-.243
Brunei	-.179	0	.006	-.002
Cambodia	-.271	0	-.022	-.066
China	-.072	-.036	-.021	-.017
Hong Kong	-.072	0	-.002	0
India	-.381	0	0	-.026
Japan	-.057	-.01	-.01	-.01
Kazakhstan	.114	-.031	-.005	.005
Kyrgyzstan	-.059	-.039	-.019	-.025
Maldives	-.201	0	-.033	-.024
Myanmar	-.393	0	.012	-.007
Nepal	-.173	-.07	0	.01
Pakistan	-.304	.001	.016	0
Philippines	-.234	0	-.033	-.059
South Korea	-.093	-.015	-.024	-.025
Sri Lanka	-.201	0	-.01	-.002
Tajikistan	-.059	-.097	-.053	-.06
Thailand	-.262	0	-.041	-.071
Uzbekistan	-.121	-.065	.028	.053
Vietnam	-.151	0	-.014	-.037

Table A-7: Counterfactual Ag Net Export Share of GDP - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Argentina	-.111	0	.05	.053
Barbados	-.237	0	.013	.026
Bolivia	-.43	-.042	-.006	-.064
Brazil	-.169	0	.011	.009
Chile	-.244	-.009	.005	-.012
Colombia	-.232	0	.009	.002
Ecuador	-.288	0	.019	-.001
Honduras	-.237	-.006	-.062	-.1
Paraguay	-.43	0	.061	-.049
Peru	-.306	-.005	.006	-.024
Suriname	-.43	0	-.009	-.036
Trinidad and Tobago	-.237	0	-.042	-.039
Uruguay	-.43	-.008	.046	-.012
Venezuela	-.319	0	-.005	-.014

Table A-8: Counterfactual Ag Net Export Share of GDP - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Bahamas	-.237	0	-.008	-.022
Belize	-.237	0	.027	.014
Canada	-.022	-.007	.005	.012
Costa Rica	-.237	0	-.02	-.028
Dominican Republic	-.237	0	-.016	-.03
El Salvador	-.237	0	.027	-.019
Guatemala	-.237	-.017	.045	.021
Haiti	-.237	-.045	-.089	-.091
Jamaica	-.237	0	-.011	-.029
Mexico	-.354	0	-.029	-.055
Nicaragua	-.237	0	-.032	-.081
Panama	-.237	0	-.003	-.017
United States	-.059	.003	.012	.016

Table A-9: Counterfactual Ag Net Export Share of GDP - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Albania	-.086	-.053	.039	.033
Armenia	-.226	-.089	-.015	-.033
Austria	-.05	-.029	-.009	-.006
Belarus	.031	.012	0	.006
Belgium	-.067	-.015	.005	.005
Bosnia and Herzegovina	-.086	-.052	-.02	-.013
Bulgaria	-.086	-.054	.022	.028
Croatia	-.05	-.044	.001	.014
Cyprus	-.078	0	.045	.038
Czech Republic	-.05	-.021	-.009	-.009
Denmark	.109	.006	.019	.038
Estonia	.031	.033	.008	.028
Finland	.109	.02	-.006	-.001
France	-.067	-.027	.014	.02
Georgia	-.226	-.091	-.065	-.108
Germany	-.029	-.021	-.005	-.002
Greece	-.078	-.012	.011	.015
Hungary	-.05	-.05	-.002	.005
Iceland	.109	.009	.049	.073
Ireland	-.039	0	-.017	-.015
Israel	-.27	0	.002	-.005
Italy	-.074	-.018	.003	.005

Table A-10: Counterfactual Ag Net Export Share of GDP - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Latvia	.031	.032	.002	.031
Lithuania	.031	.018	-.001	.016
Luxembourg	-.05	-.018	.001	.006
Macedonia	-.086	-.061	.042	.046
Malta	-.074	0	-.011	.003
Moldova	-.086	-.07	.057	-.012
Montenegro	-.086	-.03	.018	.026
Netherlands	-.07	-.007	.013	.015
Norway	.109	-.006	.009	.018
Poland	-.047	-.01	-.019	-.018
Portugal	-.096	-.009	-.006	-.005
Romania	-.066	-.056	-.03	-.022
Russia	-.077	-.006	0	.001
Serbia	-.086	-.075	.004	.011
Slovakia	-.05	-.032	-.005	-.005
Slovenia	-.05	-.039	-.025	-.024
Spain	-.089	-.012	.006	.006
Sweden	.109	.011	-.005	.001
Switzerland	-.05	-.031	-.013	-.012
Turkey	-.162	-.038	.004	.007
Ukraine	-.052	-.038	.012	.035
United Kingdom	-.039	-.002	.002	.007

Table A-11: Counterfactual Ag Net Export Share of GDP - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Australia	-.266	0	.024	.013
Fiji	.022	-.004	-.01	.044
Indonesia	-.179	0	-.007	-.009
Malaysia	-.225	0	-.005	-.009
New Zealand	.022	0	.062	.077
Singapore	-.225	0	-.014	-.016

Table A-12: Counterfactual Ag Domestic Expenditure Shares - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Angola	-.258	0	.777	.684
Benin	-.327	0	.636	.462
Botswana	-.469	0	.338	.178
Burkina Faso	-.243	0	.825	.754
Cameroon	-.2	-.005	.815	.753
Cape Verde	-.327	-.012	.542	.197
Central African Republic	-.601	-.076	.709	.323
Chad	-.601	0	.978	.681
Comoros	-.217	-.075	.299	.236
Congo	-.601	0	.408	.041
Cote d’Ivoire	-.143	0	.308	.297
Democratic Republic of Congo	-.147	-.142	.866	.887
Ethiopia	-.313	-.102	.958	.929
Gabon	-.601	0	.567	.098
Gambia	-.327	0	.046	.074
Ghana	-.14	0	.814	.787
Kenya	-.054	-.044	.867	.906
Lesotho	-.469	-.055	.271	.084
Madagascar	-.262	-.067	.786	.738

Table A-13: Counterfactual Ag Domestic Expenditure Shares - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Malawi	-.313	-.111	.791	.779
Mali	-.356	0	.797	.654
Mauritania	-.327	0	.527	.376
Mauritius	-.262	0	.129	.074
Mozambique	-.217	-.104	.758	.792
Namibia	-.469	0	.259	.1
Niger	-.341	0	.843	.598
Nigeria	-.185	0	.944	.925
Rwanda	-.601	-.058	.848	.538
Senegal	-.519	0	.571	.139
Seychelles	-.262	0	.001	.001
Sierra Leone	-.327	-.071	.598	.733
Somalia	-.166	-.125	.673	.695
South Africa	-.334	0	.67	.472
Sudan	-.561	0	.914	.573
Swaziland	-.469	-.006	.336	.172
Tanzania	-.242	-.057	.867	.827
Togo	-.327	-.042	.522	.521
Uganda	-.168	-.057	.918	.885
Zambia	-.396	0	.838	.756
Zimbabwe	-.379	-.099	.563	.498

Table A-14: Counterfactual Ag Domestic Expenditure Shares - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Algeria	-.36	-.019	.713	.444
Bahrain	-.219	0	.104	.084
Egypt	.113	0	.735	.821
Iran	-.289	.002	.864	.794
Iraq	-.411	0	.545	.332
Jordan	-.27	-.019	.402	.286
Kuwait	-.219	0	.141	.117
Lebanon	-.27	-.012	.8	.655
Libya	-.124	0	.601	.615
Morocco	-.39	-.038	.689	.452
Oman	-.219	0	.098	.071
Qatar	-.219	0	.581	.501
Saudi Arabia	-.219	0	.521	.43
Syria	-.27	-.049	.715	.598
Tunisia	-.36	-.019	.604	.324
United Arab Emirates	-.219	0	.287	.22
Yemen	-.282	-.031	.664	.589

Table A-15: Counterfactual Ag Domestic Expenditure Shares - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Afghanistan	-.247	-.086	.742	.77
Azerbaijan	-.226	-.044	.748	.68
Bangladesh	-.217	-.016	.832	.802
Bhutan	-.381	-.034	.61	.439
Brunei	-.179	0	.407	.413
Cambodia	-.271	0	.795	.664
China	-.072	-.036	.711	.757
Hong Kong	-.072	0	.345	.366
India	-.381	0	.943	.862
Japan	-.057	-.01	.506	.553
Kazakhstan	.114	-.031	.892	.942
Kyrgyzstan	-.059	-.039	.694	.622
Maldives	-.201	0	.112	.11
Myanmar	-.393	0	.937	.887
Nepal	-.173	-.07	.937	.948
Pakistan	-.304	.001	.946	.89
Philippines	-.234	0	.636	.558
South Korea	-.093	-.015	.254	.268
Sri Lanka	-.201	0	.65	.68
Tajikistan	-.059	-.097	.799	.782
Thailand	-.262	0	.373	.264
Uzbekistan	-.121	-.065	.956	.95
Vietnam	-.151	0	.482	.483

Table A-16: Counterfactual Ag Domestic Expenditure Shares - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Argentina	-.111	0	.919	.925
Barbados	-.237	0	.058	.068
Bolivia	-.43	-.042	.758	.512
Brazil	-.169	0	.896	.883
Chile	-.244	-.009	.506	.4
Colombia	-.232	0	.837	.806
Ecuador	-.288	0	.539	.503
Honduras	-.237	-.006	.219	.163
Paraguay	-.43	0	.485	.195
Peru	-.306	-.005	.717	.628
Suriname	-.43	0	.413	.074
Trinidad and Tobago	-.237	0	.016	.01
Uruguay	-.43	-.008	.592	.288
Venezuela	-.319	0	.793	.63

Table A-17: Counterfactual Ag Domestic Expenditure Shares - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Bahamas	-.237	0	.245	.122
Belize	-.237	0	.077	.061
Canada	-.022	-.007	.268	.32
Costa Rica	-.237	0	.01	.005
Dominican Republic	-.237	0	.547	.455
El Salvador	-.237	0	.546	.41
Guatemala	-.237	-.017	.487	.425
Haiti	-.237	-.045	.628	.593
Jamaica	-.237	0	.57	.541
Mexico	-.354	0	.442	.182
Nicaragua	-.237	0	.255	.194
Panama	-.237	0	.477	.359
United States	-.059	.003	.835	.853

Table A-18: Counterfactual Ag Domestic Expenditure Shares - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Albania	-.086	-.053	.84	.814
Armenia	-.226	-.089	.79	.662
Austria	-.05	-.029	.088	.094
Belarus	.031	.012	.846	.862
Belgium	-.067	-.015	.099	.093
Bosnia and Herzegovina	-.086	-.052	.53	.526
Bulgaria	-.086	-.054	.382	.383
Croatia	-.05	-.044	.576	.639
Cyprus	-.078	0	.644	.611
Czech Republic	-.05	-.021	.141	.148
Denmark	.109	.006	.163	.214
Estonia	.031	.033	.372	.405
Finland	.109	.02	.545	.678
France	-.067	-.027	.549	.581
Georgia	-.226	-.091	.39	.294
Germany	-.029	-.021	.147	.178
Greece	-.078	-.012	.588	.585
Hungary	-.05	-.05	.225	.236
Iceland	.109	.009	.057	.069
Ireland	-.039	0	.002	.002
Israel	-.27	0	.664	.457
Italy	-.074	-.018	.571	.592

