

© 2021 Noé J Nava

THREE ESSAYS ON AGRICULTURAL AND CONSUMER ECONOMICS

BY

NOÉ J NAVA

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Agricultural and Applied Economics
in the Graduate College of the
University of Illinois Urbana-Champaign, 2021

Urbana, Illinois

Doctoral Committee:

Assistant Professor William Ridley, Chair and Director of Research
Professor Sandy Dall’Erba
Associate Professor Yilan Xu
Associate Professor Rafael Garduño Rivera

Abstract

In this dissertation, I employ different methodologies to study current issues that affect the food system in the U.S. and México. From farm to fork, this research explores problems of obesity derived from the consumption of sugary soft beverages, the impacts of supply chain disruptions on the availability of food, and the role of climate change in agriculture. I contribute to a broad literature in agricultural economics by expanding the set of methodologies available to researchers and providing new insights in the existing literature. This dissertation is consituted of three chapters as follows.

The theoretical foundations and procedures to estimate demand systems to recover income- and price-elasticities have long been established; however, some streams of the literature do not account for all of the empirical implications of their datasets, leading to inconclusive results. The analysis of sin taxes on sugary soft beverages (SSB) is an exemplary case that has carried controversy in Latin-American countries. In Mexico, for instance, the federal government imposed a 20% price hike on SSB to curb its obesity epidemic in 2014, but obesity has continued increasing. This manuscript revisits the demand estimation applications aim to study the effectiveness of the sin tax with a censored Quadratic Almost Ideal Demand (QUAID) system.

Previous studies find that SSB is an elastic good (a 10% price increase is associated with a greater than 10% reduction in quantity) with an own-price elasticity between -1.11 and -1.61, but our study finds that SSB is an inelastic good (a 10% price increase is associated with a 9.5% decrease in quantity). Our results suggest that SSB taxes are not as effective as previous research indicates, explaining the gap between actual obesity rates and expected obesity rates after the sin tax was imposed in Mexico.

This paper expands a Ricardian trade general equilibrium model to assess the U.S. state vulnerabilities to production shocks and supply chain disruptions under the existing food system. We find that the distribution of welfare losses caused by disruption in the existing supply chain depends on the state's agricultural capacity and its operating expenses. Conversely, the existing supply chain capability to mitigate the impacts from productivity shocks is hindered by reliance on in-state crop production. Our extension of the Ricardian trade formulation is a comprehensive but simple model of the U.S. market of crops. We incorporate the heterogeneity of farmers' production and isolate food processing plants as the main consumer of domestic crop production. Our general equilibrium model allows us to derive conclusions that are state specific, so they can be employed to craft state level policies to strengthen the national food system.

The Ricardian framework has been widely used to study the impact of climate change on agriculture across US counties over the past few decades. While the spatial heterogeneity of climate change is well-accepted, the literature struggles to reach an agreement on how to model it, hence leading to a wide range of forecasted impacts. This paper employs Multiscale Geographically-Weighted-Regression

(MGWR) to avoid setting an a priori definition of heterogeneity and to generate county-specific marginal effects of climate change impacts. This manuscript tests the predictive power of our MGWR application with other functional forms found in the literature on a homogenized dataset of historical climate, demographic and soil quality controls. Our cross-validation exercise indicates that our MGWR approach has a higher predictive power than studies that cluster spatial units, and that other approaches have a downward bias. We attribute the divergence in results to unspecified heterogeneity. Our place-specific marginal effects will help guide the development of place-tailored mitigation and adaptation strategies to climate change.

Para mamá, papá y mis hermanas.

Acknowledgments

At the University of Illinois, I found great support to specialize my interests. I thank my advisor, William Ridley, for his excellent mentorship. Regardless that my interests did not always align with his expertise, Bill always provided thoughtful feedback to each one of these thesis chapters. I can only imagine the countless hours spent on my work. I also thank Sandy Dall’Erba for always pushing me to keep working, keep meeting and keep moving forward. My dissertation would not have been completed on time if he had given up on me. I deeply appreciate the feedback and comments from Rafael Garduño-Rivera and Yilan Xu, and specially their willingness to meet twice this semester to listen to me talk about these chapters for hours. I also want to take advantage of this space to say that the second-year paper class was instrumental in my growth and how I see research. Phil Garcia and Carl Nelson taught me that research is a collective process: We stand on the shoulders of others. Thus, I thank their comments and those from my classmates during this class. Why would you even read and comment every single one of my drafts? Because of you, I have the confidence to call myself an economist.

This dissertation, and my decision to pursue graduate school, would not have been possible without the people who believed in me and encouraged me to move

forward: My economics professors at California State University, San Bernardino. Mayo Toruño taught me the value of seeing standard economics in a different light to not only understand the world, but to make it better. His dedication to teaching made me change mathematics for economics. I remember when Thomas Pierce stopped me after intermediate macroeconomics to ask me about my intentions to go to graduate school; because of him, I decided to pursue a Ph.D. Kazim Konyar introduced me to a large array of quantitative methods in economics including general equilibrium models and structural econometrics, and now his lessons echo in my work. It was during his classes that I discovered what type of economist I wanted to be. We are proud coyotes!

I made it through my Ph.D. for the ones that bring me joy during this journey. During my time in Chambana, I made many friends with whom I shared enjoyable memories at Murphy's, Soma, Britt's house, game nights at Ian's , Legends, downtown Champaign, downtown Urbana, and countless other places. My friends in DC, you know you are part of this too, and you were there when I started this journey. My friends in California, I miss you dearly. My friends in Méxcio, you make me feel like I never left. I found balance in my life because of you.

But most importantly, this dissertation is for my father, my mother and my sisters: Mamá, a ti que dejaste de trabajar, tomaste el rol de madre cuando nací y tu vida fue de mis logros y de mis derrotas. Lo logramos. Papá, tu simplicidad y dedicación a tu familia son ejemplares. Eres mi modelo a seguir y la persona que más admiro. Lucy, por todas esas historias que tenemos juntos desde el Nintendo hasta la Universidad. Diana, llegaste tarde al grupo, pero llegaste llenándonos de alegría. Ustedes son mis

raíces y son buenas raíces.

I did not make this dissertation. We all made it. Thanks.

TABLE OF CONTENTS

CHAPTER 1 THE IMPACT OF TAXING SUGARY SOFT BEVER-	
AGES IN MEXICO: A CENSORED QUAI DEMAND SYSTEM AP-	
PROACH	1
1.1 Introduction	2
1.2 National Household Survey of Expenditures and Incomes	7
1.3 Empirical strategy	9
1.4 Discussion of demand elasticities	24
1.5 Conclusion	30
1.6 Chapter 1 References	31
CHAPTER 2 A MODEL OF THE U.S. FOOD SYSTEM: WHAT ARE	
THE DETERMINANTS OF THE STATE VULNERABILITIES TO	
PRODUCTION SHOCKS AND SUPPLY CHAIN DISRUPTIONS?	36
2.1 Introduction	36
2.2 A Ricardian trade model of U.S. domestic crop production	40
2.3 Empirical Strategy	46
2.4 Food vulnerability simulation results	56
2.5 Conclusions	64
2.6 Chapter 2 References	65
CHAPTER 3 REVISITING THE IMPACT OF CLIMATE CHANGE ON	
AGRICULTURE THROUGH SPATIALLY-VARYING AND PLACE-	
TAILORED RICARDIAN ESTIMATES	69
3.1 Introduction	69
3.2 Spatial non-stationary Ricardian-approach	72
3.3 Data	79
3.4 Results	82
3.5 Conclusion	91
3.6 Chapter 3 References	93

APPENDIX A	SUPPLEMENTARY MATERIAL FOR CHAPTER 1	97
A.1	Households exhibit rank 3 Engel curve behavior	97
A.2	Item aggregation from ENIGH 2018	98
A.3	Marshallian elasticities across methodologies	99
A.4	Marshallian elasticities across years	100
APPENDIX B	SUPPLEMENTARY MATERIAL FOR CHAPTER 3	101
B.1	Inference in the MGWR	101

CHAPTER 1

THE IMPACT OF TAXING SUGARY SOFT BEVERAGES IN MEXICO: A CENSORED QUAI DEMAND SYSTEM APPROACH

Abstract: The theoretical foundations and procedures to estimate demand systems to recover income- and price-elasticities have long been established; however, some streams of the literature do not account for all of the empirical implications of their datasets, leading to inconclusive results. The analysis of sin taxes on sugary soft beverages (SSB) is an exemplary case that has carried controversy in Latin-American countries. In Mexico, for instance, the federal government imposed a 20% price hike on SSB to curb its obesity epidemic in 2014, but obesity has continued increasing. This manuscript revisits the demand estimation applications aim to study the effectiveness of the sin tax with a censored Quadratic Almost Ideal Demand (QUAID) system. Previous studies find that SSB is an elastic good (a 10% price increase is associated with a greater than 10% reduction in quantity) with an own-price elasticity between -1.11 and -1.61, but our study finds that SSB is an inelastic good (a 10% price increase is associated with a 9.5% decrease in quantity). Our results suggest that SSB taxes are not as effective as previous research indicates, explaining the gap between actual obesity rates and expected obesity rates after the sin tax was imposed in Mexico.

1.1 Introduction

To curb its obesity epidemic, the federal government of Mexico proposed an excise tax of one peso per liter on the price of sugary soft beverages (SSB) (the equivalent of a 20% price hike in 2012) that was successfully implemented on the first day of 2014 according to the Mexican National Institute of Statistics and Geography (INEGI) (INEGI, 2021). In spite of the legislative efforts, Mexico continued among the top members with the highest percentage of obese population of the Organisation for Economic Co-operation and Development (OECD) countries, and in 2018 the country occupied the first place with one in three adults being obese (OECD, 2021). Inconclusive evidence about the policy-induced behavioral change and health outcome improvements have brought an ongoing controversy among SSB companies that oppose the tax and policy officials that advocate for stricter legislation (Haines, 2017). Despite the lack of consensus surrounding the role that the 2014 sin tax had on SSB consumption, the political process that concluded in the legislative passage of the tax has been praised and paved the road for other Latin American countries to adopt similar measures (James, Lajous, and Reich, 2020). This manuscript revisits the demand system applications that provided with evidence about the effectiveness of the sin tax by addressing the methodological weaknesses in previous studies.

An ideal policy evaluation requires two identical sub-populations that do not interact with each other. One group is exposed to a price hike and the other one is not. Temporal and group differences provide with an accurate policy-induced Average Treatment Effect (ATE) on health outcomes such as body mass index and caloric intake. Experimental approaches, however, are not feasible for the Mexican case due to the uniform implementation of the sin tax and policy anticipation on the part of consumers. Figure 1.1 demonstrates that the price increase affected all of the regions in Mexico at the same time, leaving no place for a treatment group. The temporal dimension is also compromised since civil and non-profit organizations in Mexico promoted the one peso tax to consumers before it was proposed in congress (Donaldson, 2015).

Therefore, the ideal policy experiment can not be performed, but researchers can

circumvent the empirical challenges with structural approaches that aim to model consumer behavior. For example, Colchero, Salgado, Unar-Munguía, Molina, et al. (2015) focus on the passed-through effect of the sugar tax and find that the one peso per liter tax was almost completely absorbed by consumers. In a similar study, Campos-Vazquez and Medina-Cortina (2019) investigate the pass-through effect accounting for competition at the supermarket level, finding significant heterogeneity depending on the type of beverage that ranges from 0.24 to 1.52 pesos. Both studies suggest that consumers face higher prices that can lead to reductions in SSB consumption, but no study is able to provide the magnitude at which demand falls. Colchero, Guerrero-López, et al. (2016) addresses this gap in the literature, exploiting monthly surveys on the food manufacturing industry in Mexico before and after the sugar tax was implemented and find a 7.3% decrease in per capita sales of SSB. The works of Colchero, Salgado, Unar-Munguía, Molina, et al. (2015), Colchero, Guerrero-López, et al. (2016) and Campos-Vazquez and Medina-Cortina (2019) all find that SSB price hikes led to reductions in SSB consumption, using a single product approach. However, price hikes could lead to a decline of SSB consumption, but also to an increase in the consumption of other caloric-intensive goods if both goods are substitutes of each other. Thus, a demand system analysis of the problem has emerged in the literature for a relative complete analysis of SSB.

Demand systems provide with theoretically consistent parameters that are relevant for policy evaluation i.e., own- and cross-price elasticities, but the structural parameters must be recovered from a theoretically consistent approach. The Almost Ideal Demand (AID) system of Deaton and Muellbauer (1980a) is the first model that allows for the estimation of elasticities that are consistent with the behavior described by the Cournot and Engel aggregations.¹ The AID system, however, is limited to consumer behavior that is loglinear in income i.e., elasticities are monotonic across the income spectrum. Later, Banks, Blundell, and Lewbel (1997) expanded the AID

¹Elasticities that are not restricted by the Engel and Cournot aggregations do not reflect rational consumer behavior in general. Consumer choices must reflect substitutability and complementarity across goods with respect to price changes. Consumer preferences are also budget dependent. The restrictions imposed in demand system estimations guarantee that consumer choices reflect substitutability, complementarity and budget dependency.




Figure 1.1: SSB price index across all economic regions in Mexico.

Note: Figure shows SSB price index (index = 100 in January 2010) across all the economic regions in Mexico over the last decade. A one peso per liter of SSB was proposed in 2012 and was successfully implemented on the first day of 2014. The figure demonstrates that price increases affected all of the regions in Mexico at the same time, leaving no place for a treatment group.

Source: Calculations are our own with data from INEGI (2021): Sub-components of the National Consumer Price Index at the economic region level.

system into the Quadratic Almost Ideal Demand (QUAID) system to account for consumer behavior that is quadratic in income. Several studies report employing the AID and QUAID systems to estimate price elasticities associated with SSB and infer conclusions about the impacts of a sugar tax. In Mexico, for instance, Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) estimate the AID system for a basket of drinks that includes SSB, finding that the one peso tax per liter contributed to a decline in demand for SSB of 22.6%. In a similar study for Colombian households, Caro et al. (2017) employed the QUAID system to estimate price elasticities for SSB and conclude that Colombian consumers respond to SSB price hikes by substituting SSB products for dairy and grain products. Guerrero-López, Unar-Munguía, and Colchero (2017) find similar results after estimating both the AID and QUAID systems for Chilean consumers. All of these studies report that SSB is an elastic good.

The works of Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015), Caro et al. (2017) and Guerrero-López, Unar-Munguía, and Colchero (2017) are theoretically rigorous in their attempt to connect consumer behavior with the policy-induced improvements on health outcomes, but they do not address different empirical shortcomings present in their datasets. The work of Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) for the Mexican case reports employing an AID system even though households in their dataset exhibit non-linear Engel curves. On the other hand, the works of Caro et al. (2017) and Guerrero-López, Unar-Munguía, and Colchero (2017) recognize that households in their datasets exhibit quadratic Engel curves, so they employ the QUAID system instead. However, only Caro et al. (2017) reports addressing the additional concern associated with the use of unit-values i.e., the ratio of expenditure to quantity, to proxy for prices. The employment of unit-values to proxy for prices carries endogeneity concerns since unit values not only reflect prices, but also household preferences (Deaton, 1988; Cox and Wohlgenant, 1986). Finally, none of the studies appropriately addresses the zero purchase observations present in their datasets. Caro et al. (2017) report employing a common Heckit correction introduced by Shonkwiler and Yen (1999) while Guerrero-López, Unar-Munguía, and Colchero (2017) report employing a similar approach

first introduced by Heien and Wesseils (1990). Both studies ignored the budget constraint in the observed demand system which leads to incorrect parameter estimates (Dong, Gould, and Kaiser, 2004; Garcíea-Enríquez and Echevarriéa, 2016).

To our knowledge, no study reports estimating a demand system for SSB along with correcting all empirical challenges previously described. Thus, the first contribution of this paper is the estimation of a QUAID system for SSB that addresses the unit-value endogeneity and zero purchase observations to recover own- and cross-price elasticities. We address the quality/price issue with a simple two-step correction proposed in Capacci and Mazzocchi (2011) and also adopted in Caro et al. (2017). The two-step approach requires that households within a municipality face the same price for the basket of goods, so differences in what households pay are driven by preferences. We address the zero purchase observations following Dong, Gould, and Kaiser (2004) by imposing budget constraint (adding-up) to the observed demand system. This manuscript expands the work of Dong, Gould, and Kaiser (2004) for an AID system to a QUAID system for a basket of four goods that includes SSB, juice, milk and an “other” category that encompasses other drinks consumed by households.

Our results find evidence that SSB is inelastic with an own-price elasticity of -0.95 and that SSB is a normal good when the budget elasticity is evaluated at the average of our data with a value of 1.01. Our findings for the own-price elasticity of SSB are substantially lower to the findings in previous studies that report that SSB is an elastic good (Colchero, Salgado, Unar-Munguiéa, Molina, et al., 2015; Guerrero-López, Unar-Munguiéa, and Colchero, 2017; Caro et al., 2017). We attribute the difference to the previously listed empirical complications in previous studies.

Our manuscript is organized as follows. Section 2 describes our dataset construction. We provide evidence that our dataset exhibits the empirical complications previously described, so section 2 also serves as an introduction for our empirical strategy described in section 3. In section 3, we describe the empirical complications present in our data analytically, and then we discuss how we address each issue separately. Our manuscript’s description of data censorship correction focuses on the derivation of the likelihood functions used in our model estimation for the four-item

basket of goods. Section 4 reports our price and budget elasticities, and discusses our results in the context of other demand estimations of SSB for Latin American households. Section 5 offers concluding remarks.

1.2 National Household Survey of Expenditures and Incomes

The National Household Survey of Incomes and Expenditures (ENIGH) is a biannual survey of Mexican households that covers several areas of expenditure collected by INEGI. The survey has expanded in scope to cover detailed demographic and expenditure information of surveyed households since it was first collected in 1984. Household members keep weekly diaries of their daily expenses, and INEGI visits surveyed households to collect their data in the middle of the year (over a period of 10 weeks during the third quarter of the year). Despite its granularity, ENIGH is not longitudinal. Instead, INEGI collects a representative sample of Mexican households and reports expansion factors based on the Mexican population census for each surveyed household. ENIGH has been utilized in several demand estimations (Barquera et al., 2008; Urzua, 2008), methodological advancements (Dong, Shonkwiler, and Capps Jr, 1998; Golan, Perloff, and Shen, 2001; Dong, Gould, and Kaiser, 2004; Wood, Nelson, and Nogueira, 2012) and policy evaluations (Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al., 2015). We focus on the ENIGH 2018 cross-section that includes data from over 70,000 households, representing almost the 120 million Mexicans.