Table A-19: Counterfactual Ag Domestic Expenditure Shares - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Latvia	.031	.032	.284	.29
Lithuania	.031	.018	.081	.089
Luxembourg	-.05	-.018	.065	.085
Macedonia	-.086	-.061	.598	.611
Malta	-.074	0	.067	.073
Moldova	-.086	-.07	.452	.392
Montenegro	-.086	-.03	.646	.659
Netherlands	-.07	-.007	.088	.091
Norway	.109	-.006	.285	.354
Poland	-.047	-.01	.308	.302
Portugal	-.096	-.009	.196	.185
Romania	-.066	-.056	.572	.625
Russia	-.077	-.006	.753	.764
Serbia	-.086	-.075	.726	.745
Slovakia	-.05	-.032	.288	.289
Slovenia	-.05	-.039	.186	.195
Spain	-.089	-.012	.418	.422
Sweden	.109	.011	.298	.421
Switzerland	-.05	-.031	.021	.025
Turkey	-.162	-.038	.894	.88
Ukraine	-.052	-.038	.681	.689
United Kingdom	-.039	-.002	.478	.531

Table A-20: Counterfactual Ag Domestic Expenditure Shares - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	Baseline	Counterfactual
Australia	-.266	0	.766	.617
Fiji	.022	-.004	.388	.528
Indonesia	-.179	0	.77	.732
Malaysia	-.225	0	.208	.175
New Zealand	.022	0	.606	.69
Singapore	-.225	0	.01	.008

Table A-21: Counterfactual Ag GDP Shares - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Angola	-.258	0	.05	.065	.055
Benin	-.327	0	.156	.213	.154
Botswana	-.469	0	.062	.149	.044
Burkina Faso	-.243	0	.24	.289	.233
Cameroon	-.2	-.005	.216	.254	.226
Cape Verde	-.327	-.012	.163	.242	.076
Central African Republic	-.601	-.076	.299	.562	.21
Chad	-.601	0	.257	.438	.258
Comoros	-.217	-.075	.083	.123	.067
Congo	-.601	0	.092	.257	.011
Cote d'Ivoire	-.143	0	.17	.188	.183
Democratic Republic of Congo	-.147	-.142	.421	.457	.506
Ethiopia	-.313	-.102	.359	.437	.409
Gabon	-.601	0	.073	.187	.018
Gambia	-.327	0	.008	.06	.021
Ghana	-.14	0	.181	.195	.193
Kenya	-.054	-.044	.156	.16	.185
Lesotho	-.469	-.055	.107	.191	.028
Madagascar	-.262	-.067	.404	.481	.436

Table A-22: Counterfactual Ag GDP Shares - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Malawi	-.313	-.111	.436	.54	.55
Mali	-.356	0	.197	.278	.205
Mauritania	-.327	0	.249	.336	.224
Mauritius	-.262	0	.038	.048	.016
Mozambique	-.217	-.104	.367	.426	.451
Namibia	-.469	0	.099	.157	.04
Niger	-.341	0	.325	.432	.277
Nigeria	-.185	0	.068	.078	.077
Rwanda	-.601	-.058	.409	.678	.351
Senegal	-.519	0	.153	.267	.05
Seychelles	-.262	0	.016	.026	.025
Sierra Leone	-.327	-.071	.139	.204	.246
Somalia	-.166	-.125	.334	.373	.382
South Africa	-.334	0	.056	.073	.052
Sudan	-.561	0	.12	.197	.095
Swaziland	-.469	-.006	.102	.224	.062
Tanzania	-.242	-.057	.298	.354	.321
Togo	-.327	-.042	.316	.372	.316
Uganda	-.168	-.057	.416	.459	.433
Zambia	-.396	0	.36	.496	.41
Zimbabwe	-.379	-.099	.302	.419	.322

Table A-23: Counterfactual Ag GDP Shares - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Algeria	-.36	-.019	.034	.049	.025
Bahrain	-.219	0	.015	.021	.02
Egypt	.113	0	.077	.071	.098
Iran	-.289	.002	.051	.065	.057
Iraq	-.411	0	.05	.081	.034
Jordan	-.27	-.019	.038	.05	.046
Kuwait	-.219	0	.005	.011	.005
Lebanon	-.27	-.012	.079	.083	.068
Libya	-.124	0	.067	.073	.066
Morocco	-.39	-.038	.1	.141	.077
Oman	-.219	0	.023	.029	.02
Qatar	-.219	0	.008	.011	.009
Saudi Arabia	-.219	0	.016	.021	.016
Syria	-.27	-.049	.09	.114	.084
Tunisia	-.36	-.019	.055	.076	.033
United Arab Emirates	-.219	0	.016	.018	.019
Yemen	-.282	-.031	.11	.142	.121

Table A-24: Counterfactual Ag GDP Shares - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Afghanistan	-.247	-.086	.278	.33	.321
Azerbaijan	-.226	-.044	.088	.108	.095
Bangladesh	-.217	-.016	.189	.228	.212
Bhutan	-.381	-.034	.267	.392	.239
Brunei	-.179	0	.027	.03	.022
Cambodia	-.271	0	.19	.24	.189
China	-.072	-.036	.064	.068	.074
Hong Kong	-.072	0	.018	.019	.021
India	-.381	0	.161	.224	.194
Japan	-.057	-.01	.012	.013	.014
Kazakhstan	.114	-.031	.088	.08	.091
Kyrgyzstan	-.059	-.039	.156	.162	.152
Maldives	-.201	0	.021	.03	.039
Myanmar	-.393	0	.209	.288	.266
Nepal	-.173	-.07	.283	.317	.328
Pakistan	-.304	.001	.133	.167	.148
Philippines	-.234	0	.105	.133	.104
South Korea	-.093	-.015	.012	.014	.013
Sri Lanka	-.201	0	.109	.129	.137
Tajikistan	-.059	-.097	.232	.239	.233
Thailand	-.262	0	.072	.099	.061
Uzbekistan	-.121	-.065	.108	.115	.14
Vietnam	-.151	0	.158	.178	.155

Table A-25: Counterfactual Ag GDP Shares - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	.065	.067	.07
Barbados	-.237	0	.035	.04	.054
Bolivia	-.43	-.042	.125	.186	.115
Brazil	-.169	0	.063	.071	.068
Chile	-.244	-.009	.053	.063	.045
Colombia	-.232	0	.066	.077	.07
Ecuador	-.288	0	.098	.12	.098
Honduras	-.237	-.006	.082	.111	.064
Paraguay	-.43	0	.151	.194	.065
Peru	-.306	-.005	.104	.133	.099
Suriname	-.43	0	.028	.046	.005
Trinidad and Tobago	-.237	0	.004	.014	.011
Uruguay	-.43	-.008	.091	.113	.048
Venezuela	-.319	0	.027	.037	.026

Table A-26: Counterfactual Ag GDP Shares - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Bahamas	-.237	0	.02	.025	.007
Belize	-.237	0	.104	.12	.104
Canada	-.022	-.007	.019	.019	.026
Costa Rica	-.237	0	.021	.03	.017
Dominican Republic	-.237	0	.046	.059	.043
El Salvador	-.237	0	.108	.126	.075
Guatemala	-.237	-.017	.151	.173	.146
Haiti	-.237	-.045	.161	.208	.203
Jamaica	-.237	0	.074	.092	.073
Mexico	-.354	0	.031	.052	.014
Nicaragua	-.237	0	.11	.139	.082
Panama	-.237	0	.06	.074	.056
United States	-.059	.003	.023	.024	.028

Table A-27: Counterfactual Ag GDP Shares - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Albania	-.086	-.053	.129	.134	.127
Armenia	-.226	-.089	.104	.126	.104
Austria	-.05	-.029	.012	.012	.015
Belarus	.031	.012	.05	.049	.056
Belgium	-.067	-.015	.02	.021	.02
Bosnia and Herzegovina	-.086	-.052	.047	.051	.057
Bulgaria	-.086	-.054	.06	.062	.068
Croatia	-.05	-.044	.042	.043	.058
Cyprus	-.078	0	.06	.061	.054
Czech Republic	-.05	-.021	.018	.019	.018
Denmark	.109	.006	.033	.032	.051
Estonia	.031	.033	.036	.036	.056
Finland	.109	.02	.013	.012	.017
France	-.067	-.027	.026	.027	.033
Georgia	-.226	-.091	.086	.113	.062
Germany	-.029	-.021	.011	.011	.015
Greece	-.078	-.012	.037	.038	.043
Hungary	-.05	-.05	.028	.029	.037
Iceland	.109	.009	.063	.062	.086
Ireland	-.039	0	.002	.003	.006
Israel	-.27	0	.016	.019	.011
Italy	-.074	-.018	.019	.019	.021

Table A-28: Counterfactual Ag GDP Shares - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Latvia	.031	.032	.037	.036	.064
Lithuania	.031	.018	.035	.034	.052
Luxembourg	-.05	-.018	.013	.014	.019
Macedonia	-.086	-.061	.104	.108	.112
Malta	-.074	0	.006	.007	.022
Moldova	-.086	-.07	.169	.175	.104
Montenegro	-.086	-.03	.054	.057	.065
Netherlands	-.07	-.007	.027	.028	.029
Norway	.109	-.006	.02	.019	.027
Poland	-.047	-.01	.031	.033	.034
Portugal	-.096	-.009	.026	.028	.029
Romania	-.066	-.056	.06	.063	.073
Russia	-.077	-.006	.03	.032	.033
Serbia	-.086	-.075	.069	.073	.08
Slovakia	-.05	-.032	.022	.023	.022
Slovenia	-.05	-.039	.017	.018	.019
Spain	-.089	-.012	.023	.024	.024
Sweden	.109	.011	.009	.008	.015
Switzerland	-.05	-.031	.003	.004	.005
Turkey	-.162	-.038	.047	.053	.055
Ukraine	-.052	-.038	.105	.108	.132
United Kingdom	-.039	-.002	.015	.015	.02

Table A-29: Counterfactual Ag GDP Shares - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Australia	-.266	0	.036	.039	.027
Fiji	.022	-.004	.126	.124	.185
Indonesia	-.179	0	.112	.129	.126
Malaysia	-.225	0	.042	.051	.045
New Zealand	.022	0	.078	.077	.093
Singapore	-.225	0	.007	.011	.007

Table A-30: Counterfactual GDP Losses (Share of GDP) - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Angola	-.258	0	-.018	-.027	-.025
Benin	-.327	0	-.07	-.101	-.084
Botswana	-.469	0	-.096	-.196	-.147
Burkina Faso	-.243	0	-.065	-.091	-.082
Cameroon	-.2	-.005	-.053	-.097	-.091
Cape Verde	-.327	-.012	-.108	-.216	-.077
Central African Republic	-.601	-.076	-.331	-.527	-.45
Chad	-.601	0	-.182	-.363	-.332
Comoros	-.217	-.075	-.102	-.112	-.047
Congo	-.601	0	-.165	-.323	-.121
Cote d'Ivoire	-.143	0	-.025	-.032	-.033
Democratic Republic of Congo	-.147	-.142	-.132	-.174	-.174
Ethiopia	-.313	-.102	-.163	-.218	-.217
Gabon	-.601	0	-.111	-.247	-.117
Gambia	-.327	0	-.065	-.063	-.096
Ghana	-.14	0	-.018	-.023	-.018
Kenya	-.054	-.044	-.037	-.038	-.034
Lesotho	-.469	-.055	-.131	-.18	-.109
Madagascar	-.262	-.067	-.144	-.209	-.2

Table A-31: Counterfactual GDP Losses (Share of GDP) - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Malawi	-.313	-.111	-.209	-.317	-.318
Mali	-.356	0	-.1	-.144	-.127
Mauritania	-.327	0	-.112	-.203	-.167
Mauritius	-.262	0	-.012	-.015	-.008
Mozambique	-.217	-.104	-.14	-.192	-.199
Namibia	-.469	0	-.064	-.11	-.063
Niger	-.341	0	-.142	-.231	-.177
Nigeria	-.185	0	-.013	-.015	-.015
Rwanda	-.601	-.058	-.334	-.557	-.508
Senegal	-.519	0	-.122	-.233	-.123
Seychelles	-.262	0	-.013	-.011	-.005
Sierra Leone	-.327	-.071	-.123	-.142	-.168
Somalia	-.166	-.125	-.126	-.151	-.15
South Africa	-.334	0	-.02	-.028	-.023
Sudan	-.561	0	-.078	-.144	-.125
Swaziland	-.469	-.006	-.14	-.293	-.215
Tanzania	-.242	-.057	-.108	-.152	-.149
Togo	-.327	-.042	-.093	-.109	-.132
Uganda	-.168	-.057	-.094	-.131	-.118
Zambia	-.396	0	-.175	-.328	-.314
Zimbabwe	-.379	-.099	-.202	-.328	-.312

Table A-32: Counterfactual GDP Losses (Share of GDP) - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Algeria	-.36	-.019	-.031	-.037	-.033
Bahrain	-.219	0	-.007	-.01	-.01
Egypt	.113	0	.008	.01	.011
Iran	-.289	.002	-.017	-.027	-.026
Iraq	-.411	0	-.035	-.058	-.042
Jordan	-.27	-.019	-.028	-.029	-.027
Kuwait	-.219	0	-.007	-.01	-.012
Lebanon	-.27	-.012	-.015	-.014	-.016
Libya	-.124	0	-.008	-.011	-.008
Morocco	-.39	-.038	-.076	-.11	-.099
Oman	-.219	0	-.007	-.009	-.009
Qatar	-.219	0	-.003	-.003	-.003
Saudi Arabia	-.219	0	-.006	-.006	-.005
Syria	-.27	-.049	-.066	-.082	-.077
Tunisia	-.36	-.019	-.039	-.054	-.045
United Arab Emirates	-.219	0	-.003	-.002	-.003
Yemen	-.282	-.031	-.062	-.072	-.064

Table A-33: Counterfactual GDP Losses (Share of GDP) - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Afghanistan	-.247	-.086	-.131	-.167	-.169
Azerbaijan	-.226	-.044	-.058	-.073	-.066
Bangladesh	-.217	-.016	-.059	-.09	-.086
Bhutan	-.381	-.034	-.181	-.277	-.233
Brunei	-.179	0	-.004	-.003	0
Cambodia	-.271	0	-.065	-.085	-.068
China	-.072	-.036	-.043	-.045	-.045
Hong Kong	-.072	0	-.001	-.002	-.002
India	-.381	0	-.074	-.131	-.127
Japan	-.057	-.01	-.011	-.012	-.009
Kazakhstan	.114	-.031	-.036	-.03	-.031
Kyrgyzstan	-.059	-.039	-.061	-.063	-.049
Maldives	-.201	0	-.012	-.016	-.037
Myanmar	-.393	0	-.094	-.144	-.14
Nepal	-.173	-.07	-.093	-.12	-.114
Pakistan	-.304	.001	-.041	-.062	-.061
Philippines	-.234	0	-.036	-.053	-.051
South Korea	-.093	-.015	-.023	-.024	-.028
Sri Lanka	-.201	0	-.026	-.044	-.048
Tajikistan	-.059	-.097	-.081	-.085	-.08
Thailand	-.262	0	-.034	-.055	-.041
Uzbekistan	-.121	-.065	-.064	-.066	-.063
Vietnam	-.151	0	-.028	-.04	-.042

Table A-34: Counterfactual GDP Losses (Share of GDP) - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.002	0	.001
Barbados	-.237	0	-.006	-.007	-.032
Bolivia	-.43	-.042	-.098	-.135	-.116
Brazil	-.169	0	-.01	-.015	-.013
Chile	-.244	-.009	-.02	-.026	-.024
Colombia	-.232	0	-.015	-.022	-.02
Ecuador	-.288	0	-.027	-.04	-.043
Honduras	-.237	-.006	-.039	-.058	-.039
Paraguay	-.43	0	-.049	-.058	-.036
Peru	-.306	-.005	-.037	-.06	-.058
Suriname	-.43	0	-.021	-.014	-.012
Trinidad and Tobago	-.237	0	-.012	-.016	-.016
Uruguay	-.43	-.008	-.031	-.039	-.031
Venezuela	-.319	0	-.012	-.014	-.014

Table A-35: Counterfactual GDP Losses (Share of GDP) - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Bahamas	-.237	0	-.007	-.008	-.002
Belize	-.237	0	-.021	-.021	-.026
Canada	-.022	-.007	-.018	-.018	-.016
Costa Rica	-.237	0	-.011	-.011	-.004
Dominican Republic	-.237	0	-.017	-.023	-.026
El Salvador	-.237	0	-.022	-.028	-.018
Guatemala	-.237	-.017	-.039	-.056	-.05
Haiti	-.237	-.045	-.09	-.123	-.121
Jamaica	-.237	0	-.023	-.035	-.033
Mexico	-.354	0	-.026	-.043	-.022
Nicaragua	-.237	0	-.038	-.047	-.036
Panama	-.237	0	-.017	-.029	-.027
United States	-.059	.003	0	0	.001

Table A-36: Counterfactual GDP Losses (Share of GDP) - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Albania	-.086	-.053	-.05	-.052	-.046
Armenia	-.226	-.089	-.099	-.106	-.103
Austria	-.05	-.029	-.032	-.032	-.029
Belarus	.031	.012	-.011	-.012	-.012
Belgium	-.067	-.015	-.018	-.018	-.017
Bosnia and Herzegovina	-.086	-.052	-.052	-.054	-.054
Bulgaria	-.086	-.054	-.051	-.051	-.049
Croatia	-.05	-.044	-.041	-.042	-.046
Cyprus	-.078	0	-.001	-.001	.002
Czech Republic	-.05	-.021	-.031	-.032	-.026
Denmark	.109	.006	0	0	.005
Estonia	.031	.033	.009	.008	.002
Finland	.109	.02	.003	.003	-.001
France	-.067	-.027	-.027	-.028	-.027
Georgia	-.226	-.091	-.1	-.127	-.107
Germany	-.029	-.021	-.025	-.025	-.025
Greece	-.078	-.012	-.012	-.013	-.017
Hungary	-.05	-.05	-.048	-.048	-.049
Iceland	.109	.009	.005	.005	.014
Ireland	-.039	0	-.001	-.001	-.001
Israel	-.27	0	-.004	-.005	-.009
Italy	-.074	-.018	-.017	-.017	-.018

Table A-37: Counterfactual GDP Losses (Share of GDP) - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Latvia	.031	.032	.006	.006	.009
Lithuania	.031	.018	-.005	-.005	-.004
Luxembourg	-.05	-.018	-.016	-.016	-.018
Macedonia	-.086	-.061	-.059	-.059	-.054
Malta	-.074	0	-.001	-.001	-.004
Moldova	-.086	-.07	-.071	-.073	-.051
Montenegro	-.086	-.03	-.032	-.031	-.035
Netherlands	-.07	-.007	-.01	-.01	-.014
Norway	.109	-.006	-.004	-.004	-.002
Poland	-.047	-.01	-.024	-.025	-.023
Portugal	-.096	-.009	-.011	-.013	-.013
Romania	-.066	-.056	-.054	-.057	-.061
Russia	-.077	-.006	-.026	-.027	-.026
Serbia	-.086	-.075	-.069	-.07	-.063
Slovakia	-.05	-.032	-.039	-.039	-.038
Slovenia	-.05	-.039	-.039	-.04	-.036
Spain	-.089	-.012	-.011	-.012	-.009
Sweden	.109	.011	-.002	-.002	-.005
Switzerland	-.05	-.031	-.031	-.031	-.035
Turkey	-.162	-.038	-.039	-.042	-.041
Ukraine	-.052	-.038	-.045	-.046	-.042
United Kingdom	-.039	-.002	-.004	-.004	-.003

Table A-38: Counterfactual GDP Losses (Share of GDP) - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Australia	-.266	0	-.004	-.005	-.006
Fiji	.022	-.004	.001	.003	-.013
Indonesia	-.179	0	-.023	-.034	-.027
Malaysia	-.225	0	-.012	-.013	-.012
New Zealand	.022	0	0	0	.009
Singapore	-.225	0	-.005	-.006	-.01

Table A-39: Equivalent Variation Willingness-to-Pay (Share of GDP) - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Angola	-.258	0	-.083	-.019	-.018
Benin	-.327	0	-.276	-.078	-.069
Botswana	-.469	0	-.423	-.116	-.087
Burkina Faso	-.243	0	-.226	-.07	-.063
Cameroon	-.2	-.005	-.18	-.055	-.052
Cape Verde	-.327	-.012	-.344	-.12	-.046
Central African Republic	-.601	-.076	-.723	-.428	-.356
Chad	-.601	0	-.67	-.25	-.226
Comoros	-.217	-.075	-.222	-.102	-.065
Congo	-.601	0	-.66	-.225	-.079
Cote d'Ivoire	-.143	0	-.093	-.025	-.025
Democratic Republic of Congo	-.147	-.142	-.209	-.131	-.129
Ethiopia	-.313	-.102	-.364	-.171	-.169
Gabon	-.601	0	-.572	-.15	-.069
Gambia	-.327	0	-.261	-.072	-.133
Ghana	-.14	0	-.07	-.018	-.014
Kenya	-.054	-.044	-.052	-.035	-.031
Lesotho	-.469	-.055	-.432	-.147	-.072
Madagascar	-.262	-.067	-.324	-.153	-.146

Table A-40: Equivalent Variation Willingness-to-Pay (Share of GDP) - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Malawi	-.313	-.111	-.4	-.225	-.225
Mali	-.356	0	-.359	-.115	-.106
Mauritania	-.327	0	-.36	-.126	-.101
Mauritius	-.262	0	-.057	-.013	-.007
Mozambique	-.217	-.104	-.279	-.143	-.147
Namibia	-.469	0	-.328	-.076	-.044
Niger	-.341	0	-.402	-.163	-.121
Nigeria	-.185	0	-.054	-.013	-.012
Rwanda	-.601	-.058	-.725	-.434	-.387
Senegal	-.519	0	-.513	-.155	-.08
Seychelles	-.262	0	-.064	-.014	-.023
Sierra Leone	-.327	-.071	-.33	-.13	-.164
Somalia	-.166	-.125	-.22	-.126	-.123
South Africa	-.334	0	-.103	-.022	-.019
Sudan	-.561	0	-.436	-.101	-.087
Swaziland	-.469	-.006	-.504	-.172	-.121
Tanzania	-.242	-.057	-.268	-.112	-.109
Togo	-.327	-.042	-.289	-.099	-.125
Uganda	-.168	-.057	-.205	-.096	-.085
Zambia	-.396	0	-.481	-.208	-.199
Zimbabwe	-.379	-.099	-.464	-.223	-.212