ENIGH aggregates households' purchased commodities across large aggregations such as SSB that include different brands and flavors of soda drinks, but also food items such as food away from home, fruits and vegetables. Thus, an initial empirical consideration is the creation of a basket of goods from which we can express demand as a function of aggregate expenditure. That is, demand for a category of commodities must be expressed as a function of the expenditure on that category as supposed to individual prices. We rely on the concepts of weak separability and two-stage bud-

getting for the construction of our basket of goods² (Deaton and Muellbauer, 1980b): the SSB composite good includes all sugary soft beverages, the juice composite good includes processed fruit beverages and water with fruits, the milk composite good includes dairy beverages, and an other composite good includes a variety of other drinks but is mostly comprised of water.³ To calculate households' beverage budgets, we sum total expenditure across all composite goods. Then, beverage total budget and the composite quantity are used to calculate households' budget shares and unit-values for each composite good. Unit-values, the ratio of expenditure to quantity, is often employed as proxy to composite prices when prices are not observed; the approach, however, carries empirical concerns since unit values reflect prices and household preferences (Cox and Wohlgenant, 1986; Dong, Shonkwiler, and Capps Jr, 1998). We address the quality/price issue with a simple correction proposed in Capacci and Mazzocchi (2011) and further discussed in our empirical strategy section.

We focus our attention to a subset of 63,867 households that reported buying non-alcoholic beverages during the collection of ENIGH. Table 1.1 describes our data across several beverage budget bins that reflect households' incomes among other household characteristics. The budget bins include households in the bottom three quartiles, households between the 75th and the 90th percentiles, and households in the top decile. Expenditure on non-alcoholic beverages ranges from about \$2.46 dollars among the poorest households to \$26.51 among the richest households. The share of the beverage budget allocated to SSB is the largest across all budget bins, suggesting a high level of Mexican households' preference for SSB regardless of income levels. Nevertheless, the share allocated to SSB by households in each budget category first increases as households become richer, but then decreases for the top decile; SSB is a normal good for the bottom nine deciles, but an inferior good for the top decile. Bud-

²The appendix reports item descriptions and their ENIGH codes for all items used in our aggregations. ENIGH codes can be used to trace back our aggregation to the ENIGH documentation to facilitate replication. ENIGH data is publicly available from INEGI in English. ENIGH documentation is available in Spanish from INEGI.

³We exclude non-liquid beverages such as powders and ice, so our composite goods only reflect liquids.

get shares for the rest of our composite goods exhibit similar heterogeneous behavior across households. To investigate this behavior, appendix [A.1](#) fits non-parametric regressions of the log of beverage expenditure (horizontal axes) on each budget share (vertical axes). We find evidence of rank 3 Engel curve behavior among Mexican households in our dataset, which carries econometric and modelling complications discussed in our empirical strategy section (Banks, Blundell, and Lewbel, [1997](#)). Table [1.1](#) also describes composite prices and household demographics across all budget bins. Unit-values have been processed to obtain composite prices. Small standard deviations suggest that prices within the budget bins remain stable, but they vary across budget bins with the exception of milk. We also find that households in our dataset tend to be headed by males in their late 40s, but also richer households tend to be younger and more educated than poorer households. As expected, household size is also a determinant of households' beverage expenditure.

Our data also exhibit a significant number of zero observations where the households report zero expenditure to one or more composite goods i.e., the dependent variable is zero for the household. The SSB composite good has a share of 29% zeros; the juice composite good has a share of 76% zeros; the milk composite good has a share of 34% zeros; and the other composite good has a share of 65% zeros. A significant number of zeros in the dependent variable carries econometric concerns due to the censorship of our dependent variable (Wales and Woodland, [1983](#)). We address the concern in the empirical strategy section.

1.3 Empirical strategy

We divide this section into four sub-sections in which we address each empirical challenge present in our dataset. Each subsection builds upon the previous one, so the last couple of subsections provide with the specifications employed in our estimation procedures. In the first subsection, we discuss the two-step approach proposed in Capacci and Mazzocchi ([2011](#)) to obtain composite prices from unit-values. The quality/price correction requires the average of households within the

Table 1.1: Summary of statistics across different expenditure levels

	Beverage budget bins				
	[Min, 25%]	(25%, 50%]	(50%, 75%]	(75%, 90%]	(90%, Max]
Budget	24.78 (0.454)	64.07 (0.176)	104.59 (0.147)	162.39 (0.109)	265.07 (0.248)
<i>Shares</i>					
SSB	0.38 (0.461)	0.43 (0.438)	0.45 (0.326)	0.464 (0.302)	0.43 (0.284)
Juice	0.07 (0.227)	0.07 (0.163)	0.07 (0.144)	0.08 (0.139)	0.09 (0.43)
Milk	.35 (0.436)	.37 (0.352)	0.36 (0.304)	0.35 (0.277)	0.38 (0.271)
Other	0.19 (0.359)	0.13 (0.239)	0.12 (0.192)	.10 (0.166)	0.10 (0.153)
<i>Prices</i>					
SSB	15.80 (1.181)	16.28 (1.159)	16.78 (1.141)	17.12 (1.126)	17.81 (1.122)
Juice	13.74 (1.336)	14.59 (1.298)	15.03 (1.269)	15.49 (1.242)	15.64 (1.239)
Milk	15.96 (1.161)	16.28 (1.146)	16.61 (1.133)	16.61 (1.130)	16.78 (1.142)
Other	8.81 (2.433)	9.68 (2.330)	10.18 (2.237)	10.91 (2.217)	11.36 (2.179)
<i>Demographic controls</i>					
Sex	1.31 (0.462)	1.29 (0.452)	1.26 (0.441)	1.24 (0.426)	1.21 (0.407)
Age	51.64 (16.490)	49.85 (16.044)	48.24 (15.283)	47.73 (14.665)	48.56 (14.354)
Education	5.16 (2.548)	5.51 (2.501)	5.78 (2.448)	5.99 (2.462)	6.19 (2.540)
Size	3.24 (1.770)	3.45 (2.501)	3.71 (2.448)	4.07 (1.761)	4.53 (1.919)
Observations	16,460	16,071	16,303	9,686	6,453

Note: Monetary values are means within each budget bin measured in 2018 Mexican pesos. Demographic variables are based on the household head, except for household size. Sex is 1 for male and 2 for female. Standard deviations are in parentheses. Budget shares reflect heterogenous preferences for the composite goods: Rank 3 Engel curve behavior. Prices are obtained from unit-values following the work of (Capacci and Mazzocchi, 2011).

municipality to face the same price for the basket of goods, so differences in what households pay are driven by preferences. The next subsection describes the QUAID system of Banks, Blundell, and Lewbel (1997). In contrast to the AID system, the QUAID system allows for goods that can be both inferior and normal depending on the households' income. The third subsection expands the work of Dong, Gould, and Kaiser (2004) to derive likelihood functions used to address the zero observation issue that our dataset exhibits. In contrast to Dong, Gould, and Kaiser (2004)'s general case for M products, our manuscript focuses on a four-product system and provides easy to follow steps to estimate a censored Quadratic AIDS model using standard statistical software like GAUSS, R, or Stata. Our final sub-section evaluates the demand elasticities.

1.3.1 Recovering composite prices from unit-values

Proxying composite prices with unit-values carries significant empirical concerns since unit-values reflect prices and household preferences for quality that are not observed by the researcher. Quality differentials can be manifested in, but are not limited to, the choice of where to buy the item (e.g. supermarket or local shop) or whether the household buys in bulk. Therefore, the unit-value is endogenous and will cause biased model estimates if not appropriately addressed (Cox and Wohlgemant, 1986; Deaton, 1988; Deaton, 1990; Dong, Shonkwiler, and Capps Jr, 1998; Capacci and Mazzocchi, 2011). In this manuscript, we follow Capacci and Mazzocchi (2011)'s approach to correct for unit-value endogeneity from a structural decomposition of unit-values at the municipal and household levels.

Capacci and Mazzocchi (2011) demonstrate that the log of unit-values $\ln v_{ih}^m$, is a function of household demographics z_{ih}^m , a vector of logged prices $\ln p_i^m$, the logged amount of goods included in the composite good $\ln q_{ih}^m$, and a stochastic error u_{ih}^m :

$$\ln v_{ih}^m = \eta_{i0} + \sum_l^L \eta_{lh} z_{lh}^m + \mu_i \ln q_{ih}^m + \psi \ln p_i^m + u_{ih}^m \quad (1.1)$$

where η_{lh} , ψ and μ_i are coefficients associated with variables, η_{i0} is the regression

constant, the subscript l is used for each demographic, the superscript m indicates that a household h belongs to a municipality m , and i refers to a composite good.

Two assumptions are necessary to recover composite prices from unit-values. First, the average of the households within the municipality. This assumption is not a concern because our municipal averages are constructed with the observed households within the municipality. The second assumption requires that all households within a municipality perceive the same composite price. Caro et al. (2017) also report employing the same correction in their study of SSB consumption for Colombian households.

Consider the average values of the variables in equation 1.1, but at the municipality level: $\ln \bar{v}_i^m$, $\bar{z}_l^m$, $\ln \bar{q}_i^m$, $\ln \bar{p}_i^m$ and $\bar{u}_i^m$. The assumption that consumers face the same prices within a municipality allows us to recover the coefficients in equation 1.1 without the need for a vector of prices with a simple first-difference approach:

$$(\ln v_{ih}^m - \ln \bar{v}_i^m) = \sum_l^L \eta_{lh}(z_{lh}^m - \bar{z}_l^m) + \mu_i(\ln q_{ih}^m - \bar{q}_i^m) + (u_{ih}^m - \bar{u}_i^m) \quad (1.2)$$

with the estimated coefficients from equation 1.2, a further re-arrangement allows us to recover a vector of composite prices at the municipality level: Equation 1.3. The vector of prices recovered by equation 1.3 is the price variable employed in our final specification.

$$\ln p_i^m = \eta_{i0} + \ln \bar{p}_i^m = \ln \bar{v}_i^m - \sum_l^L \hat{\eta}_{lh} \bar{z}_l^m - \hat{\mu}_i \bar{q}_i^m \quad (1.3)$$

1.3.2 Quadratic Almost Ideal Demand System

The AID system was brought to the forefront as the first approach to derive elasticities that exhibit the behavioral restrictions of the Cournot and Engel aggregations (Deaton and Muellbauer, 1980a). The AID system, however, is limited to rank 2 Engel curve behavior such that budget elasticities are monotonic across the income spectrum. Later, Banks, Blundell, and Lewbel (1997) expand the AID system into

the QUAID system to account for households that exhibit rank 3 Engel curve behavior where budget elasticities with respect to the goods can be both inferior and normal. The QUAID system can be recovered from the indirect utility function described in equation 1.4:

$$\ln V(\mathbf{p}, w) = \left[\left\{ \frac{\ln w - \ln a(\mathbf{p})}{b(\mathbf{p})} \right\}^{-1} + \lambda(\mathbf{p}) \right]^{-1} \quad (1.4)$$

where $\ln a(\mathbf{p})$ is a price index defined in equation 1.5,

$$\ln a(\mathbf{p}) = \alpha_0 + \sum_{i=1}^M \alpha_i \ln p_i + \frac{1}{2} \sum_{i=1}^M \sum_{j=1}^M \gamma_{ij} \ln p_i \ln p_j \quad (1.5)$$

$b(\mathbf{p})$ is a Cobb-Douglas price aggregator defined in equation 1.6,

$$b(\mathbf{p}) = \prod_{i=1}^M p_i^{\beta_i} \quad (1.6)$$

and $\lambda(\mathbf{p})$ is defined in equation 1.7 as introduced by Banks, Blundell, and Lewbel (1997) to allow the demand system to exhibit Engel curves of rank 3:

$$\lambda(\mathbf{p}) = \sum_{i=1}^M \lambda_i \ln p_i. \quad (1.7)$$

Here, $\ln w$ is the log of expenditure, $\ln p_i$ is the log of price of each composite good $i = \{1, \dots, M\}$, where M is the total number of composite goods in the basket, and α_i , β_i , λ_i and γ_{ij} are estimable parameters. Applying Roy's identity to equation 1.4, an estimable demand system can be obtained:

$$U_i = \alpha_i + \theta'_i \mathbf{z} + \sum_{j=1}^M \gamma_{ij} \ln p_j + \beta_i \ln \left\{ \frac{w}{a(\mathbf{p})} \right\} + \frac{\lambda_i}{b(\mathbf{p})} \left[\ln \left\{ \frac{w}{a(\mathbf{p})} \right\} \right]^2 \quad i = 1, \dots, M. \quad (1.8)$$

To incorporate demand effects of household demographics, equation 1.8 includes a vector of demographic variables $\mathbf{z}$ with their associated vector of parameters θ_i . This

way of accounting for household demographic effects is called demographic translating, which allows demographics to change the level of demand but not to alter the slope. Note that this makes demographic heterogeneity enter the demand system linearly via the intercept in equation 1.8, as well as non-linearly through the households' expenditure via the price aggregator in equation 1.5 (Pollak and Wales, 1981; Lecocq and Robin, 2015). Additionally, notice that if $\lambda_i = 0$, the resulting demand system reduces to the AID system proposed by Deaton and Muellbauer (1980a).

Additional econometric adjustments are required to ensure that the estimated parameters in equation 1.8 are theoretically consistent. Thus, they are constrained as follows:

$$\sum_{i=1}^M \alpha_i = 1 \quad \sum_{i=1}^M \beta_i = 0 \quad \sum_{j=1}^M \gamma_{ij} = 0 \quad \sum_{i=1}^M \lambda_i = 0 \quad \sum_{i=1}^M \theta_i = 0 \quad \gamma_{ij} = \gamma_{ji} \quad \forall i, j$$

The above constraints are for homogeneity, adding-up, and symmetry. They are important to restrict the demand system being from a utility maximization and also to allow demand elasticities to exhibit the behavior depicted by the Cournot and Engel aggregations.

1.3.3 Likelihood function derivation for a censored QUAID system

Wales and Woodland (1983) discuss two procedures to estimate valid demand systems that provide elasticities that reflect the behavior predicted by the Cournot and Engel aggregations in the presence of data censoring. One is the Kuhn-Tucker approach that maximizes an explicitly specified random utility function to derive the demand system. The other one is the Amemiya-Tobin approach that derives the demand system from a non-random utility function then adds a random error term to the derived system for its model estimation. Our work adapts the Amemiya-Tobin procedure to estimate a censored QUAID system, following the work of Dong, Gould, and Kaiser (2004). For the remainder of this section, we focus on the derivation of

the likelihood functions used in the estimation of our four-product system and the derivation of the demand elasticities.

Following Wales and Woodland (1983), we add an error term to equation 1.8 to form the share system of four products as:

$$S^* = U + \epsilon \quad (1.9)$$

where we define U as the stacked system of equations formed by U_i in equation 1.8, and ϵ is the stacked vector for the stochastic terms ϵ_i . Also, ϵ is assumed to have a joint normal distribution with a mean vector of zero and a variance covariance matrix as:

$$\Sigma^* = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} & \sigma_{14} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} & \sigma_{24} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} & \sigma_{34} \\ \sigma_{14} & \sigma_{24} & \sigma_{34} & \omega_{44} \end{pmatrix}. \quad (1.10)$$

The adding-up restriction described in the previous sub-section implies that the joint density function of ϵ is singular (i.e., equation 1.10 is a singular matrix). Therefore, we must drop one of the share equations from the estimation. The dropped parameters, including the dropped element in ϵ can be recovered from the theoretical restrictions discussed in the previous subsections. For example, if we drop the fourth equation, the adding-up condition would allow us to recover the fourth error element with $\epsilon_4 = -(\epsilon_1 + \epsilon_2 + \epsilon_3)$ and the fourth α coefficient in equation 1.8 with $\alpha_4 = 1 - (\alpha_1 + \alpha_2 + \alpha_3)$. Therefore, the estimation procedure we describe here aims to estimate the parameters in a system of three equations depicted in equation 1.8 and the variance-covariance elements of three equations.

S_i^* in equation 1.9 are not bounded between 0 and 1. S_i^* could be greater than 1 or negative if the data are censored. We call them the latent shares and they are not observable. The first step to estimate equation 1.9 in the Amemiya-Tobin approach is to link the latent shares with the actual shares we observe (the observed shares). Following Wales and Woodland (1983), we map the latent shares to the observed

shares as below:

$$S_i = \begin{cases} \frac{S_i^*}{\sum_{j \in \Psi} S_j^*}, & S_i^* > 0, \\ 0, & S_i^* \leq 0 \end{cases} \quad i = 1, \dots, 4 \quad (1.11)$$

where Ψ is the set of all positive latent shares. Wales and Woodland (1983) note that equation 1.11 is not the only way to map latent shares into observed shares; however, equation 1.11 is simple and has the property that the resulting density function is independent of whatever set of S^* is used in its derivation. Notice that the mapping of latent shares S^* to the observed shares S described in equation 1.11 guarantees that all observed shares sum to 1 and that the observed shares lie between 0 and 1. As an example, suppose we have $S_1^* = 1.7$, $S_2^* = -0.3$, $S_3^* = 0.5$ and $S_4^* = -0.9$. The observed shares from the mapping rule in equation 1.11 will be $S_1 = 1.7/(1.7 + 0.5) = 0.77$, $S_2 = 0$, $S_3 = 0.5/(1.7 + 0.5) = 0.23$, and $S_4 = 0$. Thus, we have the latent shares (S^*) and the observed shares (S) both sum to 1 (adding-up).

Existing Amemiya-Tobin type estimation procedures in the literature, including the most used two-stage (e.g., Heien and Wesseils (1990); Shonkwiler and Yen (1999)), for a censored demand system impose adding-up only to the latent shares and ignore the budget constraint in the observed shares, which is obviously wrong and incurs biased model estimates. More recently, Garcíea-Enríquez and Echevarría (2016) proposed a two-step estimation procedure for a censored demand system based on Tauchmann (2010)'s two-step estimators for a censored system of equations. Their method addressed the budget constraint issue for observed shares by using only observations that are non-zero for all the goods in the system in the second step. This could be an issue because a large number of observations can be excluded from the second step.⁴

We start our maximum likelihood estimation procedure by building up the likeli-

⁴In our data, only 5% of the households purchased all of the 4 goods in our system. Therefore, 95% of the households have to be excluded from the second step estimation. Furthermore, multivariate probability integrals still need to be evaluated in their first step, which could involve probability simulation as mentioned in Dong, Gould, and Kaiser (2004).

hood function for the model depicted in equations 1.8 to 1.11 for each of the households. Each household follows a purchase pattern. For example, some households purchase only one product and some others purchased two of the four products. Thus, the total unique purchase patterns, often referred to as purchase regimes, among households equals $2^4 - 1 = 15$ for a four-product system. We build up the likelihood functions for each of the regimes. For the sake of exposition, assume that the purchased products are arranged such as they are the first ones in our notation. Suppose after the rearrangement, the error terms of the first three equations have a variance covariance matrix as:

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma_{22} & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma_{33} \end{pmatrix}. \quad (1.12)$$

After the rearrangement, we are able to define the following four regimes that cover all the possible 15 regimes the households could have and derive the likelihood functions for each of them.