Table A-41: Equivalent Variation Willingness-to-Pay (Share of GDP) - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Algeria	-.36	-.019	-.104	-.031	-.028
Bahrain	-.219	0	-.034	-.007	-.008
Egypt	.113	0	.028	.008	.009
Iran	-.289	.002	-.085	-.018	-.016
Iraq	-.411	0	-.192	-.04	-.031
Jordan	-.27	-.019	-.082	-.028	-.026
Kuwait	-.219	0	-.032	-.007	-.01
Lebanon	-.27	-.012	-.037	-.014	-.015
Libya	-.124	0	-.033	-.008	-.005
Morocco	-.39	-.038	-.252	-.08	-.076
Oman	-.219	0	-.033	-.007	-.007
Qatar	-.219	0	-.013	-.003	-.003
Saudi Arabia	-.219	0	-.029	-.006	-.005
Syria	-.27	-.049	-.164	-.065	-.063
Tunisia	-.36	-.019	-.141	-.041	-.035
United Arab Emirates	-.219	0	-.014	-.003	-.002
Yemen	-.282	-.031	-.191	-.063	-.056

Table A-42: Equivalent Variation Willingness-to-Pay (Share of GDP) - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Afghanistan	-.247	-.086	-.28	-.133	-.134
Azerbaijan	-.226	-.044	-.139	-.057	-.053
Bangladesh	-.217	-.016	-.189	-.062	-.058
Bhutan	-.381	-.034	-.466	-.208	-.173
Brunei	-.179	0	-.019	-.004	-.004
Cambodia	-.271	0	-.237	-.07	-.057
China	-.072	-.036	-.057	-.04	-.04
Hong Kong	-.072	0	-.006	-.001	-.001
India	-.381	0	-.311	-.085	-.082
Japan	-.057	-.01	-.015	-.011	-.008
Kazakhstan	.114	-.031	-.01	-.034	-.035
Kyrgyzstan	-.059	-.039	-.078	-.058	-.047
Maldives	-.201	0	-.052	-.012	-.017
Myanmar	-.393	0	-.367	-.109	-.105
Nepal	-.173	-.07	-.192	-.092	-.084
Pakistan	-.304	.001	-.179	-.045	-.044
Philippines	-.234	0	-.145	-.038	-.038
South Korea	-.093	-.015	-.032	-.021	-.024
Sri Lanka	-.201	0	-.105	-.027	-.03
Tajikistan	-.059	-.097	-.1	-.078	-.074
Thailand	-.262	0	-.144	-.036	-.028
Uzbekistan	-.121	-.065	-.088	-.061	-.057
Vietnam	-.151	0	-.101	-.028	-.029

Table A-43: Equivalent Variation Willingness-to-Pay (Share of GDP) - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Argentina	-.111	0	-.008	-.002	0
Barbados	-.237	0	-.029	-.006	-.029
Bolivia	-.43	-.042	-.343	-.108	-.095
Brazil	-.169	0	-.041	-.01	-.008
Chile	-.244	-.009	-.068	-.02	-.019
Colombia	-.232	0	-.067	-.015	-.013
Ecuador	-.288	0	-.123	-.028	-.031
Honduras	-.237	-.006	-.153	-.041	-.032
Paraguay	-.43	0	-.255	-.057	-.037
Peru	-.306	-.005	-.162	-.04	-.04
Suriname	-.43	0	-.125	-.024	-.008
Trinidad and Tobago	-.237	0	-.057	-.013	-.012
Uruguay	-.43	-.008	-.151	-.034	-.028
Venezuela	-.319	0	-.064	-.013	-.012

Table A-44: Equivalent Variation Willingness-to-Pay (Share of GDP) - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Bahamas	-.237	0	-.035	-.008	-.002
Belize	-.237	0	-.091	-.021	-.024
Canada	-.022	-.007	-.018	-.016	-.014
Costa Rica	-.237	0	-.052	-.011	-.004
Dominican Republic	-.237	0	-.075	-.017	-.022
El Salvador	-.237	0	-.095	-.023	-.015
Guatemala	-.237	-.017	-.128	-.04	-.036
Haiti	-.237	-.045	-.237	-.093	-.091
Jamaica	-.237	0	-.099	-.024	-.024
Mexico	-.354	0	-.132	-.028	-.015
Nicaragua	-.237	0	-.151	-.039	-.031
Panama	-.237	0	-.076	-.017	-.018
United States	-.059	.003	-.002	0	.001

Table A-45: Equivalent Variation Willingness-to-Pay (Share of GDP) - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Albania	-.086	-.053	-.067	-.047	-.04
Armenia	-.226	-.089	-.176	-.095	-.087
Austria	-.05	-.029	-.033	-.029	-.026
Belarus	.031	.012	-.006	-.011	-.011
Belgium	-.067	-.015	-.02	-.017	-.016
Bosnia and Herzegovina	-.086	-.052	-.065	-.049	-.049
Bulgaria	-.086	-.054	-.058	-.047	-.045
Croatia	-.05	-.044	-.045	-.038	-.041
Cyprus	-.078	0	-.005	-.001	.002
Czech Republic	-.05	-.021	-.034	-.029	-.024
Denmark	.109	.006	.003	0	.005
Estonia	.031	.033	.011	.008	.002
Finland	.109	.02	.008	.003	-.001
France	-.067	-.027	-.029	-.025	-.024
Georgia	-.226	-.091	-.193	-.098	-.084
Germany	-.029	-.021	-.025	-.023	-.023
Greece	-.078	-.012	-.018	-.011	-.015
Hungary	-.05	-.05	-.05	-.045	-.045
Iceland	.109	.009	.009	.005	.014
Ireland	-.039	0	-.003	-.001	-.002
Israel	-.27	0	-.022	-.005	-.009
Italy	-.074	-.018	-.02	-.016	-.016

Table A-46: Equivalent Variation Willingness-to-Pay (Share of GDP) - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Latvia	.031	.032	.009	.006	.008
Lithuania	.031	.018	-.001	-.004	-.003
Luxembourg	-.05	-.018	-.017	-.015	-.017
Macedonia	-.086	-.061	-.07	-.055	-.049
Malta	-.074	0	-.006	-.001	-.005
Moldova	-.086	-.07	-.09	-.067	-.048
Montenegro	-.086	-.03	-.04	-.03	-.032
Netherlands	-.07	-.007	-.012	-.009	-.012
Norway	.109	-.006	-.001	-.003	-.002
Poland	-.047	-.01	-.029	-.022	-.02
Portugal	-.096	-.009	-.02	-.01	-.011
Romania	-.066	-.056	-.065	-.051	-.053
Russia	-.077	-.006	-.031	-.024	-.023
Serbia	-.086	-.075	-.08	-.064	-.057
Slovakia	-.05	-.032	-.041	-.036	-.035
Slovenia	-.05	-.039	-.043	-.036	-.033
Spain	-.089	-.012	-.016	-.011	-.008
Sweden	.109	.011	.002	-.002	-.005
Switzerland	-.05	-.031	-.032	-.028	-.032
Turkey	-.162	-.038	-.06	-.036	-.036
Ukraine	-.052	-.038	-.054	-.042	-.037
United Kingdom	-.039	-.002	-.005	-.004	-.003

Table A-47: Equivalent Variation Willingness-to-Pay (Share of GDP) - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	No Reallocation	Autarky	With Trade
Australia	-.266	0	-.019	-.004	-.005
Fiji	.022	-.004	.008	.001	-.007
Indonesia	-.179	0	-.092	-.023	-.018
Malaysia	-.225	0	-.054	-.012	-.012
New Zealand	.022	0	.001	0	.009
Singapore	-.225	0	-.025	-.005	-.009

Table A-48: Counterfactual Change in Food Prices - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Angola	-.258	0	34.771	25.884
Benin	-.327	0	48.593	28.345
Botswana	-.469	0	88.326	46.253
Burkina Faso	-.243	0	32.099	27.752
Cameroon	-.2	-.005	25	23.809
Cape Verde	-.327	-.012	48.589	30.641
Central African Republic	-.601	-.076	150.624	36.303
Chad	-.601	0	150.634	117.717
Comoros	-.217	-.075	27.717	22.011
Congo	-.601	0	150.629	32.426
Cote d'Ivoire	-.143	0	16.691	18.443
Democratic Republic of Congo	-.147	-.142	17.231	10.576
Ethiopia	-.313	-.102	45.554	30.082
Gabon	-.601	0	150.632	34.985
Gambia	-.327	0	48.586	16.877
Ghana	-.14	0	16.282	20.566
Kenya	-.054	-.044	5.708	9.776
Lesotho	-.469	-.055	88.331	48.302
Madagascar	-.262	-.067	35.499	21.645

Table A-49: Counterfactual Change in Food Prices - Sub-Saharan Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Malawi	-.313	-.111	45.559	21.76
Mali	-.356	0	55.278	34.724
Mauritania	-.327	0	48.587	23.488
Mauritius	-.262	0	35.502	18.92
Mozambique	-.217	-.104	27.716	21.893
Namibia	-.469	0	88.323	38.591
Niger	-.341	0	51.744	50.943
Nigeria	-.185	0	22.699	22.24
Rwanda	-.601	-.058	150.628	51.389
Senegal	-.519	0	107.898	37.364
Seychelles	-.262	0	35.505	13.96
Sierra Leone	-.327	-.071	48.58	8.697
Somalia	-.166	-.125	19.904	17.825
South Africa	-.334	0	50.144	28.951
Sudan	-.561	0	127.79	69.284
Swaziland	-.469	-.006	88.323	46.552
Tanzania	-.242	-.057	31.926	23.751
Togo	-.327	-.042	48.589	15.661
Uganda	-.168	-.057	20.191	21.826
Zambia	-.396	0	65.563	45.63
Zimbabwe	-.379	-.099	61.035	35.921

Table A-50: Counterfactual Change in Food Prices - Middle East and North Africa

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Algeria	-.36	-.019	56.251	32.033
Bahrain	-.219	0	28.033	19.754
Egypt	.113	0	-10.152	2.943
Iran	-.289	.002	40.648	33.444
Iraq	-.411	0	69.775	32.181
Jordan	-.27	-.019	36.986	18.564
Kuwait	-.219	0	28.04	19.881
Lebanon	-.27	-.012	36.987	28.029
Libya	-.124	0	14.156	11.629
Morocco	-.39	-.038	63.943	27.302
Oman	-.219	0	28.039	17.587
Qatar	-.219	0	28.04	25.392
Saudi Arabia	-.219	0	28.041	20.696
Syria	-.27	-.049	36.986	23.632
Tunisia	-.36	-.019	56.25	28.447
United Arab Emirates	-.219	0	28.041	17.641
Yemen	-.282	-.031	39.278	25.125

Table A-51: Counterfactual Change in Food Prices - Asia

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Afghanistan	-.247	-.086	32.806	21.242
Azerbaijan	-.226	-.044	29.197	16.818
Bangladesh	-.217	-.016	27.714	25.683
Bhutan	-.381	-.034	61.553	40.819
Brunei	-.179	0	21.806	21.636
Cambodia	-.271	0	37.173	35.462
China	-.072	-.036	7.765	7.913
Hong Kong	-.072	0	7.758	10.497
India	-.381	0	61.556	47.244
Japan	-.057	-.01	6.041	9.157
Kazakhstan	.114	-.031	-10.235	-5.317
Kyrgyzstan	-.059	-.039	6.268	10.307
Maldives	-.201	0	25.163	20.722
Myanmar	-.393	0	64.742	41.112
Nepal	-.173	-.07	20.924	19.834
Pakistan	-.304	.001	43.678	40.005
Philippines	-.234	0	30.548	19.969
South Korea	-.093	-.015	10.254	12.046
Sri Lanka	-.201	0	25.158	17.877
Tajikistan	-.059	-.097	6.269	6.705
Thailand	-.262	0	35.503	17.386
Uzbekistan	-.121	-.065	13.766	12.432
Vietnam	-.151	0	17.786	17.872