Regime 1: Only one product is purchased

Following Wales and Woodland (1983), the likelihood of this regime is:

$$L_1(S_1 > 0, S_2 = S_3 = S_4 = 0) = \int_{S_1}^{\infty} \int_{(1-\frac{S_1^*}{S_1})}^0 \int_{(1-\frac{S_1^*}{S_1}-S_2^*)}^0 \phi(S_1^*, S_2^*, S_3^*) dS_3^* dS_2^* dS_1^* \quad (1.13)$$

where $\phi(S_1^*, S_2^*, S_3^*)$ is the trivariate normal probability density function (PDF) with a mean vector of (U_1, U_2, U_3) and a variance-covariance matrix given in equation 1.12. The U s are defined in equation 1.8. Equation 1.13 is a plane truncated multinormal distribution (Tallis, 1965). Following Dong, Gould, and Kaiser (2004), we transform equation 1.13 to a rectangular truncated standard normal distribution as:

$$L_1(S_1 > 0, S_2 = S_3 = S_4 = 0) = \Phi_3(b, R_c) \quad (1.14)$$

where

$$b = \begin{pmatrix} 1 - \frac{(\frac{U_1}{S_1} + U_2 + U_3)}{V} \\ -\frac{U_2}{\sqrt{\sigma_{22}}} \\ -\frac{U_3}{\sqrt{\sigma_{33}}} \end{pmatrix},$$

$$R_c = \begin{pmatrix} 1 & \frac{-(\frac{\sigma_{12}}{S_1} + \sigma_{22} + \sigma_{23})}{\sqrt{\sigma_{22}}V} & \frac{-(\frac{\sigma_{13}}{S_1} + \sigma_{23} + \sigma_{33})}{\sqrt{\sigma_{33}}V} \\ \frac{-(\frac{\sigma_{12}}{S_1} + \sigma_{22} + \sigma_{23})}{\sqrt{\sigma_{22}}V} & 1 & \frac{\sigma_{23}}{\sqrt{\sigma_{22}}\sqrt{\sigma_{33}}} \\ \frac{-(\frac{\sigma_{13}}{S_1} + \sigma_{23} + \sigma_{33})}{\sqrt{\sigma_{33}}V} & \frac{\sigma_{23}}{\sqrt{\sigma_{22}}\sqrt{\sigma_{33}}} & 1 \end{pmatrix},$$

and

$$V = \left(\frac{\sigma_{11}}{S_1^2} + \sigma_{22} + \sigma_{33} + \frac{2\sigma_{12}}{S_1} + \frac{2\sigma_{13}}{S_1} + 2\sigma_{23} \right)^{\frac{1}{2}}.$$

$\Phi_3(\cdot)$ is a trivariate standard normal cumulative density function (CDF) evaluated at b with a correlation matrix of R_c . Most statistical software like Stata, R, and GAUSS, have build-in modules to calculate $\Phi_3(\cdot)$. When the dimension of the CDF is greater than 3, as shown in Dong, Gould, and Kaiser (2004), the calculation of the CDF may be a problem. However, as proposed by Dong, Gould, and Kaiser (2004), simulated probability procedure can be used to evaluate the high-dimensional CDF.

Regime 2: Two products are purchased

Similar to regime 1, we transform the likelihood function from a plane truncation to a rectangular truncation:

$$\begin{aligned}
& L_2(S_1 > 0, S_2 > 0, S_3 = S_4 = 0) \\
&= \int_{S_1}^{\infty} \int_{(1-\frac{S_1^*}{S_1})}^0 \phi(S_1^*, S_2^*, S_3^*) dS_3^* dS_1^* \\
&= B\Phi_2(b, R_c)
\end{aligned} \tag{1.15}$$

where

$$B = (2\pi)^{-\frac{1}{2}} |\Sigma|^{-\frac{1}{2}} |\Omega_{11}|^{-\frac{1}{2}} e^{-\frac{1}{2}} \left\{ \begin{pmatrix} U_1 \\ U_3 \end{pmatrix}' \Omega_{00} \begin{pmatrix} U_1 \\ U_3 \end{pmatrix} - \begin{pmatrix} U_1^* \\ U_3^* \end{pmatrix}' \Omega_{11} \begin{pmatrix} U_1^* \\ U_3^* \end{pmatrix} \right\},$$

$$b = EG,$$

and

$$R_c = ECR C' E'.$$

Further,

$$\begin{aligned}
E &= \begin{pmatrix} (C_1 R C_1')^{-\frac{1}{2}} & 0 \\ 0 & (C_3 R C_3')^{-\frac{1}{2}} \end{pmatrix}, \\
G &= \begin{pmatrix} 1 - (\frac{1}{s_1} U_1^* + U_3^*) \\ -U_3^* \end{pmatrix}, \\
C &= \begin{pmatrix} \frac{1}{S_1} & 1 \\ 0 & -1 \end{pmatrix} D^{-\frac{1}{2}} = \begin{pmatrix} C_1 \\ C_3 \end{pmatrix},
\end{aligned}$$

and

$$U^* = \begin{pmatrix} U_1^* \\ U_3^* \end{pmatrix} = \Omega_{11}^{-1} \Omega_{10} \begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = U.$$

Also, $I = \begin{pmatrix} 1 & 1 \end{pmatrix}$, and $J = \begin{pmatrix} 1 & (\frac{U_2}{U_1}/\frac{S_2}{S_1}) \end{pmatrix}$ such that,

$$\Omega_{11}^{-1} = \begin{pmatrix} IM_{11}I' & IM_{10} \\ M'_{10}I' & M_{00} \end{pmatrix},$$

$$\Omega_{00}^{-1} = \begin{pmatrix} JM_{11}J' & JM_{10} \\ M'_{10}J' & M_{00} \end{pmatrix},$$

$$\Omega_{10}^{-1} = \begin{pmatrix} IM_{11}J' & IM_{10} \\ M'_{10}J' & M_{00} \end{pmatrix},$$

and

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{S_2}{S_1} & 0 \\ 0 & 0 & 1 \end{pmatrix} \Sigma^{-1} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{S_2}{S_1} & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} M_{11} & M_{10} \\ M'_{10} & M_{00} \end{pmatrix}.$$

M_{11} is the (2×2) submatrix of M , M_{10} is the (2×1) submatrix of M , and M_{00} is the last element in M . R is the correlation coefficient matrix of Ω_{11} . D , a vector of (2×1) , is the diagonal elements of Ω_{11} .

In equation 1.15, $\Phi_2(\cdot)$ is a bivariate standard normal CDF and can be easily evaluated. Dong, Gould, and Kaiser (2004) combined regimes 1 and 2 for a regime when the purchased products between 1 and $M - 1$ and provided the formula for a general of M product system. When $M > 4$, as mentioned above, simulated probability procedure may be needed to evaluate the likelihood function described in equation 1.15.

Regime 3: Three products are purchased

The likelihood function of this regime is:

$$\begin{aligned}
L_3(S_1 > 0, S_2 > 0, S_3 > 0, S_4 = 0) \\
&= \int_{S_1}^{\infty} \phi(S_1^*, S_2^*, S_3^*) dS_1^* \\
&= B \int_{S_1}^{\infty} \phi(S_1^*; U_1^*, \Omega_{11}) dS_1^*
\end{aligned} \tag{1.16}$$

where

$$B = (2\pi)^{-1} |\Sigma|^{-\frac{1}{2}} |\Omega_{11}|^{-\frac{1}{2}} e^{-\frac{1}{2} \left\{ U_1' \Omega_{00} U_1 - U_1' \Omega_{11} U_1^* \right\}},$$

and

$$U_1^* = \Omega_{11}^{-1} \Omega_{10} U_1.$$

Also, $I = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}$, and $J = \begin{pmatrix} \frac{U_2}{U_1} / \frac{S_2}{S_1} & \frac{U_3}{U_1} / \frac{S_3}{S_1} \end{pmatrix}$ such that,

$$M_{11} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{S_2}{S_1} & 0 \\ 0 & 0 & \frac{S_3}{S_1} \end{pmatrix} \Sigma^{-1} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{S_2}{S_1} & 0 \\ 0 & 0 & \frac{S_3}{S_1} \end{pmatrix},$$

$\Omega_{11}^{-1} = I M_{11} I'$, $\Omega_{00}^{-1} = J M_{11} J'$, and $\Omega_{10}^{-1} = I M_{11} J'$.

This likelihood function requires only the integration of a univariate PDF. We can further simplify equation 1.16 as:

$$L_3(S_1 > 0, S_2 > 0, S_3 > 0, S_4 = 0) = B \Phi_1\left(\frac{-S_1 + U_1^*}{\sqrt{\Omega_{11}}}\right) \tag{1.17}$$

where $\Phi_1(\frac{-S_1 + U_1^*}{\sqrt{\Omega_{11}}})$ is the univariate standard normal CDF evaluated at $(\frac{-S_1 + U_1^*}{\sqrt{\Omega_{11}}})$.

Regime 4: All four products are purchased

This regime is the same as the non-censored system. Its likelihood function is:

$$L_4(S_1 > 0, S_2 > 0, S_3 > 0, S_4 > 0) = \phi(S_1^*, S_2^*, S_3^*; U, \Sigma) \quad (1.18)$$

Optimization procedure

Given its purchase pattern, i.e., how many products are purchased, each household will fall into one of the four regimes. The log likelihood function for the household i can be written as:

$$LL_i = i_1 \ln L_1 + i_2 \ln L_2 + i_3 \ln L_3 + i_4 \ln L_4 \quad (1.19)$$

where $i_1 = 1$, if household i purchased 1 product, and 0 otherwise; $i_2 = 1$, if household i purchased 2 products, and 0 otherwise; $i_3 = 1$, if household i purchased 3 products, and 0 otherwise; $i_4 = 1$, if household i purchased 4 products, and 0 otherwise.

The joint log-likelihood function for n households is:

$$LL = \sum_{i=1}^n LL_i \quad (1.20)$$

All model parameters can be obtained from maximizing equation 1.20. We use GAUSS software to estimate the model.

1.3.4 Evaluation of predicted shares and demand elasticities

Elasticity formulas derived from a non-censored demand system, for example, those formulas in Banks, Blundell, and Lewbel (1997), cannot be used for a censored

demand system. Those elasticities are for the latent demand system defined in equation 1.8. What we need are the elasticities for the expected observed demand. Approaches that derive elasticities directly from the expected observed shares defined by equation 1.11 are very complicated as showed in Dong, Gould, and Kaiser (2004). Therefore, we use simulation procedures adapted by Dong, Gould, and Kaiser (2004) from the work of Phaneuf, Kling, and Herriges (2000) to evaluate elasticities in this study.

Based on equation 1.9, we randomly generate R set of replicates for the error term ϵ_i according to its joint normal distribution. The r^{th} simulated latent share of product i , S_{ir}^* , evaluated at the sample means of the variables (indicated by a bar over a variable) is:

$$S_{ir}^* = \alpha_i + \theta'_i \bar{\mathbf{z}} + \sum_{j=1}^M \gamma_{ij} \ln \bar{p}_j + \beta_i \ln \left\{ \frac{\bar{w}}{a(\bar{\mathbf{p}})} \right\} + \frac{\lambda_i}{b(\bar{\mathbf{p}})} \left[\ln \left\{ \frac{\bar{w}}{a(\bar{\mathbf{p}})} \right\} \right]^2 + \epsilon_{ir} \quad (1.21)$$

where ϵ_{ir} is the r^{th} replicate of ϵ_i . The r^{th} replicate of the observed share of product i is then:

$$S_{ir} = \begin{cases} \frac{S_{ir}^*}{\sum_{j \in \Psi} S_{jr}^*}, & S_{ir}^* > 0, \\ 0, & S_{ir}^* \leq 0 \end{cases} \quad i = 1, \dots, 4 \quad (1.22)$$

The expected observed share of product i for R replicates is then calculated as the simple average of these simulated values:

$$E(S_i) = \frac{1}{R} \sum_{r=1}^R S_{ir} \quad (1.23)$$

Price, expenditure, and demographic elasticities can be obtained numerically based on equation 1.23. For example, suppose we have a small change in price j , ΔP_j . The Marshallian elasticity with respect to this price change is:

$$\eta_{jk} = \frac{\Delta E(S_j)}{\Delta P_k} \cdot \frac{\bar{P}_k + \frac{\Delta P_k}{2}}{E(S) + \frac{\Delta E(S)}{2}} - \mu_j \quad (1.24)$$

where $\Delta E(S_j)$ is the change in the simulated expected share after the price change, and $\mu_j = 1$ whenever $j = k$ and 0 otherwise. Notice that the function for η_{jk} resembles an elasticity of the form $\eta_{jk} = \frac{\partial \ln q_j}{\partial \ln p_k} = \frac{\partial q_j}{\partial p_k} \frac{p_k}{q_j}$, where the second element on the right hand side corresponds to the mid-point between the disturbance and the initial values.

Similarly, we can define the budget elasticities as follows:

$$\eta_j = \frac{\Delta E(S_j)}{\Delta w_k} \cdot \frac{\bar{w}_k + \frac{\Delta w_k}{2}}{E(S_j) + \frac{\Delta E(S_j)}{2}} + 1 \quad (1.25)$$

where $\Delta E(S_j)$ is the change in the simulated expected share after the expenditure change, and $\bar{w}_k$ and Δw_k are the expenditure and the expenditure change, respectively.⁵

The Hicksian elasticities can be derived from the Slutsky equation using the Marshallian elasticity given by equation 1.24, the budget elasticity in equation 1.25, and the expected observed shares given in equation 1.23. The formula is $\eta_{ij}^h = \eta_{ij}^m + \eta_i E(S_j)$, where η_{ij}^h and η_{ij}^m denote the Hicksian and the Marshallian price elasticities respectively.

1.4 Discussion of demand elasticities

The results from our estimation procedure to approximate the parameters in equations 1.8 and 1.12 are shown in table 1.2 with their respective standard errors.⁶ As we discussed in our empirical strategy section, we dropped our other category because

⁵The standard errors of the elasticities are obtained using the delta method after we simulate variations in elasticities from changes in the parameters.

⁶Parameter estimates in table 1.2 (and the demand elasticities in tables 1.3, 1.4 and 1.5 below) are statistically significant at the 1% level of confidence. We attribute the small magnitude of our standard errors to our large sample size of 63,867.

of the singularity in our system of equations. The dropped parameters can be easily recovered following the behavioral restrictions discussed in our empirical strategy section. The estimated parameters in table 1.2 lack a direct interpretation, so we focus our attention on the demand elasticities.

In equation 1.3, we describe our Marshallian and budget share elasticities.⁷ Table 1.3 provides evidence that SSB is an inelastic good with an own-price elasticity of -0.95. As noted in table 1.1, Mexican households have a strong preference for SSB with budget shares ranging from 38% to 46% across the income spectrum. Low own-price elasticities with respect to SSB are consistent with the recent findings by Dubois, Griffith, and O’Connell (2020) that show that SSB taxes are the least effective among consumers with high sugar intake. In addition, our findings for the own-price elasticity of SSB are in direct contrast to the findings in previous studies that report that SSB is an elastic good. Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) estimate the AID system for a basket of drinks that include SSB, finding that the one peso tax per liter contributed to a decline in demand for SSB of 22.6% with an own-price elasticity at -1.13.⁸ Their study and our manuscript employ the same survey (ENIGH) but for different years. In a similar study for Colombian households, Caro et al. (2017) employed the QUAID system to estimate price elasticities for SSB and a basket of other goods including sweet snacks. The authors find that the own-price elasticity of SSB is -1.61. Guerrero-López, Unar-Munguía, and Colchero (2017) find similar results after estimating both the AID

⁷These elasticities are conditional on the total expenditure of the basket of goods allocated from the first stage of a two-stage budgeting process (Gorman, 1959; Okrent and Alston, 2012). Empirical analyses of food demand typically focus on only one stage of a household’s budgeting process. In this study, we follow this general approach and focus on the second stage only. However, in this approach, the total expenditure could be endogenous and bias model estimates if not appropriately accounted for (LaFrance, 1993). To address this issue, we run an auxiliary regression of the total beverage expenditure against household income and demographics. The predicted total beverage expenditure from the auxiliary regression is then used in our share equations (Dhar, Chavas, and Gould, 2003).

⁸We test the bias across several specifications found in the literature in appendix A.2. We find that the upward bias found in Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) is associated with the unit-value endogeneity discussed in section 3.1; the own-price elasticity with respect to SSB is -1.320 before correcting for zero-purchases and censorship. Our censorship correction corrects a downward bias after correcting for the unit-value endogeneity.

Table 1.2: Description of Censored-QUAIDS estimates

Parameter	Estimate	Parameter	Estimate	Parameter	Estimate
α_1	0.543 (0.030)	θ_{22}	0.044 (0.006)	γ_{21}	-0.021 (0.006)
α_2	0.429 (0.024)	θ_{32}	-0.004 (0.001)	γ_{22}	0.004 (0.006)
α_3	0.502 (0.029)	θ_{42}	-0.012 (0.003)	γ_{31}	0.006 (0.008)
λ_1	-0.027 (0.009)	θ_{13}	-0.010 (0.004)	γ_{32}	0.022 (0.005)
λ_2	0.014 (0.007)	θ_{23}	-0.084 (0.007)	γ_{33}	-0.022 (0.008)
λ_3	0.018 (0.009)	θ_{33}	0.003 (0.001)	σ_{11}	0.239 (0.001)
θ_{11}	-0.005 (0.005)	θ_{43}	-0.002 (0.003)	σ_{21}	-0.044 (0.001)
θ_{21}	0.001 (0.008)	β_1	0.0621 (0.029)	σ_{22}	0.214 (0.001)
θ_{31}	-0.011 (0.001)	β_2	-0.081 (0.003)	σ_{31}	-0.128 (0.002)
θ_{41}	-0.012 (0.003)	β_3	-0.086 (0.026)	σ_{32}	-0.041 (0.002)
θ_{12}	-0.014 (0.004)	γ_{11}	0.028 (0.010)	σ_{33}	0.246 (0.002)

Note: Table shows the coefficient estimates and their respective standard errors for the system of equations illustrated in 1.8 and the elements in 1.12. Parameters are restricted to follow the behavioral restrictions discussed in our empirical strategy such that elasticities can follow the Engel and Cournot aggregations. The estimates σ_{ij} are the elements in 1.12 which are later used to estimate the standard errors for our Marshallian and budget elasticities using the delta method. We assume that $I_4\epsilon = 0$ is a multivariate normal, where I_4 is a 4 vector of ones, the adding-up and symmetry restrictions imply that the joint density function of ϵ is singular. We drop our other category. The dropped parameters, including the dropped element in ϵ , can be recovered from the theoretical restrictions.

Table 1.3: Marshallian and budget elasticities

With respect to	Price				Budget
	SSB	Juice	Milk	Other	
Quantity					
SSB	-0.951 (0.0001)	-0.055 (0.0000)	-0.005 (0.0000)	0.002 (0.0000)	1.009 (0.0000)
Juice	-0.039 (0.0001)	-0.948 (0.0001)	0.112 (0.0000)	-0.049 (0.0000)	0.923 (0.0000)
Milk	0.022 (0.0001)	0.075 (0.0000)	-1.031 (0.0001)	-0.029 (0.0000)	0.963 (0.0000)
Other	-0.334 (0.0001)	-0.228 (0.0001)	-0.276 (0.0001)	-0.639 (0.0000)	1.478 (0.0001)

Note: Table shows Marshallian and budget elasticities for our basket of four goods. Marshallian elasticities are calculated with equation 1.24, and budget elasticities are calculated with equation 1.25. Demand elasticities are evaluated at their mean, using the estimated parameters in table 1.2. Our expected shares (equation 1.23) are $E(S_{SSB}) = 0.4117$, $E(S_{Juice}) = 0.2338$, $E(S_{Milk}) = 0.3014$, and $E(S_{Other}) = 0.0531$, and they can be used to test the Engel and Cournot aggregations. The homogeneity in our system of equations can also be tested (i.e., the rows add-up to zero).

and QUAID systems for Chilean consumers with an own-price elasticity of -1.37.