Table A-52: Counterfactual Change in Food Prices - South America

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Argentina	-.111	0	12.486	14.87
Barbados	-.237	0	31.062	17.803
Bolivia	-.43	-.042	75.439	38.358
Brazil	-.169	0	20.34	18.347
Chile	-.244	-.009	32.277	22.685
Colombia	-.232	0	30.204	22.204
Ecuador	-.288	0	40.448	19.993
Honduras	-.237	-.006	31.063	16.769
Paraguay	-.43	0	75.435	27.023
Peru	-.306	-.005	44.091	23.927
Suriname	-.43	0	75.436	24.519
Trinidad and Tobago	-.237	0	31.057	12.666
Uruguay	-.43	-.008	75.436	31.877
Venezuela	-.319	0	46.837	34.644

Table A-53: Counterfactual Change in Food Prices - North and Central America

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Bahamas	-.237	0	31.065	13.332
Belize	-.237	0	31.065	16.946
Canada	-.022	-.007	2.25	9.276
Costa Rica	-.237	0	31.057	14.191
Dominican Republic	-.237	0	31.058	16.654
El Salvador	-.237	0	31.067	23.574
Guatemala	-.237	-.017	31.069	19.503
Haiti	-.237	-.045	31.062	17.404
Jamaica	-.237	0	31.056	13.549
Mexico	-.354	0	54.8	20.538
Nicaragua	-.237	0	31.06	16.92
Panama	-.237	0	31.06	18.066
United States	-.059	.003	6.27	9.093

Table A-54: Counterfactual Change in Food Prices - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Albania	-.086	-.053	9.409	13.132
Armenia	-.226	-.089	29.201	23.965
Austria	-.05	-.029	5.271	6.526
Belarus	.031	.012	-3.007	4.566
Belgium	-.067	-.015	7.18	9.457
Bosnia and Herzegovina	-.086	-.052	9.409	7.077
Bulgaria	-.086	-.054	9.408	9.092
Croatia	-.05	-.044	5.262	4.027
Cyprus	-.078	0	8.46	11.172
Czech Republic	-.05	-.021	5.263	7.433
Denmark	.109	.006	-9.829	4.623
Estonia	.031	.033	-3.007	4.376
Finland	.109	.02	-9.828	-2.956
France	-.067	-.027	7.181	8.486
Georgia	-.226	-.091	29.198	14.232
Germany	-.029	-.021	2.986	9.173
Greece	-.078	-.012	8.461	8.65
Hungary	-.05	-.05	5.263	6.162
Iceland	.109	.009	-9.829	5.09
Ireland	-.039	0	4.056	6.307
Israel	-.27	0	36.987	23.532
Italy	-.074	-.018	7.99	8.789

Table A-55: Counterfactual Change in Food Prices - Europe

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Latvia	.031	.032	-3.007	6.174
Lithuania	.031	.018	-3.002	5.743
Luxembourg	-.05	-.018	5.263	7.257
Macedonia	-.086	-.061	9.405	7.019
Malta	-.074	0	7.991	7.718
Moldova	-.086	-.07	9.416	11.68
Montenegro	-.086	-.03	9.41	7.236
Netherlands	-.07	-.007	7.527	10.548
Norway	.109	-.006	-9.83	3.725
Poland	-.047	-.01	4.931	7.639
Portugal	-.096	-.009	10.619	10.082
Romania	-.066	-.056	7.067	6.077
Russia	-.077	-.006	8.343	10.246
Serbia	-.086	-.075	9.41	6.876
Slovakia	-.05	-.032	5.263	6.094
Slovenia	-.05	-.039	5.268	7.708
Spain	-.089	-.012	9.768	10.424
Sweden	.109	.011	-9.828	.349
Switzerland	-.05	-.031	5.259	8.984
Turkey	-.162	-.038	19.33	14.43
Ukraine	-.052	-.038	5.492	9.617
United Kingdom	-.039	-.002	4.057	7.879

Table A-56: Counterfactual Change in Food Prices - Western Pacific and Oceania

Country	Ag Productivity Change	Manufacturing Productivity Change	Autarky	With Trade
Australia	-.266	0	36.24	24.799
Fiji	.022	-.004	-2.152	12.003
Indonesia	-.179	0	21.805	20.004
Malaysia	-.225	0	29.027	18.33
New Zealand	.022	0	-2.152	11.207
Singapore	-.225	0	29.03	15.615

Table A-57: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Sub-Saharan Africa

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Angola	-.019	-.018	.003
Benin	-.078	-.069	.009
Botswana	-.116	-.087	.003
Burkina Faso	-.07	-.063	-.038
Cameroon	-.055	-.052	.011
Cape Verde	-.12	-.046	-.023
Central African Republic	-.428	-.356	-.037
Chad	-.25	-.226	-.032
Comoros	-.102	-.065	-.03
Congo	-.225	-.079	.009
Cote d'Ivoire	-.025	-.025	-.016
Democratic Republic of Congo	-.131	-.129	-.12
Ethiopia	-.171	-.169	-.091
Gabon	-.15	-.069	.001
Gambia	-.072	-.133	-.091
Ghana	-.018	-.014	-.017
Kenya	-.035	-.031	-.045
Lesotho	-.147	-.072	-.085
Madagascar	-.153	-.146	-.073

Table A-58: Equivalent Variation Willingness-to-Pay (Share of GDP) Alternative Trade Cost Cases - Sub-Saharan Africa

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Malawi	-.225	-.225	-.119
Mali	-.115	-.106	-.005
Mauritania	-.126	-.101	.003
Mauritius	-.013	-.007	-.01
Mozambique	-.143	-.147	-.074
Namibia	-.076	-.044	-.003
Niger	-.163	-.121	-.056
Nigeria	-.013	-.012	-.006
Rwanda	-.434	-.387	-.086
Senegal	-.155	-.08	-.046
Seychelles	-.014	-.023	.038
Sierra Leone	-.13	-.164	-.105
Somalia	-.126	-.123	-.103
South Africa	-.022	-.019	-.008
Sudan	-.101	-.087	-.024
Swaziland	-.172	-.121	-.013
Tanzania	-.112	-.109	-.061
Togo	-.099	-.125	-.066
Uganda	-.096	-.085	-.049
Zambia	-.208	-.199	-.001
Zimbabwe	-.223	-.212	-.074

Table A-59: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Middle East and North Africa

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Algeria	-.031	-.028	-.019
Bahrain	-.007	-.008	-.001
Egypt	.008	.009	0
Iran	-.018	-.016	-.004
Iraq	-.04	-.031	-.004
Jordan	-.028	-.026	-.026
Kuwait	-.007	-.01	0
Lebanon	-.014	-.015	-.022
Libya	-.008	-.005	-.009
Morocco	-.08	-.076	-.029
Oman	-.007	-.007	-.009
Qatar	-.003	-.003	-.002
Saudi Arabia	-.006	-.005	-.001
Syria	-.065	-.063	-.047
Tunisia	-.041	-.035	-.048
United Arab Emirates	-.003	-.002	-.002
Yemen	-.063	-.056	-.028

Table A-60: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Asia

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Afghanistan	-.133	-.134	-.075
Azerbaijan	-.057	-.053	-.035
Bangladesh	-.062	-.058	-.022
Bhutan	-.208	-.173	.013
Brunei	-.004	-.004	.005
Cambodia	-.07	-.057	.002
China	-.04	-.04	-.04
Hong Kong	-.001	-.001	-.004
India	-.085	-.082	-.013
Japan	-.011	-.008	-.01
Kazakhstan	-.034	-.035	-.038
Kyrgyzstan	-.058	-.047	-.052
Maldives	-.012	-.017	-.035
Myanmar	-.109	-.105	.002
Nepal	-.092	-.084	-.064
Pakistan	-.045	-.044	-.034
Philippines	-.038	-.038	-.006
South Korea	-.021	-.024	-.017
Sri Lanka	-.027	-.03	0
Tajikistan	-.078	-.074	-.119
Thailand	-.036	-.028	-.002
Uzbekistan	-.061	-.057	-.049
Vietnam	-.028	-.029	-.01

Table A-61: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - South America

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Argentina	-.002	0	-.006
Barbados	-.006	-.029	-.01
Bolivia	-.108	-.095	-.036
Brazil	-.01	-.008	-.006
Chile	-.02	-.019	-.01
Colombia	-.015	-.013	-.006
Ecuador	-.028	-.031	-.003
Honduras	-.041	-.032	-.011
Paraguay	-.057	-.037	-.03
Peru	-.04	-.04	-.008
Suriname	-.024	-.008	-.015
Trinidad and Tobago	-.013	-.012	.004
Uruguay	-.034	-.028	-.014
Venezuela	-.013	-.012	-.001

Table A-62: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - North and Central America

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Bahamas	-.008	-.002	-.029
Belize	-.021	-.024	-.001
Canada	-.016	-.014	-.016
Costa Rica	-.011	-.004	-.021
Dominican Republic	-.017	-.022	.004
El Salvador	-.023	-.015	-.021
Guatemala	-.04	-.036	-.019
Haiti	-.093	-.091	-.048
Jamaica	-.024	-.024	.02
Mexico	-.028	-.015	-.003
Nicaragua	-.039	-.031	-.021
Panama	-.017	-.018	-.002
United States	0	.001	0

Table A-63: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Europe

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Albania	-.047	-.04	-.047
Armenia	-.095	-.087	-.077
Austria	-.029	-.026	-.031
Belarus	-.011	-.011	-.007
Belgium	-.017	-.016	-.015
Bosnia and Herzegovina	-.049	-.049	-.065
Bulgaria	-.047	-.045	-.033
Croatia	-.038	-.041	-.03
Cyprus	-.001	.002	.003
Czech Republic	-.029	-.024	-.027
Denmark	0	.005	.005
Estonia	.008	.002	.02
Finland	.003	-.001	.006
France	-.025	-.024	-.021
Georgia	-.098	-.084	-.075
Germany	-.023	-.023	-.021
Greece	-.011	-.015	-.007
Hungary	-.045	-.045	-.035
Iceland	.005	.014	.009
Ireland	-.001	-.002	-.003
Israel	-.005	-.009	-.012
Italy	-.016	-.016	-.015

Table A-64: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Europe

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Latvia	.006	.008	.005
Lithuania	-.004	-.003	-.015
Luxembourg	-.015	-.017	-.008
Macedonia	-.055	-.049	-.053
Malta	-.001	-.005	.006
Moldova	-.067	-.048	-.05
Montenegro	-.03	-.032	-.016
Netherlands	-.009	-.012	-.006
Norway	-.003	-.002	-.004
Poland	-.022	-.02	-.025
Portugal	-.01	-.011	-.007
Romania	-.051	-.053	-.042
Russia	-.024	-.023	-.025
Serbia	-.064	-.057	-.071
Slovakia	-.036	-.035	-.031
Slovenia	-.036	-.033	-.044
Spain	-.011	-.008	-.007
Sweden	-.002	-.005	.004
Switzerland	-.028	-.032	-.022
Turkey	-.036	-.036	-.03
Ukraine	-.042	-.037	-.046
United Kingdom	-.004	-.003	0