It is important to notice several distinctions between our work and previous studies that make positive conclusions about the effectiveness of SSB price hikes with their estimated own-price elasticity of SSB in Latin-American countries. We attribute the difference between our own-price elasticity with respect to SSB and the findings in the previously cited studies to several unaddressed empirical shortcomings in the previous studies. The work of Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) for the Mexican case employs an AID system even though households in their dataset exhibit non-linear Engel curves. The failure to account for non-linear Engel curves implies that estimated-elasticities are unable to represent income heterogeneity common in health studies i.e., poor households exhibit normal income elasticities with respect to SSB while rich households exhibit inferior income elasticities with respect to SSB. In their study, Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) used the same ENIGH survey. We find evidence of quadratic Engel curves for the years of 2012, 2014 and 2016 in the ENIGH survey

for the four goods considered in our study. Because of methodological differences in the collection of ENIGH, we do not consider previous years.⁹

The work of Caro et al. (2017) recognizes that households in their datasets exhibit quadratic Engel curves, so they employ the QUAID system instead. However, only the work of Caro et al. (2017) addresses the additional concern associated with the use of unit-values i.e., the ratio of expenditure to quantity, to proxy for prices. As mentioned above, the employment of unit-values to proxy for prices carries endogeneity concerns. Finally, none of the studies appropriately addresses the zero purchase observations i.e., households choose not to consume any of the goods, present in their datasets. Caro et al. (2017) employ a common two-stage correction introduced by Shonkwiler and Yen (1999) while Guerrero-López, Unar-Munguía, and Colchero (2017) employ a similar approach first introduced by Heien and Wesseles (1990). As we note in our empirical strategy section, these two-stage corrections impose the adding-up restriction only to the latent share, ignoring the budget constraint in the observed shares.

In table 1.3, we can also study complementarity and substitutability with respect to SSB price hikes. Our results indicate that increases in the price of SSB will increase the consumption of milk, but such price hike will decrease the consumption of juice and our other category. The remainder of our own-price elasticities indicate that juice is inelastic (-.95), milk is elastic (-1.03) and our other category is inelastic (-.64). Our results in table 1.3 also indicate that SSB is a normal good when the budget elasticity is evaluated at the average of our data with a value of 1.01. The remainder of our budget elasticities indicate that juice is also an inferior good (0.92), milk is an inferior good (0.96) and other is a normal good (1.48).

Marshallian elasticities give the gross price effects without compensating the income effect. To show the net price effect, we calculate the Hicksian elasticities in table 1.4. As expected, after compensating the income effect, Hicksian own-price

⁹Additionally, Colchero, Salgado, Unar-Munguía, Hernandez-Avila, et al. (2015) employ a pooled regression that includes the years of 2006, 2008 and 2010 since ENIGH does not follow the same consumer across the years, effectively creating a larger cross-section. This aggregation, however, can carry empirical complications for the structure of the error term if behavior changes across the waves. They employ demographic variables as controls.

Table 1.4: Hicksian elasticities

With respect to	Price			
	SSB	Juice	Milk	Other
Quantity				
SSB	-0.535 (0.0001)	0.181 (0.0000)	0.299 (0.0000)	0.056 (0.0000)
Juice	0.341 (0.0001)	-0.732 (0.0001)	0.391 (0.0001)	0.001 (0.0001)
Milk	0.418 (0.0001)	0.301 (0.0000)	-0.741 (0.0001)	0.022 (0.0000)
Other	0.275 (0.0001)	0.117 (0.0001)	0.169 (0.0001)	-0.561 (0.0000)

Note: Table shows Hicksian obtained from the Slutsky equation using the Marshallian elasticities, the budget elasticities and the expected shares described in table 1.3. The formula is $\eta_{ij}^h = \eta_{ij}^m + \eta_i E(S_j)$, where η_{ij}^h and η_{ij}^m denote the Hicksian and the Marshallian elasticities respectively.

elasticities become smaller in magnitude than their Marshallian counterparts. The Hicksian own-price elasticity with respect to SSB is -0.54, which is comparable to the one of -0.47 reported by Seale Jr, Regmi, and Bernstein (2003) in their multi-country demand elasticity analysis that includes Mexico. The difference between our Marshallian own-price elasticity and our Hicksian own-price elasticity with respect to SSB is 44%, indicating a large income compensating effect.

In table 1.5, we describe our demographic elasticities that shift our budget shares. We find that education is a major factor to decrease the consumption of SSB, and increase the consumption of our other good. While small, having a female head of household decreases the consumption of SSB. The age of the household and household size slightly increases the consumption of SSB.¹⁰

¹⁰In appendix A.3, we replicate table 1.3 for the years of 2014 and 2016 to both study consumer behavior after the sin tax was enacted in 2014, but also as robustness test for our conclusions. We observe that elasticities do not change considerably from year to year, but for the own-price elasticity, we observe that SSB is becoming more elastic i.e., it approaches -1. In contrast to the findings in Colchero, Salgado, Unar-Mungui a, Hernandez-Avila, et al. (2015), we find a more inelastic elasticity than what they found in their study.

Table 1.5: Household demographic elasticities

With respect to	Household demographic variable			
	Age	Size	Education	Female
Quantity				
SSB	0.004 (0.0000)	-0.007 (0.0000)	-0.081 (0.0000)	-0.003 (0.0000)
Juice	-0.027 (0.0001)	-0.062 (0.0000)	-0.039 (0.0000)	-0.006 (0.0000)
Milk	-0.009 (0.0000)	0.069 (0.0000)	0.085 (0.0000)	0.003 (0.0000)
Other	0.136 (0.0000)	-0.067 (0.0000)	0.325 (0.0000)	0.034 (0.0000)

Note: Table shows household demographic elasticities, which are interpreted as shift not rotation of the demand for each of our composite goods. Demand elasticities are evaluated at their mean, using the estimated parameters in table 1.2. Age is the log value of household head's age. Size is the inverse of the number of household members. Education is household head's education. Female indicates if household head is female.

1.5 Conclusion

We revisit the effectiveness of taxing SSB in Mexico to reduce consumption and ultimately curb the country's obesity epidemic. In the lack of an ideal experimental setting, the policy's evaluation requires of structural approaches that aim to model consumer behavior. Our manuscript opts for a censored QUAID system approach, which addresses several empirical shortcomings in previous SSB demand estimations in Latin America. Thus, our manuscript's main contribution is the estimation of the QUAID system for SSB that addresses the unit-value endogeneity and data censoring problem. Our main result suggests that SSB is inelastic with an own-price elasticity of -0.95 (a 10% price hike is associated with a 9.5% decline in consumption) and that SSB is a normal good when the elasticity is evaluated at the average of our data with a 1.01 budget elasticity. Our findings for the own-price elasticity of SSB are in direct contrast to the findings in previous studies that report that SSB is an inelastic good. It is important to notice several distinctions between our work and previous studies that make positive conclusions about the effectiveness of SSB price hikes with their

estimated own-price elasticity of SSB in Latin-American countries. We attribute the difference between our own-price elasticity with respect to SSB and the findings in the previous studies to the several unaddressed empirical shortcomings discussed in our manuscript. We conclude that policy makers in Mexico should re-examine the sin-tax policy to determine its effectiveness to curb the obesity epidemic.

1.6 Chapter 1 References

- Banks, James, Richard Blundell, and Arthur Lewbel (1997). “Quadratic Engel curves and consumer demand”. In: *Review of Economics and Statistics* 79(4), pp. 527–539.
- Barquera, Simon, Lucia Hernandez-Barrera, Maria Lizbeth Tolentino, Juan Espinosa, Shu Wen Ng, Juan A Rivera, and Barry M Popkin (2008). “Energy intake from beverages is increasing among Mexican adolescents and adults”. In: *The Journal of Nutrition* 138(12), pp. 2454–2461.
- Campos-Vazquez, Raymundo M and Eduardo M Medina-Cortina (2019). “Pass-through and competition: The impact of soft drink taxes as seen through Mexican supermarkets”. In: *Latin American Economic Review* 28(1), pp. 1–23.
- Capacci, Sara and Mario Mazzocchi (2011). “Five-a-day, a price to pay: An evaluation of the UK program impact accounting for market forces”. In: *Journal of Health Economics* 30(1), pp. 87–98.
- Caro, Juan Carlos, Shu Wen Ng, Ricardo Bonilla, Jorge Tovar, and Barry M Popkin (2017). “Sugary drinks taxation, projected consumption and fiscal revenues in Colombia: Evidence from a QUAIDS model”. In: *PLOS one* 12(12), e0189026.
- Colchero, M Arantxa, JC Salgado, M Unar-Munguía, M Hernandez-Avila, and JA Rivera-Dommarco (2015). “Price elasticity of the demand for sugar sweetened beverages and soft drinks in Mexico”. In: *Economics & Human Biology* 19, pp. 129–137.

- Colchero, M Arantxa, Juan Carlos Salgado, Mishel Unar-Munguía, Mariana Molina, Shuwen Ng, and Juan Angel Rivera-Dommarco (2015). “Changes in prices after an excise tax to sweetened sugar beverages was implemented in Mexico: Evidence from urban areas”. In: *PLOS one* 10(12), e0144408.
- Colchero, MA, Carlos Manuel Guerrero-López, Mariana Molina, and Juan Angel Rivera (2016). “Beverages sales in Mexico before and after implementation of a sugar sweetened beverage tax”. In: *PLOS one* 11(9), e0163463.
- Cox, Thomas L and Michael K Wohlgenant (1986). “Prices and quality effects in cross-sectional demand analysis”. In: *American Journal of Agricultural Economics* 68(4), pp. 908–919.
- Deaton, Angus (1988). “Quality, quantity, and spatial variation of price”. In: *American Economic Review* 78(3), pp. 418–430.
- Deaton, Angus (1990). “Price elasticities from survey data: Extensions and Indonesian results”. In: *Journal of Econometrics* 44(3), pp. 281–309.
- Deaton, Angus and John Muellbauer (1980a). “An almost ideal demand system”. In: *American Economic Review* 70(3), pp. 312–326.
- Deaton, Angus and John Muellbauer (1980b). *Economics and consumer behavior*. Cambridge University Press.
- Dhar, Tirtha, Jean-Paul Chavas, and Brian W Gould (2003). “An empirical assessment of endogeneity issues in demand analysis for differentiated products”. In: *American Journal of Agricultural Economics* 85(3), pp. 605–617.
- Donaldson, Elisabeth (2015). *Advocating for sugar-sweetened beverage taxation: A case study of Mexico*. Johns Hopkins Bloomberg School of Public Health.
- Dong, Diansheng, Brian W Gould, and Harry M Kaiser (2004). “Food demand in Mexico: An application of the Amemiya-Tobin approach to the estimation of a censored food system”. In: *American Journal of Agricultural Economics* 86(4), pp. 1094–1107.

- Dong, Diansheng, JS Shonkwiler, and Oral Capps Jr (1998). “Estimation of demand functions using cross-sectional household data: The problem revisited”. In: *American Journal of Agricultural Economics* 80(3), pp. 466–473.
- Dubois, Pierre, Rachel Griffith, and Martin O’Connell (2020). “How well targeted are soda taxes?” In: *American Economic Review* 110(11), pp. 3661–3704.
- García-Enríquez, Javier and Cruz A Echevarría (2016). “Consistent estimation of a censored demand system and welfare analysis: The 2012 VAT reform in Spain”. In: *Journal of Agricultural Economics* 67(2), pp. 324–347.
- Golan, Amos, Jeffrey M Perloff, and Edward Z Shen (2001). “Estimating a demand system with nonnegativity constraints: Mexican meat demand”. In: *Review of Economics and Statistics* 83(3), pp. 541–550.
- Gorman, William M (1959). “Separable utility and aggregation”. In: *Econometrica: Journal of the Econometric Society*, pp. 469–481.
- Guerrero-López, Carlos M, Mishel Unar-Munguía, and M Arantxa Colchero (2017). “Price elasticity of the demand for soft drinks, other sugar-sweetened beverages and energy dense food in Chile”. In: *BMC Public Health* 17(1), pp. 1–8.
- Haines, Anjana (2017). “Sugar tax: The winners and losers”. In: *International Tax Review*.
- Heien, Dale and Cathy Roheim Wesseils (1990). “Demand systems estimation with microdata: A censored regression approach”. In: *Journal of Business & Economic Statistics* 8(3), pp. 365–371.
- INEGI (2021). *Indice Nacional de Precios al Consumidor (INPC)*.
- James, Erin, Martián Lajous, and Michael R Reich (2020). “The politics of taxes for health: An analysis of the passage of the sugar-sweetened beverage tax in Mexico”. In: *Health Systems & Reform* 6(1), e1669122.
- LaFrance, Jeffrey T (1993). “Weak separability in applied welfare analysis”. In: *American Journal of Agricultural Economics* 75(3), pp. 770–775.

- Lecocq, Sébastien and Jean-Marc Robin (2015). “Estimating almost-ideal demand systems with endogenous regressors”. In: *The Stata Journal* 15(2), pp. 554–573.
- OECD (2021). *Non-medical determinants of health: Body weight*. Data retrieved from OECD.stat, <https://stats.oecd.org/Index.aspx?ThemeTreeId=9>.
- Okrent, Abigail and Julian Alston (2012). “The demand for disaggregated food-away-from-home and food-at-home products in the United States”. In: *USDA-ERS Economic Research Report*(139).
- Phaneuf, Daniel J, Catherine L Kling, and Joseph A Herriges (2000). “Estimation and welfare calculations in a generalized corner solution model with an application to recreation demand”. In: *Review of Economics and Statistics* 82(1), pp. 83–92.
- Pollak, Robert A and Terence J Wales (1981). “Demographic variables in demand analysis”. In: *Econometrica: Journal of the Econometric Society*, pp. 1533–1551.
- Seale Jr, James L, Anita Regmi, and Jason Bernstein (2003). *International evidence on food consumption patterns*. Tech. rep.
- Shonkwiler, J Scott and Steven T Yen (1999). “Two-step estimation of a censored system of equations”. In: *American Journal of Agricultural Economics* 81(4), pp. 972–982.
- Tallis, GM (1965). “Plane truncation in normal populations”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 27(2), pp. 301–307.
- Tauchmann, Harald (2010). “Consistency of Heckman-type two-step estimators for the multivariate sample-selection model”. In: *Applied Economics* 42(30), pp. 3895–3902.
- Urzua, Carlos M (2008). “Evaluación de los efectos distributivos y espaciales de las empresas con poder de mercado en México”. In: *Documento de trabajo, Tecnológico de Monterrey, Campus Ciudad de México*.

- Wales, Terence J and Alan Donald Woodland (1983). “Estimation of consumer demand systems with binding non-negativity constraints”. In: *Journal of Econometrics* 21(3), pp. 263–285.
- Wood, Benjamin DK, Carl H Nelson, and Lia Nogueira (2012). “Poverty effects of food price escalation: The importance of substitution effects in Mexican households”. In: *Food Policy* 37(1), pp. 77–85.

CHAPTER 2

A MODEL OF THE U.S. FOOD SYSTEM: WHAT ARE THE DETERMINANTS OF THE STATE VULNERABILITIES TO PRODUCTION SHOCKS AND SUPPLY CHAIN DISRUPTIONS?

This paper expands a Ricardian trade general equilibrium model to assess the U.S. state vulnerabilities to production shocks and supply chain disruptions under the existing food system. We find that the distribution of welfare losses caused by disruption in the existing supply chain depends on the state’s agricultural capacity and its operating expenses. Conversely, the existing supply chain capability to mitigate the impacts from productivity shocks is hindered by reliance on in-state crop production. Our extension of the Ricardian trade formulation is a comprehensive but simple model of the U.S. market of crops. We incorporate the heterogeneity of farmers’ production and isolate food processing plants as the main consumer of domestic crop production. Our general equilibrium model allows us to derive conclusions that are state specific, so they can be employed to craft state level policies to strength the national food system.

2.1 Introduction

Biedny, Malone, and Lusk (2020) conclude that despite the sharp political polarization between the two main U.S. parties, both sides agree on a specific food policy issue: shifting food production back to local markets. Proponents of policies that promote local agriculture generally raise environment and health concerns¹ (DuPuis

¹In Blay-Palmer et al. (2013), several examples of bottom-up approaches to the promotion of local agriculture are discussed.

and Goodman, 2005; Levkoe, 2011; Marsden, 2013). Their claim is that as distance between food sources and consumers increases, a disconnection between the adverse consequences of food production widens, hindering regulators’ efforts to internalize negative externalities (Clapp, 2015). Even recently in the face of COVID-19, reliance on distant food sources caused waste of resources that increased food prices (Hobbs, 2020; Thilmany et al., 2020). Long before these most recent debates on inward- versus outward-looking approaches to production, however, economists have emphasized the gains to be had from specialization to expand access to efficiently produced products (Krugman, 1981; Eaton and Kortum, 2002; Arkolakis, Costinot, and Rodríguez-Clare, 2012). Due to the heterogeneity of U.S. agriculture, the strengths and vulnerabilities of the food system depends on a combination of both.

Recent work for the U.S. considers the impact of severe agricultural productivity shocks on farming communities; namely, farmworkers impacted by the California drought of 2012-2017 (Rodríguez et al., 2015; Tortajada et al., 2017; Greene, 2018). But trade-induced resilience has been demonstrated to mitigate the impact from food shortages. Burgess and Donaldson (2010) employ rail data from the colonial era in India to show that the occurrence of famines is reduced when farming communities rely on external sources for their food supply. The extreme opposite to reliance on domestic food sources is import dependency, and a growing stream of the food security literature studies the severe welfare repercussions from relying on a limited number of food sources. For example, whenever the exporter market faces a production shock such as a drought, Puma et al. (2015) find significant reductions on the level of imports that are not necessarily substituted for domestic production. Moreover, Gephart et al. (2016) employ a forward shock-propagation model, and their results suggest that the increase of import reliance can increase food vulnerability to production shocks.

The latter two analyses of food vulnerability and trade focus on consumer behavior and food production separately. Gephart et al. (2016) account for consumers’ import substitution between goods, allowing for consumers to adjust their expenditures according to their preferences, but fail to account for producers’ responses. Since the model does not permit producers to gain market share over those im-

pacted by the weather shock, welfare simulation results in regions not affected by the weather shock are likely understated. On the other hand, Puma et al. (2015) study countries' welfare after simulating disruptions in European and Asian exports, but their simulations do not account for consumers' substitution responses, causing welfare simulation results to likely be overstated. An appropriate assessment of food vulnerabilities must consider both the producers' and consumers' responses.