Table A-65: Equivalent Variation Willingness-to-Pay (Share of GDP)
Alternative Trade Cost Cases - Western Pacific and Oceania

Country	Autarky	Estimated Trade Cost Case	Low Trade Cost Case
Australia	-.004	-.005	-.003
Fiji	.001	-.007	.008
Indonesia	-.023	-.018	-.006
Malaysia	-.012	-.012	-.001
New Zealand	0	.009	.002
Singapore	-.005	-.009	0

Table A-66: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Sub-Saharan Africa

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Angola	2.09	.05	.034	.033	-.029	-.008
Benin	1.63	.156	.121	.074	-.131	-.033
Botswana	3.28	.062	.034	.019	-.087	-.035
Burkina Faso	1.2	.24	.187	.23	-.09	-.075
Cameroon	1.68	.216	.145	.166	-.106	-.051
Cape Verde	1.43	.163	.083	.06	-.05	-.084
Central African Republic	1.47	.299	.287	.094	-.436	-.316
Chad	1.13	.257	.213	.243	-.226	-.221
Comoros	1.34	.083	.053	.006	-.081	.088
Congo	2.73	.092	.046	.005	-.093	-.046
Cote d'Ivoire	1.39	.17	.132	.14	-.064	-.033
Democratic Republic of Congo	10.19	.421	.151	.169	-.159	-.027
Ethiopia	1.23	.359	.333	.376	-.19	-.182
Gabon	1.23	.073	.065	.006	-.069	-.04
Gambia	1.24	.008	.074	.027	-.206	-.133
Ghana	1.65	.181	.143	.164	-.024	-.011
Kenya	2.59	.156	.102	.109	-.049	-.017
Lesotho	1.56	.107	.045	.03	-.081	-.121
Madagascar	1.8	.404	.299	.275	-.154	-.127

Table A-67: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Sub-Saharan Africa

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Malawi	2.84	.436	.309	.357	-.244	-.167
Mali	1.81	.197	.147	.146	-.126	-.067
Mauritania	1.98	.249	.157	.149	-.101	-.074
Mauritius	1.69	.038	.017	.028	-.009	-.008
Mozambique	1.92	.367	.267	.305	-.169	-.13
Namibia	2.29	.099	.06	.027	-.044	-.035
Niger	1.97	.325	.196	.248	-.193	-.125
Nigeria	2.16	.068	.045	.049	-.012	-.008
Rwanda	1.14	.409	.39	.322	-.394	-.366
Senegal	1.24	.153	.1	.047	-.105	-.101
Seychelles	2.03	.016	.031	.004	-.023	.038
Sierra Leone	1.49	.139	.177	.173	-.204	-.146
Somalia	2.27	.334	.201	.219	-.176	-.152
South Africa	2.7	.056	.036	.03	-.033	-.012
Sudan	1.67	.12	.092	.068	-.087	-.065
Swaziland	1.92	.102	.058	.034	-.135	-.091
Tanzania	1.19	.298	.239	.318	-.131	-.123
Togo	1.26	.316	.277	.302	-.183	-.149
Uganda	1.25	.416	.338	.404	-.109	-.115
Zambia	1.52	.36	.28	.318	-.233	-.181
Zimbabwe	4.17	.302	.122	.14	-.248	-.111

Table A-68: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Middle East and North Africa

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Algeria	1.96	.034	.022	.015	-.068	-.029
Bahrain	2	.015	.013	.006	-.008	-.008
Egypt	2.08	.077	.049	.07	.009	.006
Iran	1.62	.051	.038	.042	-.067	-.019
Iraq	2.9	.05	.023	.019	-.031	-.018
Jordan	2.53	.038	.032	.025	-.055	-.007
Kuwait	1.19	.005	.003	.004	-.01	-.005
Lebanon	1.91	.079	.058	.056	-.034	-.007
Libya	2.07	.067	.043	.048	-.026	-.01
Morocco	1.44	.1	.086	.048	-.093	-.071
Oman	1.14	.023	.018	.015	-.007	-.01
Qatar	1.92	.008	.007	.005	-.003	0
Saudi Arabia	1.43	.016	.011	.011	-.005	-.004
Syria	1.54	.09	.058	.058	-.106	-.094
Tunisia	2.77	.055	.037	.021	-.068	-.016
United Arab Emirates	1.97	.016	.011	.013	-.002	-.001
Yemen	7.17	.11	.04	.04	-.086	-.018

Table A-69: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Asia

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Afghanistan	2.19	.278	.186	.218	-.158	-.134
Azerbaijan	2.23	.088	.055	.054	-.075	-.054
Bangladesh	1.52	.189	.144	.165	-.093	-.058
Bhutan	3.4	.267	.134	.075	-.179	-.076
Brunei	.71	.027	.006	.023	-.004	-.007
Cambodia	4.62	.19	.092	.085	-.064	-.019
China	3.48	.064	.034	.038	-.065	-.03
Hong Kong	1.92	.018	.012	.013	-.001	-.002
India	3.24	.161	.087	.106	-.082	-.045
Japan	1.72	.012	.008	.009	-.017	-.007
Kazakhstan	2.23	.088	.056	.06	-.05	-.044
Kyrgyzstan	10.26	.156	.041	.062	-.062	-.069
Maldives	1.77	.021	.03	.073	-.023	-.04
Myanmar	2.13	.209	.144	.179	-.111	-.071
Nepal	1.52	.283	.223	.277	-.109	-.103
Pakistan	1.69	.133	.103	.111	-.06	-.031
Philippines	1.99	.105	.07	.064	-.038	-.022
South Korea	1.54	.012	.007	.009	-.035	-.022
Sri Lanka	1.28	.109	.105	.099	-.03	-.02
Tajikistan	10.58	.232	.075	.081	-.101	-.06
Thailand	2.61	.072	.045	.033	-.028	-.014
Uzbekistan	9.37	.108	.04	.044	-.092	-.005
Vietnam	4.55	.158	.078	.082	-.049	-.014

Table A-70: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - South America

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counter-factual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Argentina	2.1	.065	.051	.058	-.009	-.002
Barbados	1.19	.035	.055	.047	-.029	-.008
Bolivia	1.79	.125	.096	.071	-.132	-.105
Brazil	2.46	.063	.041	.043	-.008	-.009
Chile	1.89	.053	.035	.029	-.022	-.014
Colombia	1.92	.066	.047	.051	-.031	-.011
Ecuador	1.79	.098	.075	.067	-.043	-.021
Honduras	1.52	.082	.061	.045	-.072	-.04
Paraguay	2	.151	.079	.062	-.048	-.046
Peru	1.91	.104	.073	.063	-.064	-.029
Suriname	2.35	.028	.019	.006	-.008	-.019
Trinidad and Tobago	1.85	.004	.009	.002	-.012	-.004
Uruguay	2.25	.091	.062	.018	-.031	-.014
Venezuela	2.33	.027	.017	.016	-.012	-.005

Table A-71: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - North and Central America

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counter-factual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Bahamas	1.84	.02	.006	.011	-.002	-.024
Belize	1.8	.104	.083	.097	-.024	.009
Canada	1.63	.019	.014	.02	-.02	-.02
Costa Rica	1.38	.021	.015	.015	-.004	-.017
Dominican Republic	1.74	.046	.035	.027	-.022	-.008
El Salvador	1.92	.108	.059	.064	-.015	-.02
Guatemala	1.69	.151	.118	.116	-.057	-.035
Haiti	2.45	.161	.12	.134	-.123	-.067
Jamaica	3.38	.074	.041	.033	-.024	-.006
Mexico	2.62	.031	.019	.009	-.015	-.009
Nicaragua	1.9	.11	.061	.079	-.061	-.028
Panama	1.74	.06	.048	.035	-.018	-.009
United States	1.96	.023	.017	.022	-.011	0

Table A-72: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Europe

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Albania	2.08	.129	.077	.058	-.056	-.049
Armenia	4.46	.104	.047	.05	-.106	-.074
Austria	1.6	.012	.008	.01	-.039	-.039
Belarus	2.2	.05	.033	.039	-.016	-.023
Belgium	1.69	.02	.012	.015	-.02	-.019
Bosnia and Herzegovina	3.79	.047	.025	.039	-.061	-.05
Bulgaria	1.86	.06	.041	.046	-.06	-.045
Croatia	2.16	.042	.029	.037	-.06	-.038
Cyprus	2.66	.06	.033	.046	.002	0
Czech Republic	1.61	.018	.011	.013	-.036	-.038
Denmark	1.76	.033	.025	.039	.005	.001
Estonia	1.95	.036	.036	.029	.001	.01
Finland	1.81	.013	.009	.014	-.002	-.002
France	1.81	.026	.019	.025	-.035	-.024
Georgia	4.68	.086	.032	.031	-.1	-.06
Germany	1.75	.011	.007	.01	-.032	-.026
Greece	1.89	.037	.028	.028	-.033	-.008
Hungary	2.4	.028	.02	.023	-.07	-.045
Iceland	2.2	.063	.059	.075	.015	.008
Ireland	1.7	.002	.003	.003	-.002	-.002
Israel	2.19	.016	.01	.006	-.009	.003
Italy	2.21	.019	.012	.014	-.035	-.014

Table A-73: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Europe

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counter-factual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Latvia	2.37	.037	.034	.045	.006	-.008
Lithuania	2.04	.035	.029	.031	-.006	-.021
Luxembourg	2.18	.013	.01	.011	-.026	-.008
Macedonia	3.82	.104	.056	.064	-.064	-.046
Malta	2.56	.006	.009	.015	-.013	.003
Moldova	8.94	.169	.036	.062	-.06	-.052
Montenegro	2.06	.054	.041	.046	-.041	-.032
Netherlands	2.1	.027	.019	.022	-.014	-.005
Norway	1.71	.02	.012	.026	-.002	-.002
Poland	2.3	.031	.019	.024	-.028	-.029
Portugal	2.73	.026	.015	.018	-.014	.001
Romania	2.81	.06	.036	.041	-.072	-.045
Russia	2.16	.03	.02	.023	-.028	-.034
Serbia	3.55	.069	.037	.045	-.081	-.06
Slovakia	1.71	.022	.014	.015	-.05	-.047
Slovenia	1.5	.017	.01	.014	-.048	-.054
Spain	1.88	.023	.016	.018	-.023	-.003
Sweden	1.7	.009	.007	.011	-.005	-.002
Switzerland	2.02	.003	.002	.004	-.043	-.024
Turkey	1.98	.047	.034	.038	-.053	-.038
Ukraine	4.38	.105	.054	.069	-.045	-.038
United Kingdom	2.12	.015	.01	.014	-.003	-.001

Table A-74: Equivalent Variation Willingness-to-Pay (Share of GDP) Accounting for Economic Growth and Adaptation Costs and Benefits - Western Pacific and Oceania

Country	Projected GDP Per-Capita 2080 / Present	Ag GDP Share Baseline	Ag GDP Share 2080 Baseline	Ag GDP Share 2080 Counterfactual	EV WTP Losses from Present Baseline	EV WTP Losses from 2080 Baseline
Australia	1.5	.036	.03	.022	-.005	0
Fiji	3.79	.126	.08	.101	-.012	-.003
Indonesia	4.84	.112	.05	.057	-.018	-.012
Malaysia	2	.042	.031	.028	-.012	-.002
New Zealand	2	.078	.057	.08	.009	.003
Singapore	1.6	.007	.006	.005	-.009	-.003

Table A-75: Country-Level Panel Regression

	(1)	(2)	(3)	(4)
	log(GDP)	Food Share of Imports	Ag Share of GDP	Ag Labor Share
KDD X 100	-0.0223	0.00638	0.0165	0.00483
	(-0.55)	(1.80)	(3.92)	(3.14)
GDD X 100	0.00251	-0.00191	-0.00165	-0.00113
	(0.44)	(-2.87)	(-1.53)	(-1.74)
Observations	7561	5775	5522	3718
Country FE	X	X	X	X
Year FE	X	X	X	X
Ag Labor Weights				

Notes: t-statistics in parentheses. Reported Driscoll and Kraay (1998) standard errors are robust to heteroskedasticity, spatial correlation, and autocorrelation of up to 5 lags. Results come from estimating Equation 24 with crop-area weighted growing and killing degree days. Data covers 164 countries from 1960-2012 with varying coverage by country and outcome variable. Economic data from all sources above are retrieved from the World Bank Databank.