The simultaneous consideration of consumers' utility and producers' profit optimization problems connects trade between regions with supply and demand determinants through a gravity-like equation (Anderson, 1979). Approaches based on gravity models are ideal for assessing the normative implications of various counterfactuals on different policies related to production and trade. The form of the gravity equation, however, depends on the research interest. Anderson and Wincoop (2003) popularize the Armington-CES model that emphasizes demand determinants such as regions' income. In contrast, Eaton and Kortum (2002) propose a Ricardian analysis that emphasizes producers' relative comparative advantage as reflected by differences in output prices and input costs.² A fundamental difference between the Armington-CES model and the Ricardian formulation of Eaton and Kortum (2002) is that the latter structurally captures the basis of comparative advantage. The approach starts by purging the model from dyadic variation that embodies bilateral determinants of trade such as how far apart two trade partners are from each other. Then, monadic variation can be further split into a productivity term and a production costs term.

Few empirical studies of agriculture base their analyses on a Ricardian notion of comparative advantage. Closer to our work, Reimer and Li (2009); Reimer and Li (2010) exploit this formulation to study how yield variability will affect households for a selected group of countries, finding that openness to trade is positively related to countries' welfare simulation results. In a similar manner, we build on the Ricardian trade model of Eaton and Kortum (2002) to study U.S. states food vulnerability in the existing domestic food system.

²Both the Armington-CES model and the Ricardian trade model are equivalent under the assumption that the trade elasticity θ is equal to $1 - \sigma$, where σ is the elasticity of substitution between goods (Yotov et al., 2016).

Our general equilibrium application allows us to conduct counterfactual simulations to assess the implications of production shocks and disruptions to the food supply chains on consumer welfare. Production shocks are modeled through reductions in each state’s agricultural capacity—a structural parameter that governs each state’s agricultural productivity. Food supply chain disruptions are modeled through doubling the cost of each state’s cost to import. In addition, we simulate the extreme scenario in which trade is shut down such that each U.S. state is forced to rely on local agriculture. While unrealistic, our autarky counterfactual provides evidence about the mechanisms through which local agriculture benefits consumers in each state. To assess the distribution of impacts, we construct a ranking of relative comparative advantage based on trade flows to further study the determinants of food vulnerability in each of our simulations. We find that the distribution of welfare losses depends on a state’s ability to substitute consumption sources.

This article makes three primary contributions to the literature on food security and trade. First, our approach is based on a comprehensive but simple model of the U.S. market for crops that can be employed to analyze a large array of counterfactuals concerning food security in the United States. Our approach accounts for the heterogeneity of farmers’ technology through producer-specific efficiency terms. In addition, we recognize food processing plants as the main consumers of farmers’ production and isolate their substitution effect. By simulating consumer and producer interaction through crop prices, we are able to account for substitution behavior and avoid biases that arise in partial equilibrium analyses because the latter are unable to capture simultaneous responses between producers and consumers.

Secondly, we construct a ranking of relative comparative advantage using our trade data. While most Ricardian trade studies recognize competitiveness (of exporter) and openness (of importer) in their applications, researchers do not exploit the ranking and rather focus on the decomposition of the determinants of productivity and comparative advantage. For example, the canonical formulation of Eaton and Kortum (2002) uses the openness terms to infer countries’ willingness to trade, and more recently, Keller, Santiago, and Shiue (2017) uses the competitiveness terms to study technological spillovers in China. We recognize the quality of the ranking as relative

comparative advantage in the system and use it to study the mechanisms that shape welfare losses in our food vulnerability counterfactuals.

Due to the challenges associated with collecting domestic trade data, our final contribution is concerned with the expanding literature of trade analyses in domestic settings. Since U.S. domestic trade lacks of non-market barriers such as tariffs, focusing on domestic trade presents the opportunity to study market mechanisms that govern trade and welfare distributions. We decompose each U.S. state’s comparative advantage into an agricultural capacity term and an operating expenses term. To our knowledge, our study is the first to adopt a general equilibrium Ricardian approach to modeling productivity and domestic trade in agriculture at a regional (sub-national) level.

The remainder of the paper is organized as follows. The next section reviews our extension of the formulation of Eaton and Kortum (2002), and its implications for the U.S. domestic market for crops. Section 3 presents our empirical strategy. First, we describe how we process the data to capture the economic realities of the U.S. domestic crop market. Then, we estimate the model parameters necessary for our post-estimation simulations. The final subsection conducts the decomposition of each state’s comparative advantage. In section 4, we discuss our simulation strategy and results from our counterfactuals. Finally, we offer concluding remarks.

2.2 A Ricardian trade model of U.S. domestic crop production

Our conceptual framework considers the two sides of the market. On the one hand, farmers are crop producers. We denote farmers as i and refer to them as exporters, producers and the origin of the commodity interchangeably. U.S. farmers operate

under different weather, technology and land characteristics across U.S. states.³ We capture their heterogeneity with a land productivity term denoted as $a_i(l)$, where l indicates commodity, a technological productivity-augmenting factor denoted as z_i , and farmers' production costs denoted as c_i (Eaton and Kortum, 2002; Heerman, 2020). We denote final prices as $P_{ij}(l)$ and bilateral trade costs as t_{ij} , which distinguish the source of the commodity, i , and where it is consumed, j .⁴ Assuming constant returns to scale technologies, final price is related to costs of production, a farmer's efficiency and bilateral trade costs as illustrated in 2.1:

$$P_{ij}(l) = \frac{c_i}{z_i a_i(l)} t_{ij}. \quad (2.1)$$

Our producer-specific technological term, z_i , is assumed to be drawn from a Fréchet probability distribution function: $F_i(z) = e^{-T_i z^{-\theta}}$. Here $T_i > 0$ is a producer-specific parameter that governs a region's productivity outcome, so as T_i grows larger, the probability of drawing a large efficiency outcome, z_i , increases. We refer to T_i as agricultural capacity. The parameter $\theta > 1$ is common across producers and governs the dispersion of the productivity draws of individual producers. As θ tends to infinity, the probability of a single region dominating the market decreases (i.e., large θ values imply that productivity differences between regions will be minimal). The employment of the Fréchet distribution is a realistic approach since its quality of focusing on extreme values captures both individual producers' productivity and the overall distribution of comparative advantage across all producers.

In turn, food processing plants are the crop consumers. We denote the food pro-

³In the U.S., a vast tradition of economics studies evidence the heterogeneity of farmers' responses to weather events. Starting with their canonical Ricardian approach, Mendelsohn, Nordhaus, and Shaw (1994) forecast that climate change will create winners and losers in U.S. agriculture depending on the best use of farmers' land. Schlenker, Hanemann, and Fisher (2005); Schlenker, Hanemann, and Fisher (2006) incorporate space into the Ricardian approach and find heterogeneous producers' resilience to weather shocks depending on geographical characteristics.

⁴Bilateral trade costs are modeled in a multiplicative fashion. Assuming no intra-trade costs (i.e., $t_{ii} = 1$), this modelling decision implies that if region i wants region j to receive one unit of her goods, but bilateral trade costs between these two regions is $t_{ij} = 2$, then she will need to charge for the price of two units of her goods to account for freight costs as shown in 2.1. This modeling approach is known in the literature as iceberg trade costs and represents trade costs in real terms.

cessing plants as j , and we refer to them as importers, consumers and the destination of the commodity interchangeably. A continuum of goods between zero and one is assumed. Food processing plants consume farmers' crops according to a Dixit-Stiglitz technology function depicted in 2.2:

$$Q_j(l) = \left[\int_0^1 Q_i(l)^{\frac{-(1-\sigma)}{\sigma}} dl \right]^{\frac{-\sigma}{(1-\sigma)}}, \quad (2.2)$$

where $\sigma > 0$ is the elasticity of substitution, $Q_j(l)$ denotes the quantity of l consumed by j from the source denoted as $Q_i(l)$. For ease of notation, we drop l from the rest of the description of the model.

Next, the Fréchet distribution, 2.1 and 2.2 are combined to retrieve a distribution of final prices: $G_j(p) = 1 - e^{-\Phi_j p^\theta}$. Here, $\Phi_j = \sum_i^I T_i (c_i t_{ij})^{-\theta}$, where I denotes all producers, ultimately connects *all* regions' technologies, production costs and bilateral trade costs into the buyer's final price. This fact is illustrated in 2.3, where $\gamma = \Gamma\left(\frac{\theta + 1 - \sigma}{\theta}\right)$, and Γ is the gamma function:

$$p_j = \gamma \left[\sum_i^I T_i (c_i t_{ij})^{-\theta} \right]^{\frac{-1}{\theta}}. \quad (2.3)$$

The model is able to simulate a market where each profit-maximizing farmer offers a different price to each expenditure-minimizing food processing plant based on her production technology and the cost of shipping the commodity (2.1). In turn, the food processing plant with a production technology expressed by 2.2 sees all prices and chooses the lowest one. Therefore, the allocation of commodities reached by the farmers and food processing plants leads to a market equilibrium illustrated by 2.3.

To empirically study equilibrium deviations and their welfare consequences, the structural parameters T_i , θ and t_{ij} are connected with trade data. Let $X_j = \sum_i^I X_{ij}$ be expenditure on commodities from all regions, where X_{ij} is trade going from i to j . The probability distribution of prices introduced earlier implies that the fraction of a state's expenditure on crops from state i is equal to the probability that the

state i offers the lowest price. Thus, trade shares are connected to the structural parameters as in 2.4:

$$\frac{X_{ij}}{X_j} = \frac{T_i (c_i t_{ij})^{-\theta}}{\sum_i^I T_i (c_i t_{ij})^{-\theta}}. \quad (2.4)$$

2.4 has two empirical properties regarding short-run substitution behavior in the face of extreme weather events.⁵ By simulating a region's decline in crop yields by reductions in T_k , 2.4 implies that if a region is impacted by an extreme weather event, she will rely on imports. On the other hand, if the extreme weather event hits instead any of her importer partners, then she will import less from the affected state.⁶

Because crop production combines several inputs, we assume that a harvested acre combines the operating inputs labor and land with intermediate inputs coming from other agricultural producers in a Cobb-Douglass fashion (Heerman, 2020):

$$c_i = (w_i^\alpha r_i^{1-\alpha})^\beta p_i^{1-\beta} = \lambda_i^\beta p_i^{1-\beta} \quad (2.5)$$

where, w_i is the labor wage, r_i is land rent rate, and α is the share of operating expenses, λ_i , devoted to labor. Thus, p_i is intermediate input costs with a share of total production expenses of $1 - \beta$. In addition, we normalize 2.4 by multiplying it by domestic sales, $(\frac{X_j}{X_{jj}})$, and thus obtain 2.6. A final step that allows for relative prices to be eliminated from the model is combining the expression for intermediate

⁵Let $\pi_{ij} = \frac{X_{ij}}{X_j}$. Then, $\frac{\partial \pi_{ij}}{\partial T_j} < 0$, i.e., there is an inverse relationship between the expenditure on goods from i and the level of j 's agricultural capacity. And, $\frac{\partial \pi_{ij}}{\partial T_i} > 0$, i.e., there is a positive relationship between the expenditure on goods from i and the level of i 's own agricultural capacity.

⁶Dall'Erba, Chen, and Nava (2021a) provide evidence of producers substituting domestic intermediate inputs for imports when a state is impacted by drought events. Because trade is endogenous and as shown in 2.4 depends on producers' characteristics, econometric models that analyze producers' outcomes with trade variables cannot interpret the coefficient associated with the imports variable as causal. To address the endogeneity concerns, the authors exploit two streams of the literature discussed here. First, a gravity equation is employed in a first-stage. Then, the predicted imports variable from the first-stage is recovered and employed on the second-stage to analyze U.S. farmers' profits.

use with 2.3 into 2.7:

$$\frac{X_{ij}}{X_{jj}} = \frac{T_i}{T_j} \left(\frac{\lambda_i}{\lambda_j} \right)^{-\theta\beta} \left(\frac{p_i}{p_j} \right)^{-\theta(1-\beta)} t_{ij}^{-\theta} \quad (2.6)$$

$$\frac{X'_{ij}}{X'_{jj}} = \left(\frac{T_i^{\frac{1}{\beta}}}{\lambda_i^\theta} \right) \left(\frac{\lambda_j^\theta}{T_j^{\frac{1}{\beta}}} \right) t_{ij}^{-\theta} \quad (2.7)$$

where the logarithm of the left-hand side is given by $\ln X'_{ij} = \ln X_{ij} - \left[\frac{(1-\beta)}{\beta} \right] \ln \left(\frac{X_i}{X_{ii}} \right)$ and $\ln X'_{jj} = \ln X_{jj} - \left[\frac{(1-\beta)}{\beta} \right] \ln \left(\frac{X_j}{X_{jj}} \right)$, which is defined as normalized trade.

2.7 is a gravity-like equation as commonly found in the trade literature, where the expressions in parentheses are the size terms.⁷ In contrast to gravity equations derived from the demand side, normalized trade depends on regional comparative advantage. For example, the size term associated with the exporter increases as she becomes more productive (i.e., T_i increases), but decreases as her production expenses rise (i.e., λ_i increases). The technology-expense relationship is inverted for the importer. The more expensive it is for the importer to produce (i.e., λ_j increases) in her region, the more she relies on imported consumption. In contrast, the more productive she becomes (i.e., T_j increases), the more she consumes domestic production.

Having established the relationship between bilateral trade, we can define our welfare measure as changes in food processing plants' real expenditure on crops. A region's income can be obtained from crop production as $Y_i = c_i L_i$, where L_i is the number of harvested acres in i . Thus, 2.4 can be manipulated to analyze a region's income into 2.8.

$$c_i L_i = \sum_j \frac{T_i (c_i t_{ij})^{-\theta}}{\sum_K T_k (c_k t_{kj})^{-\theta}} X_j. \quad (2.8)$$

Using hat-algebra, we define an equilibrium measure in changes in the structural

⁷A further re-arrangement of parentheses allows us to define w_i^θ and $T_j^{\frac{1}{\beta}}$ as multilateral resistance terms that prevent trade between the two regions.

parameters, given in 2.9:

$$Y_i \hat{c}_i = \hat{T}_i \hat{c}_i^{-\theta} \sum_J \frac{\hat{t}_{ij}^{-\theta} \pi_{ij}}{\sum_K \pi_{kj} \hat{T}_k \hat{c}_k^{-\theta}} \hat{X}_j, \quad (2.9)$$

here, $\hat{c}_i = \frac{c'_i}{c_i}$ is the change in farmers' operating expenses per harvested acre, $\hat{T} = \frac{T'_i}{T_i}$ is the change in agricultural capacity and $\hat{X}_j$ is the change in expenditure. $\pi_{ij} = \frac{X_{ij}}{X_j}$ is a bilateral trade share that indicates how much of j 's expenditure comes from i in the baseline. $\hat{t}_{ij}^{-\theta} = e^\tau$, where $\tau = 0$ if no change in bilateral trade costs is analyzed.

Finally, we define food processing plants' change in nominal expenditure in state i as $\hat{X}_i = \frac{Y_i \hat{c}_i + D_i}{X_i}$, where D_i is an additive component to account for trade imbalances. Then, nominal expenditure change is normalized by change in prices: $\hat{P}_i = [\sum_k \pi_{ki} \hat{T}_k \hat{c}_k^{-\theta}]^{\frac{-1}{\theta}}$. Therefore, 2.9 is reduced into our final expression of welfare given by 2.10, where $\hat{W}_i$ is the change of welfare in i :⁸

$$\hat{W}_i = \frac{\hat{X}_i}{\hat{P}_i}. \quad (2.10)$$

Equilibrium deviations that affect farmers' production such as weather events can be studied by solving the previous system of equations with simulated values for $\hat{T}$. Similarly, we can study equilibrium deviations that increase the cost of trading such as those caused by disruptions in the food supply chains between states with simulated values for $\hat{t}_{ij}^{-\theta}$. We can further study the role of inter-state trade by implementing an autarky simulation. Arkolakis, Costinot, and Rodríguez-Clare (2012) demonstrate that welfare losses caused by moving to autarky can be expressed in terms of the fraction of a state's expenditure on domestic production. Thus, 2.10 is reduced into 2.11:

⁸Because our welfare measures are defined in terms of changes in real expenditures, welfare measures can be comparable across U.S. states. The theoretical model of Eaton and Kortum (2002) models regions as representative consumers, which can be thought of as the average consumer in each region. In that sense, the average consumer in a U.S. state can be compared with that of another state.

$$\hat{W}_i^A = \pi_{jj}^{\frac{1}{\theta}}. \quad (2.11)$$

This intuitive result illustrates that the higher the reliance on imports is, the higher the loss when the level of imports is reduced. Having described the theoretical setting for our analysis of domestic crop trade, we now turn to estimating the model’s key relationships using data on inter-state U.S. trade in crop products.

2.3 Empirical Strategy

The previous section describes our mathematical conceptual framework of U.S. crop producers. We incorporate assumptions about the heterogeneity of farmers’ production and characterize the optimizing behavior of the model’s agents. In this section, we first discuss the data employed in our analysis and how we process our data to mitigate threats to our empirical strategy. Then, we employ the data to recover the necessary structural parameter to run our simulations in a two-step approach. The final sub-section studies the relative comparative advantage of U.S. states in the domestic market for crops.

2.3.1 Data

U.S. crop markets are comprised of multiple producers with heterogeneous production technologies, but there are only three major consumers of U.S. crops: Food processors that buy farmers’ crops to produce their output, exporters that resell U.S. domestic production overseas, and households that buy unprocessed crops in the form of fruits, vegetables and grains. This distinction has empirical implications since substitution effects (i.e., σ in 2.2) are different depending on the consumer. For example, exporters’ demand for domestic crops is determined at the international level, while food processors’ demand for domestic crops is derived from domestic consumers at groceries stores, restaurants, schools and other institutions.

To homogenize consumption and to account for the substitution effects between crops varying across different consumer groups, we drop foreign imports and exports and instead just focus on domestic trade flows. One concern that might arise is that such a restriction limits the comprehensiveness of our analysis of crop production. 2.1 describes the consumption of our two crop aggregations considered in this manuscript as classified by the Standard Classification of Transported Goods (SCTG). SCTG 02 refers to the crop aggregation that considers cereals and grains, while SCTG 03 encompasses fruits, vegetables and oilseeds such as soybeans. The share of U.S. crop production that is exported is roughly 10% with the largest part being SCTG 02. In addition, our conceptual mathematical model considered in the previous section characterizes food processing plants as the main consumer of farmers' crops. Only 12% (13% after excluding exports) of U.S. crop production goes to households, so the larger part of U.S. crop production goes to food processing plants.

Our data come from the fourth version of the U.S. Freight Analysis Framework (FAF4) 2017 prepared for the Bureau of Transportation Statistics (BTS) and the Federal Highway Administration (FHA). FAF4 imputes domestic data flows between all U.S. states for several commodities and is constructed using the Commodity Flow Survey (CFS) for most types of commodities; however, because the CFS treats agricultural goods as outside of its scope, the FAF4 dataset uses information from the USDA's Census of Agriculture to more accurately model the true extent of domestic trade flows in farm products. Our measure of crop trade flows is constructed by aggregating SCTG 02 (cereal crops) and SCTG 03 (fruits and vegetables, and a handful of other crops such as soybeans).

2.1 describes the data for each of the variables used in our analysis. The top part of the table describes dyadic variables that capture attributes about the relationship between two states. Here, we consider distance and whether two states share the same border. We calculate distance between polygons' centroids in a shapefile, using the Geoda software. Our measure of distance is in kilometers. The bottom part of the table describes monadic variables that capture attributes specific to individual states that will be used to estimate the comparative advantage dispersion parameter θ . Imports and exports are created using FAF4 data to describe trade flows. Both




Figure 2.1: Consumption of U.S. crop production.