Appendix G: Additional RPS Tables and Figures

Figure A-35: REC Tracking Markets

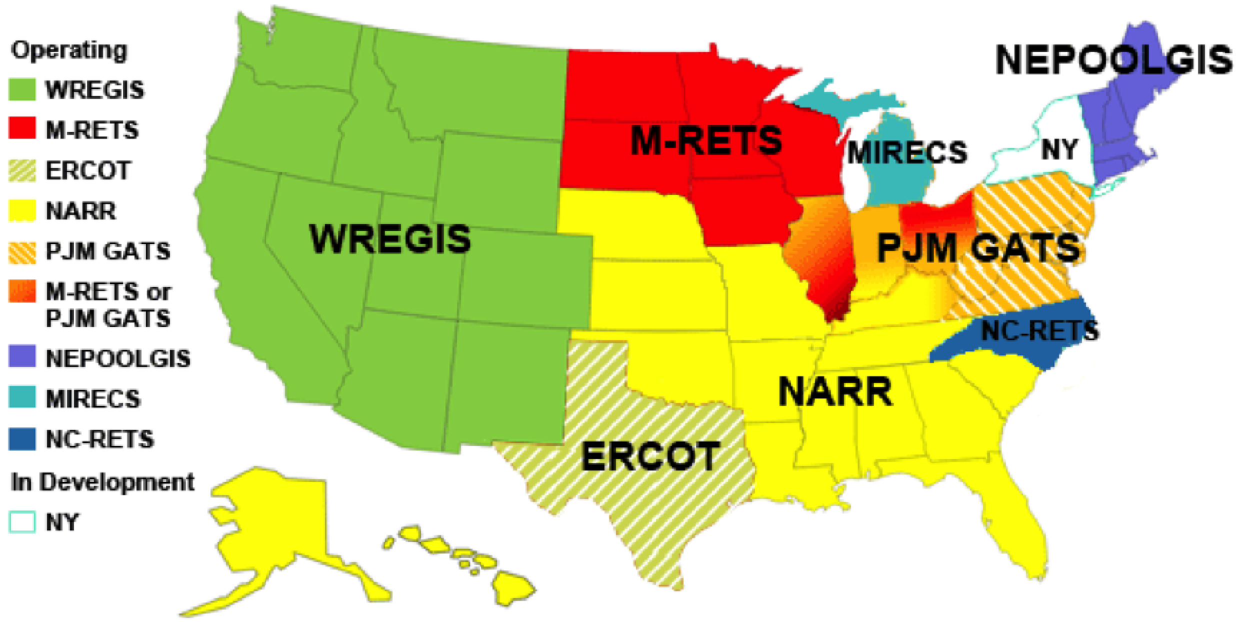

Notes: Map comes from the EPA, and shows the regions within which Renewable Energy Credits are tradable for compliance. Renewable Energy Credits may in principle be traded across regions, but most RPS compliance occurs within region.

Figure A-36: Implementation of Energy Programs by State and Year

Year	RPS	Public Benefit Fund	Net Metering	Green Power Purchasing	Energy Efficiency	Re-structure
1990	0	0	5	0	0	0
1991	1	0	6	0	0	0
1992	1	0	6	0	0	0
1993	1	0	9	0	0	0
1994	1	0	10	0	0	0
1995	1	0	10	0	0	0
1996	1	0	11	0	0	5
1997	2	1	14	0	0	11
1998	3	4	17	0	1	13
1999	7	9	21	0	3	21
2000	7	11	25	0	3	23
2001	8	13	28	0	4	24
2002	10	14	30	1	4	24
2003	10	15	31	2	4	24
2004	17	15	33	3	5	24
2005	20	17	33	4	7	24
2006	21	18	35	4	9	24
2007	26	18	36	6	13	24
2008	29	18	39	6	23	24
2009	30	18	43	7	24	24
2010	30	18	44	7	28	24
2011	30	18	44	7	28	24
2012	30	19	44	8	28	24
2013	30	20	44	8	28	24
2014	30	20	44	8	28	24
2015	30	20	44	8	28	24

Notes: Data on other programs comes from DSIRE. Data on restructuring dates comes from Fabrizio, Rose and Wolfram (2007).

Figure A-37: Electricity Prices Before and After RPS Passage, by Sector

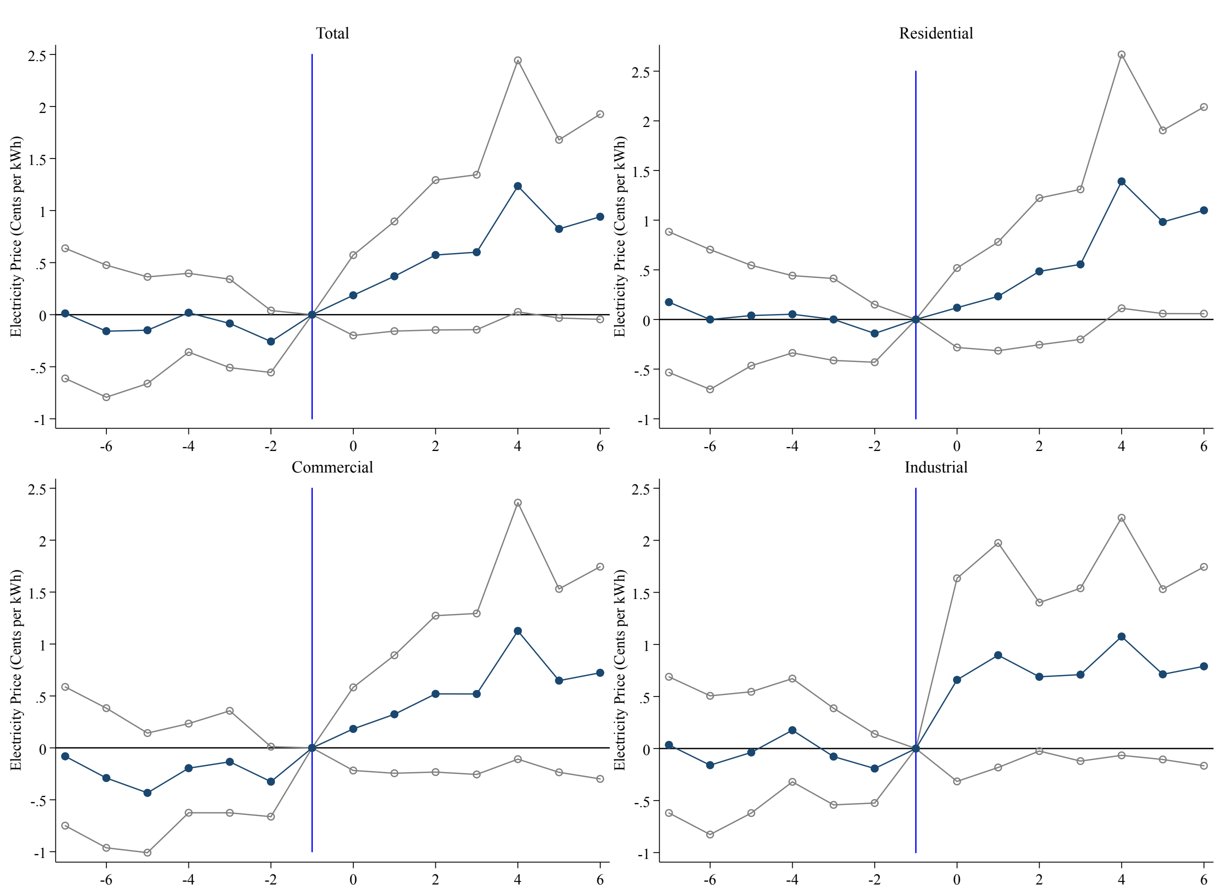

Notes: Graphs show coefficients for σ_τ for τ = -7 to τ = 6 from the event study specification in Equation (7). This specification regresses the dependent variable - retail electricity prices - on indicator variables for years relative to program passage, controlling for state, year, and other programs fixed effects. Blue lines show the point estimates and gray lines contain the 95% confidence interval. Electricity price data are from the EIA. Standard errors are clustered at the state-level.

Figure A-38: Estimated Effects of RPS Programs on Net Renewable Requirements (Extended Post Period)

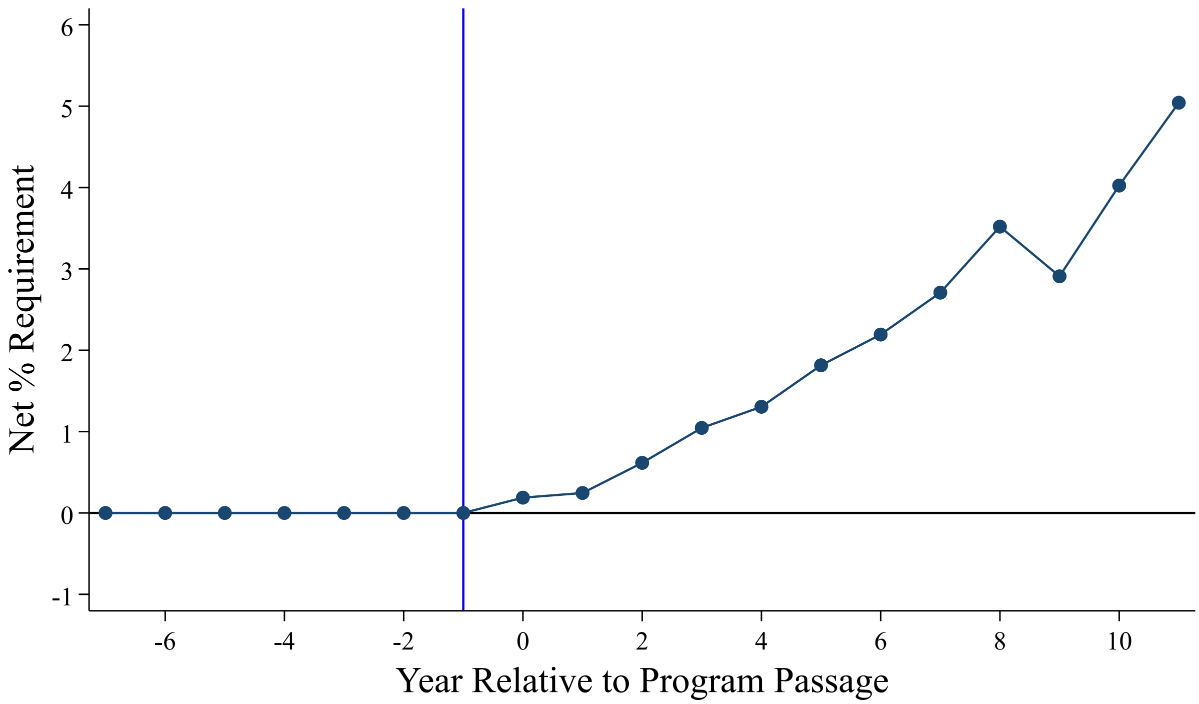

Notes: Graph shows the mean net RPS requirement percentage for event years τ = -7 to τ = 11. Gross RPS requirement data are from the LBNL. RPS program passage dates and requirements are from the Department of Energy and state government websites.