Note: SCTG 02 includes cereals such as wheat, corn and rye. SCTG 03 includes fruits, vegetables and oilseeds such as soybeans. Analysis focuses on domestic production and consumption for crops. The share of SCTG 02 production destined to exports is 8.1%, while the share of SCTG 03 production destined to exports is 2.4%. The share of domestic SCTG 02 production consumed by U.S. households is 2.4%, and the share of SCTG 03 production consumed by U.S. households is 9.2%.

Source: Calculations are our own with data from the U.S. Freight Analysis Framework (FAF4) 2017 and IMPLAN: Economic Impact Analysis for Planning.

measures are in millions of 2017 U.S. dollars. Population density, which we will use as an exogenous predictor of operating costs in an IV analysis below, is collected from the 2010 U.S. Census Bureau of Statistics and is measured in population per square kilometer. Operating expenses per harvested acre comes from the 2017 U.S. Census

of Agriculture and is measured in 2017 dollars. Climate variables that affect agricultural output during the growing season are precipitation and temperature during summer. Precipitation is measured in cubic centimeters, and temperature is measured in degrees Celsius. Weather data are collected from the Parameter-elevation Regression on Independent Slopes Model (PRISM) data base over the 20 years period between 1997 and 2017. Then, weather data are transformed into climate by taking the average of each variable over time for each location (PRISM, 2011).

Table 2.1: Trade, crop operating expenses, population density and climate variables in 2017

Dyadic Variables	Mean	S.D.	Min	Max
Distance	1827.27	1295.40	0.00	5179.18
Contiguity	0.10	0.29	0.00	1.00
Monadic Variables	Mean	S.D.	Min	Max
Imports	9,280.78	9,934.75	233.99	43,387.97
Exports	9,280.78	10,927.41	105.50	46,177.39
Operating expenses per harvested acre	708.75	492.77	75.83	2478.09
Population density	83.85	61.69	6.46	286.11
Summer temperature	24.25	3.70	18.27	32.30
Summer precipitation	42.27	16.68	3.33	82.93

Note: The number of observations for our dyadic variables is 2,304, while the number of observations for our monadic variables is 48. All monetary measures are in 2017 U.S. dollars. The value of imports and exports are created using FAF4. Population density is measured in population per square kilometers. Precipitation is measured in cubic centimeters, and temperature is measured in degrees Celsius.

Because virtually all agricultural activity occurs outside cities and our climate variables described earlier could be influenced by urbanization, we follow a two-step aggregation of our data variables. First, all data are obtained at the county level for all counties in the contiguous United States. Next, metropolitan and highly populated counties are dropped.⁹ Finally, we calculate state averages excluding metropolitan and high population counties. This aggregation technique allows us

⁹The population density cutoff at the county level is 1,600. The most populous county included in our dataset is Lake county, IL with an average farmland value of \$12,149 in 2017 according to the U.S. Department of Agriculture-National Agricultural Statistics Service.

to reduce weight from highly populated, urbanized regions with little agricultural production.

2.3.2 Recovering structural parameters

Recovering the structural parameters θ and T_i requires a two-step approach. First, 2.7 is parameterized with the log-linearization of the expression. 2.12 is the regression model, where the size terms are denoted by S_i for the exporter and S_j for the importer and are approximated by importer and exporter fixed effects. In our notation, we keep the hypothesized sign, since the estimated coefficients should conform with the direction of the sign. The error term is ϵ_{ij} . Then, size term estimates are used in a second-step as the dependent variable in the estimation of 2.13, where the size terms come from 2.7 and are denoted as $\hat{S}_i$ and $\hat{S}_j$ for the exporter and the importer respectively.

$$\ln \frac{X'_{ij}}{X'_{jj}} = -\theta \ln t_{ij} + S_i - S_j + \epsilon_{ij} \quad (2.12)$$

$$\hat{S}_k = \alpha + \frac{1}{\beta} \ln T_k - \theta \ln \lambda_k + \nu_k \quad (2.13)$$

In contrast with 2.12, θ can be recovered from 2.13 since it is associated with an observable variable. Here, ν_k is the error term. Our measure of the share of intermediate inputs is .12 on average (i.e., β in 2.7) as reported by IMPLAN: Economic Impact Analysis for Planning.¹⁰ The estimation of 2.12 is carried out through the use of dummy variables. Since bilateral trade costs are not observed, we approximate them by a contiguity dummy that indicates whether two states share the same border and six dummies corresponding to the distance bins (0,865], (865,1730], (1730,2595], (2595,3460], (3460,4325] and (4325, Max]. Since we assume that intra-trade has zero


¹⁰We consider $\beta = 0.12$, the weighted average of our crop aggregation, despite that some crops in our aggregation require higher operating expenses e.g., fruits and vegetables in SCTG 03. For robustness, we test several values for β , and we find that estimates within $\beta \in [0.08, 0.21]$ are not statistically different from each other, and simulation conclusions remain unaffected by the different estimates.

cost (i.e., $t_{ii} = 1$), then all dummies for distance are interpreted relative to intra-trade costs. Importer and exporter fixed effect dummy variables are included, but normalized to sum to zero. Notice that the estimation of 2.12 will drop all zero-trade observations (i.e., $X_{ij} = 0$). Because trade between regions is not symmetric, the variance-covariance matrix is assumed to have diagonal elements $\sigma_1^2 + \sigma_2^2$ that affect both two-way trade and one-way trade, and certain nonzero off-diagonal elements σ_2^2 that affect only two-way trade. For this reason, we estimate 2.12 by feasible generalized least squares (FGLS).


2.2 reports estimates from 2.12. In panel (a), we report the proxy estimates for bilateral trade costs (aggregated). As expected, two states sharing a contiguous border tends to be associated with more trade between the states. The trade literature has long supported an inverse relationship between trade and distance. Our results indicate that trade volumes decrease with distance, but the marginal impact of distance on trade becomes attenuated in the last two distance bins. We attribute this to freight mode decisions. While most crop production is shipped within short distances by truck, the volume shipped by truck decreases with distance but increases for rail and barges.

In panel (b), estimates of the fixed effects are depicted graphically. The estimates associated with the size terms are structurally symmetric as shown in 2.7. Deviations from symmetry represent non-market influences that affect either the imports or exports from a region. For instance, as described earlier, each set of fixed effects can be interpreted as either competitiveness or openness for the exporter and importer side respectively. The distinction reflects institutional and policy differences that affect specific countries such as weak government institutions, affecting exports, and policies that encourage citizens to consume domestically sourced goods, affecting imports. We test for symmetry of the fixed effect estimates by testing the difference between each pair of estimates. Furthermore, the imposed symmetry of our structural model prevents us from testing whether point estimates are statistically different from zero since each set of size-term fixed effects crosses zero. For this reason, the symmetry test also serves as a significance test. The point estimates are shown without confidence intervals, but the confidence interval corresponding to the test of

whether each pair is equal to each other is shown.



(a) Bilateral trade costs



(b) Size terms

Figure 2.2: Estimation results from gravity equation: Distance and importer- and exporter-size effects.

Note: R^2 is 0.97 with 1,334 observations. The imposed symmetry of our structural model prevents us from testing whether point estimates are statistically different from zero since each set of size term fixed effects crosses zero. For this reason, we instead test whether the absolute values of point estimates in each pair of size terms is equal to each other. The point estimates are shown without confidence intervals, but the confidence interval for the test of whether each pair is equal to each other is shown.

Using estimates for the size terms, we estimate 2.13 to recover θ . Since agricultural capacity (T_k) is not observed, we proxy it using historical precipitation and temperature for each U.S. state, following a quadratic polynomial with an interaction term. Because we expect high operating expenses per acre to be associated with high levels of agricultural capacity (e.g., high operating expenses might reflect more productive soil), we employ the two-stage least squares (2SLS) estimator. Our instrumental variable is average population density in rural counties within each U.S. state. We expect that the wage component in 2.5 is affected by population density, so our use of this variable ensures that our instrument affects operating expenses per acre, but not other determinants of agricultural capacity such as climate.¹¹

¹¹The threat to the exclusion restriction when population density is the aggregate of the whole state is that climate change is affected by activities related to population density such as traffic and polluting industries, which are typically concentrated in cities and their surrounding areas. Our instrument is weakly and negatively correlated with temperature in the growing season. In fact, a large part of the variation in agricultural productivity is explained by climate and land characteristics, ensuring that population density only affects operating expenses per harvested acre (Liang et al., 2017).

2.2 reports results from the estimation of 2.13. The exporter columns use the size terms related to the exporter, and the importer columns use the size terms related to the importer. The first column employs the OLS estimator and the second column the 2SLS estimator. Variance inflation factors of climate variables are above 100 in every regression, so climate estimates are highly uninformative.¹² Our conceptual framework predicts θ to be positive and greater than 1, so the coefficient associated with operating expenses per harvested acre must be negative and less than 1. The results from OLS estimations show that θ is not statistically different from zero. As expected, correcting for endogeneity brings larger and statistically different from zero estimates of θ . The exporter size terms indicates that $\theta = 2.962$, but the importer size terms indicates that $\theta = 3.804$. 2.2 also reports the first stage and demonstrates that our instrument is not a weak one.¹³

Therefore, we rely on the work of Reimer and Li (2010) to further test the robustness of our estimate. Their point estimates for θ are no lower than 2.52 and no higher than 4.96, and employ several estimation techniques. Reimer and Li (2010) estimate $\theta = 2.83$ using a generalized method of moments approach with crop data. Additionally, Reimer and Li (2010) report that their lowest $\theta = 2.52$ uses maximum likelihood and a parametrization of the Fréchet distribution. The authors also report that their largest estimate of $\theta = 4.96$ is based on a proxy estimation using the relative prices of the commodities analyzed, but Simonovska and Waugh (2014) show the use of such proxies systematically overestimates θ . Therefore, our importer side estimate of $\theta = 2.962$ falls slightly above the small range of valid estimates found by Reimer and Li (2010). We choose $\theta = 2.962$ as our preferred value for comparative advantage because it is closer to those found in their study.

¹²Because the coefficient of interest is θ , climate proxies are simply used as control variables.

¹³A concern associated with the estimations in 2.2 is the finite sample properties of the IV estimator (Heid, Larch, and Yotov, 2017). Andrews and Armstrong (2017) report that exactly identified models with low number of observations have undesired properties (i.e., IV estimator is consistent, not unbiased), but the authors propose an estimator that is unbiased. We implement their estimator and find no difference with the estimations in 2.2.

Table 2.2: Second step estimations, using OLS and 2SLS estimators

	Exporter		Importer	
	OLS	2SLS	OLS	2SLS
Labor expenses	-0.008 (1.138)	-2.962** (1.344)	-0.497 (1.175)	-3.804*** (1.399)
Temperature	2.643 (2.899)	2.301 (2.968)	3.433 (2.995)	3.049 (1.399)
Precipitation	-0.167 (0.325)	-0.032 (0.313)	-.136 (0.336)	0.015 (0.315)
Temperature ²	-0.056 (0.058)	-0.044 (0.054)	-0.068 (0.059)	-0.055 (0.054)
Precipitation ²	-0.001 (0.003)	0.001 (0.003)	0.001 (0.003)	0.002 (0.003)
Interaction	0.007 (0.015)	-0.001 (0.018)	0.002 (0.016)	-0.008 (0.018)
Constant	-29.864 (37.515)	-10.245 (37.750)	-36.514 (38.748)	-14.545 (37.9423)
	First Stage			
Population Density	0.709*** (0.085)		0.709*** (0.085)	
F-statistic	14.81		14.81	
R ²	0.10	0.59	.08	0.59
Observations	48	48	48	48

Note: Operating expenses and population density are natural logarithm values. Coefficient associated with operating expenses is θ . Values in parentheses are robust standard errors. Coefficients for climate variables are recovered up to the constant $\frac{1}{\beta}$; *p < 0.10; **p < 0.05, ***p < 0.01.

2.3.3 Comparative advantage

To analyze the comparative advantage of each U.S. state, 2.13 is re-arranged into the following equality: $\hat{T}_i = (e^{\hat{S}_i} \lambda_i^\theta)^\beta$. Using labor expenses per harvested acre, the estimates for the exporters' size term, our preferred value for $\theta = 2.962$, and $\beta = 0.12$, we calculate the value of agricultural capacity of each state. Thus, each state's exporter size terms can be decomposed into an agricultural capacity component and a labor expense component. This decomposition is shown in 2.3.




Figure 2.3: Comparative advantage (exporter size-term) decomposition.

Note: Figure reports size terms estimates obtained using 2.13: natural log of the implied level of agricultural capacity and natural log of operating expenses per harvested acre in each U.S. state multiplied by $\theta = 2.962$. Implied level of agricultural capacity is calculated as $T_i = (e^{\hat{S}_i} \lambda_i^\theta)^\beta$.

All else constant, 2.7 implies that net exporters have high levels of agricultural capacity relative to their operating expenses ($T_i^{\frac{1}{\beta}} > \lambda_i^\theta$). This is because regions with relatively high levels of agricultural capacity relative to their input costs are better able to specialize in crop production. Similarly, net importers have low levels of agricultural capacity relative to their operating expenses ($T_i^{\frac{1}{\beta}} < \lambda_i^\theta$). Therefore, the size term can be interpreted as a state's comparative advantage. 2.3 shows that a state's comparative advantage largely depends on its agricultural capacity. While operating expenses per harvested acre remain close to the national average, agricultural capacity can take extreme values. When agricultural capacity surpasses the values of operating expenses, the state becomes a net exporter.

Our results indicate that Kansas, South Dakota, Iowa, California, Nebraska, Indiana, Illinois and Florida have a comparative advantage in the U.S. market for

crops. Not surprisingly, only these eight states jointly account for over half of all domestic crop flows. However, comparative advantage in crop production is not only determined by agricultural capacity. Agricultural capacity is largest for California and Florida, but their comparative advantage is dragged down by high operating expenses within the state. This emphasizes the importance of reducing operating expenses to compete. On the other hand, Louisiana, Rhode Island, New Hampshire, Utah, Alabama and Nevada have the lowest levels of agricultural capacity and high operating expenses. These six states export less than 2% of domestic crop production and import up to four times what they export.

In the next section, we employ the structural parameter for comparative advantage and our trade data to run three simulations. First, we run an autarky simulation such that inter-state trade is shut down. Results from this simulation represent the extreme-case scenario in which U.S. states are forced to rely on local agriculture. The second simulation is with U.S. states vulnerability to food supply chain disruptions. Our simulation increases the cost to import food to each state separately to study what states are the most vulnerable to food supply disruptions. While the state’s comparative advantage can help mitigate welfare losses caused by supply chain disruptions, we conduct an additional simulation to investigate the ability of the existing food supply chain to mitigate welfare losses from events that affect each states’s ability to produce crops.

2.4 Food vulnerability simulation results

In section 2, we described our general equilibrium model that connects farmers and food processing plants in the U.S. through domestic crop trade. Then, we conclude with the sufficient conditions to analyze welfare effects from equilibrium deviations i.e., our domestic crop trade data (2.1) and our structural parameter $\theta = 2.962$ (2.2). In this section, we employ simulations to assess the food vulnerabilities of the existing U.S. domestic food system to supply chain disruptions with 2.10. Our supply chain disruption doubles the cost of importing crops for each U.S. state (i.e.,


$t_{ij} = 1.5 \forall j \neq i$). In addition, we conduct a simulation to investigate the ability of the existing food supply chain to mitigate welfare losses from events that affect each state's ability to produce crops. Our production shock simulation decreases by 50% each U.S. state's agricultural capacity (i.e., $T_i = 0.5 \forall i$). Before discussing our food vulnerability results, however, we run an autarky simulation to further study the mechanisms that shape the distribution of welfare losses with 2.11. All our state simulations are ran separately to control for spillovers from other states' counterfactuals. The descriptive statistics of all our simulation results are shown in 2.3. Notice that we decompose our welfare measure into a total expenditure term, $\hat{W}_i$, and a price index term, $\hat{P}_i$, so our welfare measure is real expenditure for crops (Arkoulakis, Costinot, and Rodri  guez-Clare, 2012). We solve our model using the tools reproduced by Baier, Yotov, and Zylkin (2019) and the Stata software.

Table 2.3: Description of simulation results


Autarky				
	Mean	(S.D.)	Minimum	Maximum
$\hat{W}_i$	0.79	(0.137)	0.43	0.97
$\hat{X}_i$	0.92	(0.371)	0.25	1.76
$\hat{P}_i$	1.12	(0.355)	0.46	1.96
Food supply chain shock				
	Mean	(S.D.)	Minimum	Maximum
$\hat{W}_i$	0.82	(0.072)	0.66	0.96
$\hat{X}_i$	0.84	(0.069)	0.67	0.96
$\hat{P}_i$	1.02	(0.010)	1.01	1.05
Production shock				
	Mean	(S.D.)	Minimum	Maximum
$\hat{W}_i$	0.89	(0.081)	0.71	0.99
$\hat{X}_i$	1.12	(0.038)	1.04	1.19
$\hat{P}_i$	1.27	(0.107)	1.07	1.45

Note: Table shows descriptive statistics from simulation results for all U.S. states. Autarky results are calculated with 2.11, while food supply disruptions and production shocks are calculated with 2.10. Our autarky simulation is the extreme case in which states are forced to consume only production within their borders. Our production shock simulation is interested with scenarios in which farmers' productivity declines by half. Our food supply chain disruption simulation is interested with scenarios in which trade costs double.

On average, our autarky simulation indicates that U.S. states' real expenditure on crops should decline by 21% when U.S. states are forced to rely on their own crop production, and exports and imports are restricted. Assuming that crop demand within the state is equal to its total crop production, the simulated restrictions have two opposing effects on prices: demand for the state's crop decreases while supply for crops in the state increases. Thus, the net effect on crop prices depends on the initial food processing plants' consumption source composition. In 2.4, we describe the heterogeneity of our welfare results both geographically in panel (a), but also re-arranged following the comparative advantage pattern illustrated in 2.3 in panel (b). States with large reliance on imports reduce their nominal expenditure for crops the most, pushing prices down; but also, states with large reliance on imports have the most severe welfare losses. For instance, Louisiana, New Hampshire, Utah and Alabama have welfare losses of more than 40% despite having price declines ranging from 47% for Louisiana and 54% for Alabama. On the other hand, the states with the lowest import dependency, which in addition are net crop exporters, have the largest increases in nominal crop expenditure, pushing prices upward. In contrast to net crop importers, these states' welfare losses are the least severe. Kansas, South Dakota and Iowa perceive price increases between 20% for Iowa to 51% for Kansas, but their welfare losses are not higher than 7%.



(a) Spatial distribution




(b) Comparative advantage distribution

Figure 2.4: Autarky simulation results.


Note: Figure describes the spatial (top) and comparative advantage (bottom) distribution of our autarky simulation welfare results. Our autarky simulation is the extreme case in which states are forced to consume only production within their borders and is calculated with [2.11](#). The horizontal axis in the bottom panel is arranged by the states' comparative advantage i.e., [2.3](#).