Figure A-39: Estimated Effects of RPS Programs on Gross Renewable Requirements (Extended Post Period)

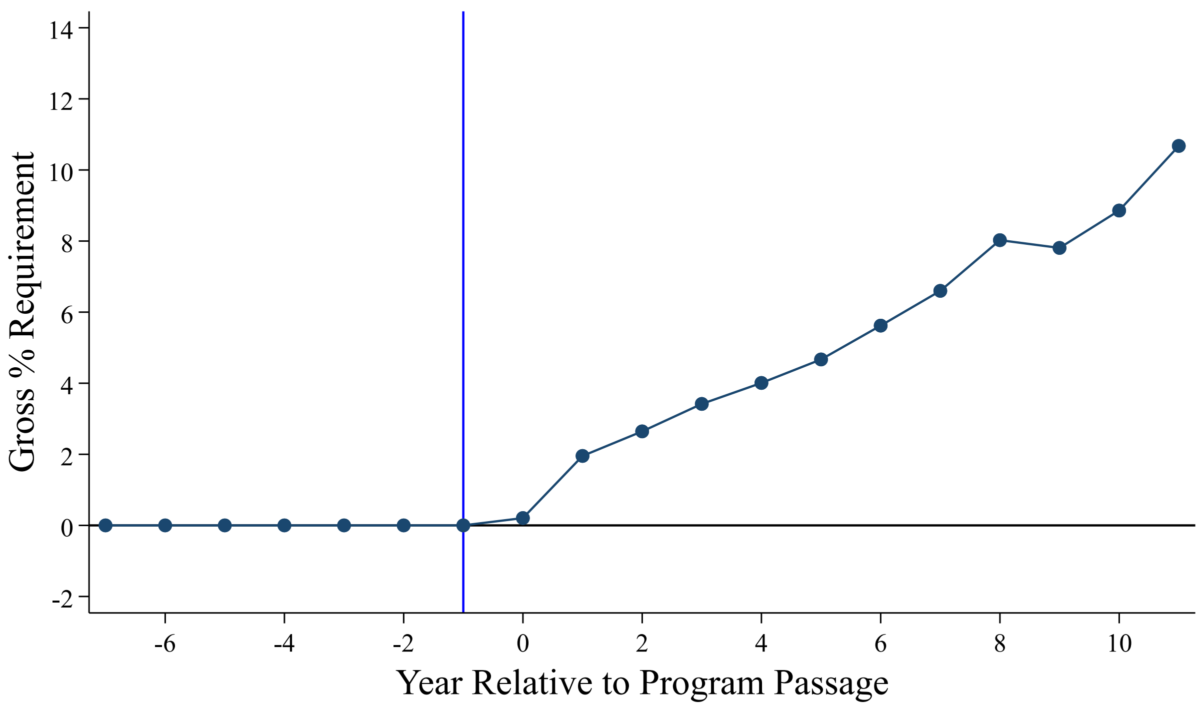

Notes: Graph shows the mean gross RPS requirement percentage for event years τ = -7 to τ = 11. Gross RPS requirement data are from the LBNL. RPS program passage dates and requirements are from the Department of Energy and state government websites.

Figure A-40: Estimated Effects of RPS Programs on Retail Electricity Prices (Extended Post Period)

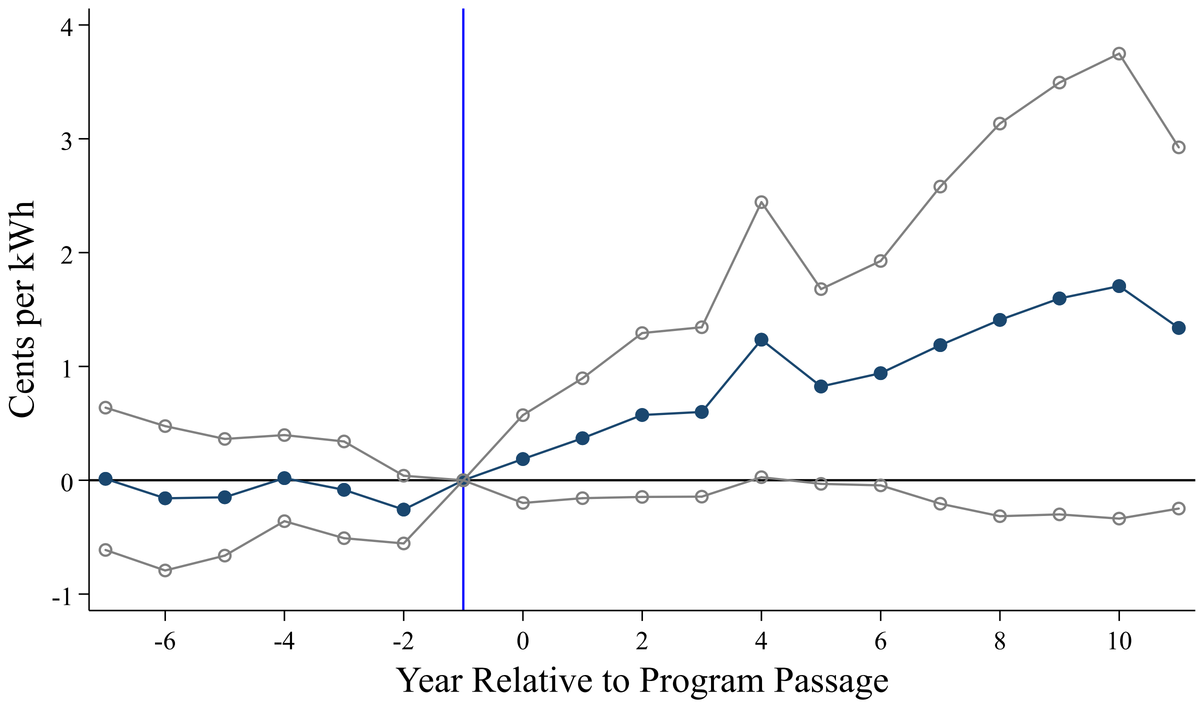

Notes: Graph shows coefficients for σ_τ for τ = -7 to τ = 11 from the event study specification in Equation (7) for retail electricity prices on indicator variables for years relative to program passage, controlling for state, year, and other programs fixed effects. Blue lines show the point estimates and gray lines contain the 95% confidence interval. Electricity price data and electricity generation for calculating net requirement are from the EIA. RPS program passage dates and requirements are from the Department of Energy and state government websites. Standard errors are clustered at the state-level.

Figure A-41: CO_2 Emissions Intensity Before and After RPS Passage

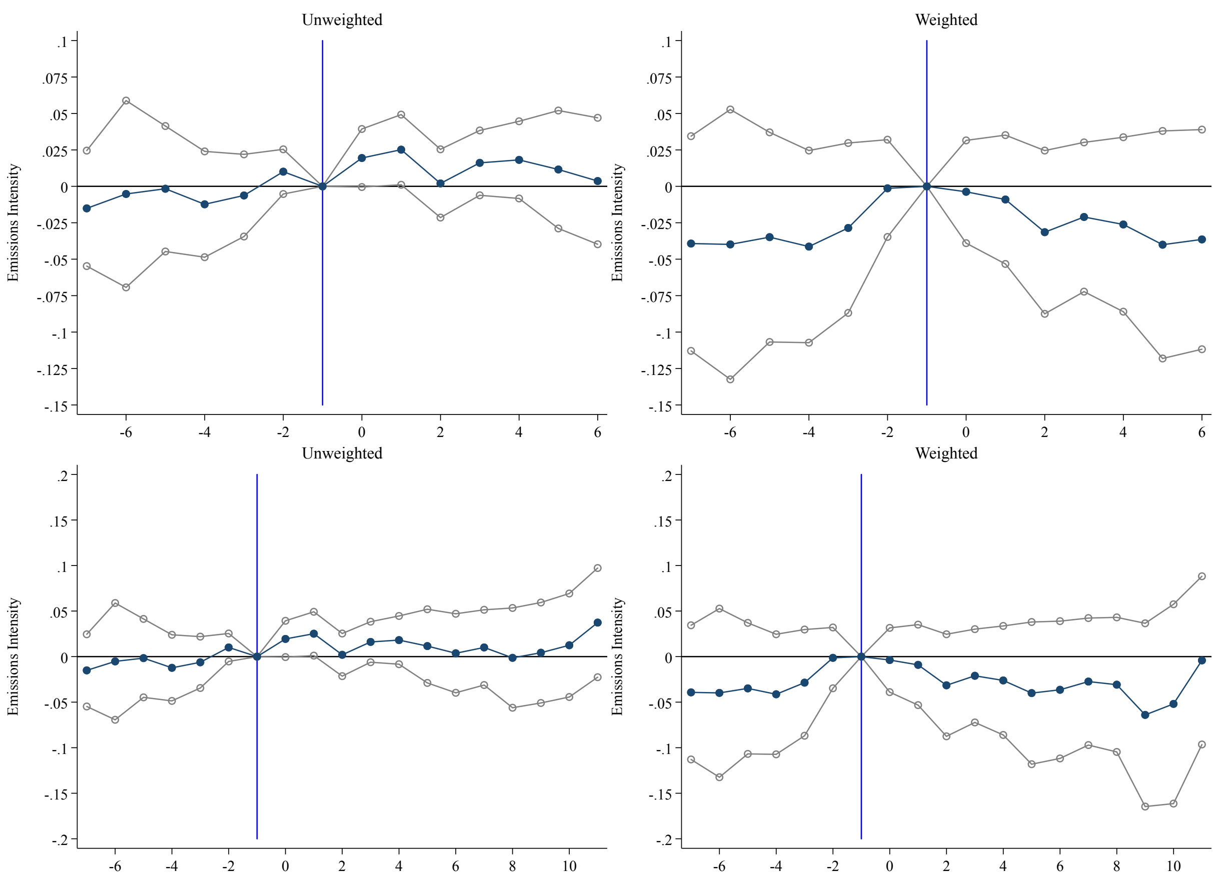

Notes: Graphs show coefficients for σ_τ from the event study specification in Equation (7). This specification regresses the dependent variable - CO2 emissions intensity - on indicator variables for years relative to program passage, controlling for REC regions and year fixed effects, as well as other programs fixed effects whose values are a generation-weighted average of the states' indicator values within a given REC region. The plots labelled "Weighted" use state-count weights, and the ones labelled "Unweighted" do not. The top two plots show a narrower time frame, from $\tau = -7$ to $\tau = 6$, where we have a balanced panel of 29 states. The bottom two plots show a larger time frame in which we have an unbalanced panel that varies from 29 to 16 states. Blue lines show the point estimates and gray lines contain the 95% confidence interval. Standard errors are clustered at the REC region level.