The results from our autarky simulation suggest that the state's ability to substitute imports for domestic consumption plays a role when determining the welfare loss from disruptions in the supply chain. In [2.5](#), we test this hypothesis with more mild but realistic restrictions to the U.S. domestic supply chain. Similarly to [2.4](#),

panel (a) shows the spatial distribution of our welfare results, and panel (b) their comparative advantage distribution. At first glance, we observe that the patterns depicted in 2.4 and 2.5 are similar; the higher the comparative advantage, the lower the welfare loss in our food vulnerability simulation. The net effect on prices in our food supply disruption is higher crop prices at the state level, which in turn increase U.S. states nominal crop expenditure. Our results indicate, however, that the lower prices are observed in places where comparative advantage is high such as Kansas, North Dakota, South Dakota and California. While food processing plants face a shortage of imports, domestic crop production is large enough to substitute import losses. In addition, states with large comparative advantage also experience the lower welfare losses. On the other hand, states on the leftmost part of our comparative advantage spectrum experience the highest welfare losses and the highest crop prices, suggesting that food processing plants within those states have a difficult time substituting imports for domestic crop production.



(a) Spatial distribution




(b) Comparative advantage distribution

Figure 2.5: Food supply chain disruption simulation results.


Note: Figure describes the spatial (top) and comparative advantage (bottom) distribution of our food supply chains disruption simulation welfare results. Our food supply chains disruption simulation consists in doubling each state's costs to import crops from other states and is calculated with 2.10. The horizontal axis in the bottom panel is arranged by the states' comparative advantage i.e., 2.3.

While the state's comparative advantage can help mitigate welfare losses caused by food supply chain disruptions, in 2.6 we analyze if the existing food supply chain can help mitigate the losses from events that affect each state's ability to produce

crops by substituting domestic consumption for imports. Similarly to 2.4, panel (a) shows the spatial distribution of our welfare results and panel (b) their comparative advantage distribution. For our production shock simulation, we expect the largest losses to be associated with the states with the highest comparative advantage; but by limiting our focus to the rightmost states only, we can study the role of the existing food supply chain on substituting domestic consumption for imports. Our results suggest that the initial level of imports is a major determinant to shape welfare results from our simulation. Food processing plants in Montana, Kansas and North Dakota rely heavily on their states' crop production for their production, and little on imports. These states differ largely on their implied comparative advantage (i.e., 2.3), but their nominal food expenditure and prices change in similar levels. On the other hand, states with a more balance composition of imports and domestic consumption have the least severe impacts, suggesting that inter- and intra-supply chains are developed at similar levels. Substitution of domestic consumption for imports is more feasible in these states.



(a) Spatial distribution



(b) Comparative advantage distribution

Figure 2.6: Production shocks simulation results.

Note: Figure describes the spatial (top) and comparative advantage (bottom) distribution of our production shocks simulation welfare results. Our production shocks simulation consists in reducing each state's agricultural productivity by half separately and is calculated with [2.10](#). The horizontal axis in the bottom panel is arranged by the states' comparative advantage i.e., [2.3](#).

2.5 Conclusions

Our extension of the Ricardian trade model of Eaton and Kortum (2002) is a comprehensive but simple model of the U.S. market for crops that can be employed to analyze a large array of counterfactuals concerning food security within the United States. We incorporate assumptions about the heterogeneity of farmers' production and characterize the optimizing behavior of the model's agents. Our conceptual framework also considers the realities of the U.S. domestic food system and models food processing plants as the main consumer of farmers' crops. An advantage of our approach is that we consider the production side, so our model permits an analysis of market mechanisms that govern crop trade, expenditure and prices through a specification of U.S. states' comparative advantage. Finally, we decompose each U.S. state's comparative advantage in the market for crops into an agricultural capacity term and an operating expenses term, which can be used to study the determinants of simulation welfare results.

We implement our general equilibrium model of interstate crop trade to assess the food vulnerabilities of the existing U.S. domestic food supply chain. Our simulation results provide insights into U.S. states food vulnerabilities through food processing plants' expenditure on crops, crop prices, and overall consumer welfare. Three counterfactual simulations are considered in this manuscript. The first counterfactual is autarky. We simulate the extreme case in which U.S. states are forced to rely on their own crop production. The second counterfactual is concerned with supply chain disruption that increases the trade costs between states. We demonstrate that U.S. states' food vulnerability to disruptions in the food supply chains or production shocks can be explained by their relative comparative advantage in the domestic market for crops. We find that the distribution of welfare losses caused by disruptions in the existing supply chains depends on a state's ability to substitute imports for domestic production. We conduct an additional simulation to investigate the ability of the existing food supply chain to mitigate welfare losses from events that affect each state's ability to produce crops, arriving to similar conclusions regardless of a state's initial comparative advantage. The policy implications of our study suggest

that food supply chains must not be unidirectional, but they must both look inward as well as outward.

2.6 Chapter 2 References

- Anderson, James E (1979). “A theoretical foundation for the gravity equation”. In: *American Economic Review* 69(1), pp. 106–116.
- Anderson, James E and Eric van Wincoop (2003). “Gravity with gravitas: A solution to the border puzzle”. In: *American Economic Review* 93(1), pp. 170–192.
- Andrews, Isaiah and Timothy B Armstrong (2017). “Unbiased instrumental variables estimation under known first-stage sign”. In: *Quantitative Economics* 8(2), pp. 479–503.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare (2012). “New trade models, same old gains?” In: *American Economic Review* 102(1), pp. 94–130.
- Baier, Scott L, Yoto V Yotov, and Thomas Zylkin (2019). “On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration”. In: *Journal of International Economics* 116, pp. 206–226.
- Biedny, Christina, Trey Malone, and Jayson L Lusk (2020). “Exploring Polarization in U.S. Food Policy Opinions”. In: *Applied Economic Perspectives and Policy* 00(00), pp. 1–21.
- Blay-Palmer, A, K Landman, I Knezevic, R Hayhurst, et al. (2013). “Sustainable local food spaces: constructing communities of food.” In: *Local Environment* 18(5), pp. 521–641.
- Burgess, Robin and Dave Donaldson (2010). “Can openness mitigate the effects of weather shocks? Evidence from India’s famine era”. In: *American Economic Review* 100(2), pp. 449–53.

- Clapp, Jennifer (2015). “Distant agricultural landscapes”. In: *Sustainability Science* 10(2), pp. 305–316.
- Dall’Erba, Sandy, Zhangliang Chen, and Noé J Nava (2021a). “US Interstate Trade Will Mitigate the Negative Impact of Climate Change on Crop Profit”. In: *American Journal of Agricultural Economics*.
- Deschenes, Olivier and Michael Greenstone (2007). “The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather”. In: *American Economic Review* 97(1), pp. 354–385.
- DuPuis, E Melanie and David Goodman (2005). “Should we go “home” to eat?: Toward a reflexive politics of localism”. In: *Journal of Rural Studies* 21(3), pp. 359–371.
- Eaton, Jonathan and Samuel Kortum (2002). “Technology, geography, and trade”. In: *Econometrica* 70(5), pp. 1741–1779.
- Gephart, Jessica A, Elena Rovenskaya, Ulf Dieckmann, Michael L Pace, and Åke Brännström (2016). “Vulnerability to shocks in the global seafood trade network”. In: *Environmental Research Letters* 11(3), p. 035008.
- Greene, Christina (2018). “Broadening understandings of drought—the climate vulnerability of farmworkers and rural communities in California (U.S.A.)” In: *Environmental Science & Policy* 89, pp. 283–291.
- Heerman, Kari ER (2020). “Technology, ecology and agricultural trade”. In: *Journal of International Economics* 123, p. 103280.
- Heid, Benedikt, Mario Larch, and Yoto Yotov (2017). “Estimating the effects of non-discriminatory trade policies within structural gravity models”. In: *CESifo Working Paper Series*.
- Hobbs, Jill E (2020). “Food supply chains during the COVID-19 pandemic”. In: *Canadian Journal of Agricultural Economics*.

- Keller, Wolfgang, Javier Andres Santiago, and Carol H Shiue (2017). “China’s domestic trade during the Treaty-Port Era”. In: *Explorations in Economic History* 63, pp. 26–43.
- Krugman, Paul R (1981). “Intraindustry specialization and the gains from trade”. In: *Journal of Political Economy* 89(5), pp. 959–973.
- Levkoe, Charles Zalman (2011). “Towards a transformative food politics”. In: *Local Environment* 16(7), pp. 687–705.
- Liang, Xin-Zhong, You Wu, Robert G Chambers, Daniel L Schmoldt, Wei Gao, Chaoshun Liu, Yan-An Liu, Chao Sun, and Jennifer A Kennedy (2017). “Determining climate effects on U.S. total agricultural productivity”. In: *Proceedings of the National Academy of Sciences* 114(12), E2285–E2292.
- Marsden, Terry (2013). “Sustainable place-making for sustainability science: The contested case of agri-food and urban–rural relations”. In: *Sustainability Science* 8(2), pp. 213–226.
- PRISM (2011). “PRISM climate data”. In: *Oregon State University*.
- Puma, Michael J, Satyajit Bose, So Young Chon, and Benjamin I Cook (2015). “Assessing the evolving fragility of the global food system”. In: *Environmental Research Letters* 10(2), p. 024007.
- Reimer, Jeffrey J and Man Li (2009). “Yield variability and agricultural trade”. In: *Agricultural and Resource Economics Review* 38(2), pp. 258–270.
- Reimer, Jeffrey J and Man Li (2010). “Trade costs and the gains from trade in crop agriculture”. In: *American Journal of Agricultural Economics* 92(4), pp. 1024–1039.
- Rodriguez, Lizette, Marcel Horowitz, Dorina Espinoza, Alberto Aguilera, Adela de la Torre, and Lucia L Kaiser (2015). “The impact of the California drought on food security among rural families of Mexican origin”. In: *Journal of Applied Research on Children: Informing Policy for Children at Risk* 6(2), p. 11.

- Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher (2005). “Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach”. In: *American Economic Review* 95(1), pp. 395–406.
- Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher (2006). “The impact of global warming on U.S. agriculture: an econometric analysis of optimal growing conditions”. In: *Review of Economics and statistics* 88(1), pp. 113–125.
- Simonovska, Ina and Michael E Waugh (2014). “The elasticity of trade: Estimates and evidence”. In: *Journal of International Economics* 92(1), pp. 34–50.
- Thilmany, Dawn, Elizabeth Canales, Sarah A Low, and Kathryn Boys (2020). “Local Food Supply Chain Dynamics and Resilience during COVID-19”. In: *Applied Economic Perspectives and Policy*.
- Tortajada, Cecilia, Matthew J Kastner, Joost Buurman, and Asit K Biswas (2017). “The California drought: Coping responses and resilience building”. In: *Environmental Science & Policy* 78, pp. 97–113.
- Yotov, Yoto V, Roberta Piermartini, José-Antonio Monteiro, and Mario Larch (2016). *An advanced guide to trade policy analysis: The structural gravity model*. World Trade Organization Geneva.

CHAPTER 3

REVISITING THE IMPACT OF CLIMATE CHANGE ON AGRICULTURE THROUGH SPATIALLY-VARYING AND PLACE-TAILORED RICARDIAN ESTIMATES

The Ricardian framework has been widely used to study the impact of climate change on agriculture across US counties over the past few decades. While the spatial heterogeneity of climate change is well-accepted, the literature struggles to reach an agreement on how to model it, hence leading to a wide range of forecasted impacts. This paper employs Multi-scale Geographically-Weighted-Regression (MGWR) to avoid setting an a priori definition of heterogeneity and to generate county-specific marginal effects of climate change impacts. This manuscript tests the predictive power of our MGWR application with other functional forms found in the literature on a homogenized dataset of historical climate, demographic and soil quality controls. Our cross-validation exercise indicates that our MGWR approach has a higher predictive power than studies that cluster spatial units, and that other approaches have a downward bias. We attribute the divergence in results to unspecified heterogeneity. Our place-specific marginal effects will help guide the development of place-tailored mitigation and adaptation strategies to climate change.

3.1 Introduction

The estimation of damages from global warming, particularly in the agricultural sector, is receiving an increasing amount of attention. The Ricardian approach, a

reduced-form hedonic analysis developed by Mendelsohn, Nordhaus, and Shaw (1994) to calibrate such an impact over a single cross-section of data, has been extensively employed to predict climate-change-induced damages on agriculture across spatial units (farms, municipalities, regions) of the country under study, including the U.S. (Schlenker, Hanemann, and Fisher, 2005; Schlenker and Roberts, 2006; Dall’erba and Domínguez, 2016), China (Liu et al., 2004; Wang et al., 2009) and South-Africa (Gbetibouo and Hassan, 2005).¹ For a single country, however, the consequences of climate change can diverge widely across studies. Cai and Dall’Erba (2021) investigate the role of group-membership i.e., clustering spatial units that are thought to respond similarly to climate change, in the U.S. Their findings suggest that a priori definitions of group-membership are the single determinant that explains the divergence in results across Ricardian approach studies for U.S. counties.

The global nature of the marginal effect of climate on spatial units has motivated the development of such grouping strategies. Early contributions of the Ricardian literature (see the work of Mendelsohn, Nordhaus, and Shaw (1994) and Mendelsohn and Dinar (2003)) have been criticized for not fully accounting for region specific features. Subsequent studies, such as Schlenker, Hanemann, and Fisher (2005), Deschenes and Greenstone (2007), and Dall’erba and Domínguez (2016), have remedied this shortcoming by grouping counties based on irrigation status, state boundaries and elevation. In addition, studies add a quadratic term in climate such that the marginal effect is dependent on space; but as we demonstrate in the next section, a quadratic term does not fully incorporate the role of geographic determinants into the estimated impact of climate. These approaches allow the generation of between two and fifty different marginal effects for each covariate. However, Cai and Dall’Erba (2021) indicates that among the county groupings chosen in the literature, there is little evidence that one grouping outperforms the other in terms of predictive power.

Ex-ante grouping definitions raise endogeneity concerns, so a potential solution to this confusion is to generate spatially varying estimates without any a priori clustering of counties via local regression techniques. Then, clear county-specific marginal

¹For an exhaustive review of these Ricardian applications see Mendelsohn and Massetti (2017).

effects can be calculated and place-tailored mitigation strategies can be derived from them. Several empirical applications of local regression techniques have outperformed global approaches in terms of their predictive power (see, among others Páez, Long, and Farber (2008), Zhang et al. (2011) and Soler and Gemar (2018)). However, few applications of local regression models focus on agriculture and, when they do, their outcome variable is crop yield.² Consequently, spatial heterogeneity in the marginal effects of climate change has been underexplored. We focus on the United States, a country that has been subject to a large number of previous county-level studies with which we can compare our approach using the newly-developed Multiscale Geographically Weighted Regression (MGWR) framework (Fotheringham, Yang, and Kang, 2017).³ By letting the data speak for themselves, this paper seeks to contribute to both the locally weighted regression and the Ricardian literature with the first hedonic study of U.S. agriculture that appropriately incorporates spatial heterogeneity.

Our primary contribution, therefore, is to demonstrate the predictive gains from using the MGWR estimator compared to the global approach, with or without spatial heterogeneity, used in the works of Mendelsohn, Nordhaus, and Shaw (1994), Dall’erba and Domínguez (2016) and Schlenker, Hanemann, and Fisher (2005). In all four cases, we rely on the same data, measured in 2012, and the same functional form and covariate choices as defined in Mendelsohn, Nordhaus, and Shaw (1994). We are aware that the more recent contributions of Dall’erba and Domínguez (2016) and Schlenker, Hanemann, and Fisher (2005) use more and different covariates, and that other contributions, such as Deschenes and Greenstone (2007), have extended the Ricardian approach to a panel context. Nevertheless, it is only by comparing models for which only difference is the definition of spatial heterogeneity that we can assess the relative advantage of our MGWR compared to global approaches. We do

²See for example Shiu and Chuang (2019).

³In contrast to other GWR methods, the MGWR is more flexible in how spatial heterogeneity is modelled since it allows the calibration of coefficients to have different bandwidths depending on the variable i.e., each covariate is regressed across different subsets of observations. For example, if the nature of one variable is global and the rest are not, MGWR will assign a bandwidth equal to the total of observations to the former covariate and a lower bandwidth to the latter one.

so by comparing the capacity of each model to accurately predict the 2017 observed values when relying on marginal effects calibrated in 2012. We find that our MGWR estimator outperforms previous handling of the spatial heterogeneity. For example, previous approaches predict losses when we observe gains, and only our MGWR predicts gains. We attribute this difference to a miss-specification of clustering in previous studies that can be corrected with the use of MGWR.

In the next section, we develop the conceptual framework of our spatial non-stationary Ricardian-approach that extends the original idea of Mendelsohn, Nordhaus, and Shaw (1994) to a local, multiscale-structure. We describe in section 3 the data employed in our estimations and cross-validation exercise. Contrary to the a priori clustering approaches adopted by previous studies, MGWR allows us to be agnostic about any possible clustering of counties in terms of their response to climate change. In order to show the appeal of this approach to our case study, we employ in section 4 a cross-validation approach on the data observed in 2017. Results provide evidence of the relative greater predictive power of our local approach compared to any previous specification.

3.2 Spatial non-stationary Ricardian-approach

The idea at the core of the Ricardian approach is that landowners adapt their production choices to new climate conditions. The decisions made vis-à-vis the optimal use of the land are reflected in the land price which is commonly modeled as the discounted sum of future net returns (Mendelsohn, Nordhaus, and Shaw, 1994; Plantinga, Lubowski, and Stavins, 2002):

$$V_i = \int \pi_i^*(t) e^{-rt} dt \quad (3.1)$$

where V_i is farmland value, π_i^* is the optimal net returns to land at location i , and r is the discount rate. Net returns are a function of the product of the land's output market price P_k and a vector of input prices w , and they depend also on a set of exogenous variables and different land uses denoted as k (Hsiang, 2016);

$\mathbf{z}_i = (\mathbf{C}_i, \mathbf{G}_i, \mathbf{S}_i)$ that consists of climate attributes $\mathbf{C}_i$, socio-economic characteristics $\mathbf{G}_i$, and soil quality indicators $\mathbf{S}_i$:


$$\pi_i^* = \max_{k \in \{1, \dots, K\}} \pi_{ik}(P_k, \mathbf{w}, \mathbf{z}_i) \quad (3.2)$$

In this framework, the effect of climate change on farmland values can be conceptualized in panel (a) of 3.1. Note that the x-axis focuses on changes in temperature, but a similar approach could be used for changes in precipitation. The three overlapping parabolas represent the envelope formed by three different types of land use reflecting that as temperature increases, farmers would choose to switch their inputs and outputs to maximize the land's net return (y-axis left) for a given temperature level. The land value this process leads to (y-axis right) is reflected in the outer parabola for the entire range of temperature. Traditionally, 3.1 is parametrized through the econometric model depicted in 3.3 where $(\beta'_1, \beta'_2, \mathbf{G}, \boldsymbol{\theta})$ is a set of exogenous coefficients (Fezzi and Bateman, 2015).⁴


$$V_i = \alpha_0 + \beta'_1 \mathbf{C}_i + \beta'_2 \mathbf{C}_i^2 + \gamma' \mathbf{G}_i + \boldsymbol{\theta}' \mathbf{S}_i + \epsilon_i \quad (3.3)$$

The average marginal effect of climate on farmland value, $\frac{\partial V_i}{\partial \mathbf{C}_i} = \beta'_1 + 2\beta'_2 \mathbf{C}_i$, is a function of $\mathbf{C}_i$. However, in this expression, both linear terms are a-spatial. Their magnitude is assumed to be valid for any county in the sample, from Florida to Minnesota, so that the only source of spatially varying marginal effects lies in the value of the covariate $\mathbf{C}_i$. Several authors have challenged the assumption that the marginal effect of climate change is constant for all locations (Mendelsohn and Dinar, 2003; Schlenker, Hanemann, and Fisher, 2005; Timmins, 2006). Panel (b) of 3.1 illustrates their rationale for the simplest case of two different responses. If the econometric model omits the different responses to climate change, future farmland predictions will fall between the green and the black outer parabolas.

⁴Based on their approach, Mendelsohn, Nordhaus, and Shaw (1994) forecast that global warming may bring average gains to U.S. agriculture. Since then, their pioneering work has led to a surge of studies using a variation of the Ricardian approach, either in the U.S. or abroad (Mendelsohn and Massetti, 2017).



(a) Spatial homogeneous responses



(b) Spatial heterogeneous responses

Figure 3.1: Hedonic approach under different spatial assumptions.

Note: For each color, the three overlapping paraboles represent three different land uses, so the upper envelop maximizes net returns with respect to the different activities and temperature levels. The outer parabole represents the land value over the different temperature and illustrates the farmers' net return expectations.

The slope heterogeneity illustrated in panel (b) of 3.1 can be treated by disaggregating the estimations to a finer spatial scale. The Ricardian literature offers different grouping strategies that consist of clustering counties that are thought to be subject to similar effects from climate change. Some of the earliest investigations in this direction find that climate change has little or no effect on irrigated farmlands, suggesting that in earlier studies the gains associated with climate change are underestimated (Mendelsohn and Dinar, 2003). Once irrigation is accounted for, Schlenker, Hanemann, and Fisher (2005) find that the projected climate impact estimates for rain-fed areas converge to a national annual loss of \$9.62 billion. In their panel data estimation, Deschenes and Greenstone (2007) allow the climate coefficients to fully interact with irrigation and state fixed effects. Their results indicate that climate change will benefit farmers by \$3.24 billion.⁵ More recently, Dall’erba and Domínguez (2016) split the U.S. Southwestern counties into two groups based on the median value of elevation. Their results show that the probability of a negative future impact of climate change is larger in highland counties which are forecasted to experience heat waves more frequently than in the lowland counties.

When Cai and Dall’Erba (2021) investigate the effect of the above grouping specifications on the marginal impact of climate change on farmlands, they conclude that the predicted change in farmland values varies only slightly when either irrigation or elevation is used as the criterion for grouping spatial units. Contradictory findings about the direction and/or magnitude of the impact of climate change in U.S. agriculture arising from different grouping strategies of spatial units suggest that the clustering of spatial units in such studies has an effect on the analytical results – a problem often referred to as the modifiable areal unit problem (MAUP) (Fotheringham and Wong, 1991). This effect has been recognized in the context of studies of climate-induced agricultural productivity by (Deschenes and Greenstone, 2007).

Given the unsuitability of a priori groupings of spatial units with which to demonstrate spatially heterogeneous effects of climate change on agricultural production,

⁵Fisher et al. (2012) find and correct several irregularities in the paper of Deschenes and Greenstone (2007), hence concluding with a negative impact of climate change on U.S. agriculture across different model specifications.

we take a different approach and rely on local regressions to relax the assumption that the data generating process is constant over space without having to specify any a priori regionalization of units. Several types of statistical models exist that generate locally varying responses to the same stimuli, such as the Bayesian spatially varying coefficients model (Gelfand et al., 2003; Fotheringham, Yang, and Kang, 2017; Fotheringham, Li, and Wolf, 2021) and the frequentist multiscale geographically weighted regression approach (Yu, Fotheringham, Li, Oshan, Kang, et al., 2020; Yu, Fotheringham, Li, Oshan, and Wolf, 2020), although comparative studies suggest they produce similar results (Wolf, Oshan, and Fotheringham, 2018).

We selected the GWR framework because of its ease-of-use, familiarity and scalability. 3.4 depicts a spatially varying model where y_i is the dependent variable, and x_{ik} is the k^{th} explanatory variable. This model is known as the basic-GWR. The local coefficients are denoted as $\beta_0(u_i, v_i)$ for the intercept and $\beta_k(u_i, v_i)$ for the coefficients associated with the k^{th} explanatory variable. In turn, coefficients are a function of the i^{th} observation's latitude u_i and longitude v_i . The stochastic error is denoted by ϵ_i . To retrieve the local coefficients associated with 3.4, GWR employs a variation of the Generalized Least Squares: $\beta = [\mathbf{X}'\mathbf{W}(u_i, v_i)\mathbf{X}]^{-1}\mathbf{X}'\mathbf{W}(u_i, v_i)\mathbf{Y}$, where the diagonal weighted matrix is denoted by $\mathbf{W}(u_i, v_i)$. Because its elements depend on the location of the i^{th} observation, each parameter is location-specific.

$$y_i = \beta_0(u_i, v_i) + \sum_k \beta_k(u_i, v_i)x_{ik} + \epsilon_i \quad (3.4)$$

One problem with basic-GWR is that it imposes the same bandwidth (i.e. the same degree of spatial heterogeneity) on each set of local parameter estimates. To remove this problem, Fotheringham, Yang, and Kang (2017) derive a generalization of the basic-GWR, known as Multiscale Geographically Weighted Regression (MGWR), that permits multiple local bandwidths. This model is depicted in 3.5 where the superscript bw indicates that each set of local parameter may have a different bandwidth, distinguishing it from 3.4.

$$y_i = \beta_0^{bw}(u_i, v_i) + \sum_k \beta_k^{bw}(u_i, v_i)x_{ik} + \epsilon_i \quad (3.5)$$

MGWR is estimated with a back-fitting algorithm that maximizes the expected log-likelihood of each term in an additive model: $\mathbf{y} = \sum_{\mathbf{k}} \mathbf{f}_{\mathbf{k}}$, where each $\mathbf{f}_{\mathbf{k}} = \beta_k^{bw}(u_i, v_i)x_{ik}$ in 3.5 and $\mathbf{y}$ is its dependent variable. In each iteration, the back-fitting algorithm calculates the residual $\epsilon = \mathbf{y} - \sum_{\mathbf{k}} \mathbf{f}_{\mathbf{k}}$ and regresses each term of the form $\epsilon + \mathbf{f}_{\mathbf{k}}$ against its corresponding data vector x_{ik} using the basic-GWR estimator described above until a pre-specified convergence criterion is reached.⁶ Estimating the Ricardian model of 3.3 with the capabilities of the MGWR estimator allows the estimation of the localized marginal effects of the climate conditions on agricultural production without pre-specifying any grouping of spatial units and without imposing any restriction on optimized covariate-specific bandwidths. In this framework, the global model in 3.3 relating agricultural land value to a series of covariates is depicted in 3.6.

$$V_i = \alpha_0^{bw}(u_i, v_i) + \beta_1^{bw}(u_i, v_i)'C + \beta_2^{bw}(u_i, v_i)'C^2 + \gamma^{bw}(u_i, v_i)'G + \theta^{bw}(u_i, v_i)'S + \epsilon_i \quad (3.6)$$

To illustrate how local models improve upon 3.3, panel (a) in 3.2 shows the extreme case in which the coefficient β varies across all spatial units. Panels (b) – (d) employ different grouping techniques such that β is allowed to vary across clusters, and panel (e) is the basic-GWR estimator. As more interactions are included global specifications improve but they still present severe biases. The set of β values recovered with the GWR specification captures closely the true specification of β without imposing any dummy interaction. Panels (b) – (d) reflect the bias that ad-hoc specifications would have on local estimates while panel (e) illustrates our proposed correction. The key feature of GWR methods is therefore to allow the researcher to uncover the presence of local clusters and to include them in the estimation process without prior knowledge of the number of existing groups, group location and group membership.

⁶MGWR inference is discussed in the appendix A.




Figure 3.2: Spatially varying data generating process across different estimators. Note: True $\beta = 1 + (u_i + v_i)^2$. Panels (a) is the simulated spatial data set. Panels (b) - (d) employ different grouping techniques such that β is allowed to vary across clusters, and panel (e) is the basic-GWR estimator. As more interactions are included, global specifications improve, but they still present severe biases. The set of β values recovered with the GWR specification captures closely the true specification of β without imposing any dummy interaction.

3.3 Data

Following closely the canonical work of Mendelsohn, Nordhaus, and Shaw (1994), our model relates average farmland value per acre with seasonal climate data while controlling for demographic and soil quality characteristics as shown in 3.6. We initially considered calibrating our model with data from all 3,008 non-metropolitan contiguous U.S. counties, but we followed Schlenker, Hanemann, and Fisher (2005) and Dall’erba and Domínguez (2016) by removing urban counties, 207 of them, defined as counties with more than 1,600 inhabitants per square mile. A further 34 counties were removed because of missing data for at least one of our variables. Our final sample is therefore composed of 2,687 U.S. counties.

3.1 contains basic summary statistics related to the dependent and the main explanatory variables employed in our study. Our variable of interest is the average county farmland value per acre. The U.S. Department of Agriculture prepares a census every five years in which it collects farmers’ land prices. The two most recent censuses were in 2012 and 2017, so we will use the former for calibration purposes and the later to test the predictive accuracy of each model compared to observed values. Our historical climate variable are precipitation and temperature. Precipitation is measured in cubic inches and temperature is measured in degree Celsius. The climate data are collected from the Parameter-elevation Regression on Independent Slopes Model (PRISM) database (PRISM Climate Group, 2021). We process PRISM data to estimate a historical average for precipitation and temperature for the U.S. counties over the 1993-2012 and 1993-2017 periods. The Ricardian literature raises concerns about the non-concavity in the farmland value’s responses to seasonal climate. Indeed, one should expect farmland value to increase with respect to climate at a decreasing rate as in 3.1. This limitation prevents researchers from interpreting extremum as maximum and is the results of aggregating monthly climate into seasonal climate (Darwin, 1999). Disaggregating climate into monthly values, however, would cause serious multicollinearity issues, and dropping some arbitrary variable may cause an omitted variable bias (Mendelsohn and Massetti, 2017).

Table 3.1: Descriptive statistics: Dependent variable and historical climate variables

Variable	Mean	s.d.	Max.	Min.
2012				
Farmland value per acre	\$3,216	\$2,029	\$ 21,801	\$ 192
<i>Temperature</i>				
Winter (Dec. - Feb.)	1.94	5.94	19.56	-11.99
Spring (Mar. - May.)	12.50	4.93	24.42	0.87
Summer (Jun. - Aug.)	23.50	3.14	32.61	13.45
Autumn (Sep. - Nov.)	13.33	4.31	25.15	2.56
<i>Precipitation</i>				
Winter (Dec. - Feb.)	2.79	1.87	17.29	0.06
Spring (Mar. - May.)	3.34	1.27	10.55	0.04
Summer (Jun. - Aug.)	3.649	1.44	9.83	0.01
Autumn (Sep. - Nov.)	2.86	1.28	10.77	0.08
2017				
Farmland value per acre	\$3,165	\$2,095	\$ 28,569	\$ 162
<i>Temperature</i>				
Winter (Dec. - Feb.)	2.23	6.03	19.65	-12.03
Spring (Mar. - May.)	12.88	4.90	24.74	1.21
Summer (Jun. - Aug.)	23.44	3.13	32.77	13.53
Autumn (Sep. - Nov.)	13.11	4.33	24.99	2.28
<i>Precipitation</i>				
Winter (Dec. - Feb.)	2.79	1.90	17.65	0.03
Spring (Mar. - May.)	3.47	1.29	10.96	0.00
Summer (Jun. - Aug.)	3.61	1.56	9.82	0.00
Autumn (Sep. - Nov.)	2.76	1.22	12.15	0.02
Observations				2,689

Note: Temperature is reported in degrees Celsius. Precipitation is reported in cubic millimeters per day.

3.1 contains basic summary statistics related to our control variables. The county’s population density per square mile is obtained from the 2010 U.S. Census Bureau of Statistics, and the county’s per-capita income is obtained from the U.S. Bureau of Economic Analysis of the same year. Controlling for socio-economic characteristics allows us to remove unwanted variation that can bias our forecast. Soil quality is a measure of a land’s productivity, so we include information about salinity, flood frequency ratio, wetland, slope steepness, erosion index, sand contents, clay contents, permeability, and moisture capacity. Soil indicators are obtained from the U.S. Geological Survey. Because we do not have reliable data for our control variables in 2017 and those for 2012 are proxied with data from 2010, we follow Deschenes and Greenstone (2007); Deschênes and Greenstone (2012) and assume they are constant. Thus, we can forecast changes in the value of farmland prices in function of climate change (first-degree difference in our climate variables).

In the subsequent sections, we test the robustness of our model relative to previous specifications forecasting the impacts of climate change on agricultural land values, namely: Mendelsohn, Nordhaus, and Shaw (1994) (no grouping of spatial units); Schlenker, Hanemann, and Fisher (2005) (two-way division of spatial units based on irrigated ratio ≥ 0.2); and Dall’erba and Doménguez (2016) (two-way division based on median elevation). For each approach we employ the same 2012 dataset to predict the 2017 farmland values. This allows us to assess the predictive power of the local regression approach described here relative to other approaches for a year for which the data are observed. We then employ our updated data in 2017 to estimate the marginal effect of yearly climate variables on farmland values to further study the capabilities of our estimator choice.

Table 3.2: Descriptive statistics for the value of land and control variables

Variable	Mean	s.d.	Max.	Min.
Socioeconomic Variables				
Population density per square mile in 2010	73.22	99.79	1582.74	0.13
Per capita income in 2010	\$39,209	\$10,531	\$52,380	\$18,404
Soil Quality Variables				
Salinity	0.20	0.56	7.19	0.00
Flood frequency ratio	0.08	0.12	0.73	0.00
Wetland ratio	0.55	0.38	1.00	0.00
Slope steepness	11.46	11.75	68.41	0.54
Erosion index	0.27	0.08	0.49	0.00
Sand contents	38.37	19.68	96.33	0.98
Clay content	18.54	8.15	55.12	1.12
Permeability	20.46	18.68	113.63	0.91
Moisture capacity	0.16	0.04	0.38	0.01
Observations				2,687

3.4 Results

3.4.1 Cross-Validation Results

Cai and Dall’Erba (2021) recognize the shortcomings associated with the wrongful grouping of U.S. counties in the Ricardian approach, so the authors perform cross-validation exercises across several groupings for the same year. Our manuscript expands their approach and calibrates all the models on 2012 data to predict the changes in farmland values by 2017. To our knowledge, no previous Ricardian study has used a cross-validation approach across years to demonstrate the better predictive power of one model over another. The results of our cross-validation exercise are

depicted in 3.3 where the observed 2012-2017 changes are displayed on a map and in a histogram and the 2012-2017 predicted changes based on the various models are also displayed.

The table at the bottom of the figure reports certain prediction diagnostics. All the global regression models significantly underestimate the changes in farmland values. Many of the predictions from the global models are centered around zero, but with some high-valued cluster of losses, which is reflected in the average estimated county loss of -\$412.08 to -\$280.79 per acre. Only MGWR predicts gains between 2012-2017 with an average of \$109.86 per acre, which is still lower than the observed values in 2017 at \$573.92 per acre. Despite this difference, our local estimator significantly outperforms any of the global specifications as they each predict losses by 2017. The better predictive performance of MGWR is also evidenced by its low Mean Squared Error (MSE)⁷ value compared to other models.⁸ Further examination of 3.3 indicates that the results from the local model reproduces a pattern similar to the isotherms formed when temperature is mapped across latitudes in the U.S. In contrast, none of the global specifications shows this pattern, but instead they highlight ill-defined clusters of high losses.⁹

Two additional elements can be drawn from these results. First, the local capabilities of the non-stationary Ricardian approach do not impose any kurtosis in the distribution of gains and losses. 3.3 shows that the tails from the global predictions are thinner and longer than those from the local predictions. As a result, most global predictions fall around the mean, hence creating high peaks in their distribution. Because MGWR Ricardian approach does not restrict the kurtosis of our predictions, it is able to uncover important unspecified heterogeneity. The second feature we note is that local specifications ignore outliers, such as the land appreciation in the

⁷ $MSE = \frac{1}{n} \sum_N (V_i - \hat{V}_i)^2$, where $\hat{V}_i$ is the predicted value.

⁸Our approach also ignores some outliers along the West Coast which every global estimator predict would increase slightly in 2017, suggesting that our estimator exchanges complexity in predictions for generality in the model.

⁹The MGWR predictions to the middle of the century do not follow the same temperature/latitude approach, demonstrating that our MGWR approach incorporates unspecified clusters beyond weather.




Figure 3.3: Spatial and density distribution of observed and predicted values for 2017 across specification.

Note: MGWR is our local specification. MNS employs no grouping strategy following Mendelsohn, Nordhaus, and Shaw (1994). SHF splits counties by irrigation ratio (0.2) following Schlenker, Hanemann, and Fisher (2005). DD splits the counties by median elevation following Dall’erba and Domínguez (2016).

Californian counties which is likely to be a result of the California Drought of 2012. This suggests that the local model underestimates the dependent variable in the locations that are impacted by extreme weather events but still does a much better job in capturing the impacts of these events than any global model.

We investigate further the differences between our MGWR estimator and the global estimators in 3.4 by dividing forecasts by climate regions (Karl and Koss, 1984). This figure shows that our local estimator outperforms all global estimators across the different climate regions except for the Southern and Southeastern regions. The underperformance of global estimators is clearly observed in the Southwest and in the Northwest and Central regions where all global estimators predict significantly lower values than the observed 2017 farmland values. MGWR predictions in the Northwest and Central regions surpassed the observed 2017 farmland values. The absolute value of the difference between our local and the global regressions favors our MGWR application. For the Southern and Southeastern regions, the median predictions from the global estimators are close to the median values of those regions with values above the average predictions, suggesting that the local estimator outperforms the global estimators.

We acknowledge that there are two caveats about our MGWR estimation that also affect the studies this manuscript examines. The functional form is assumed to be constant through time. While this is a necessary assumption to make predictions, adaptation strategies to mitigate the impact of climate change can take place, affecting the model's parameters. The second caveat is the role of omitted variable bias in our estimations. Studies of farmland prices are constrained by cross-sectional data, so practitioners are unable to control for time-invariant influences. Despite these limitations, the nature of our study is not to develop the *best* functional form, but to propose an estimator to treat the spatial heterogeneity in the study of farmland prices. As this section demonstrates, our MGWR Ricardian approach incorporates the spatial heterogeneity present in our dataset. This having been said, the next section is concerned with the estimation of county-specific marginal effects of climate that can be employed to develop place-tailored mitigation strategies to climate change.




Figure 3.4: Observed and forecasted farmland values in 2017.


Note: Figure investigates the regional differences across estimators by dividing forecast by climate regions (Karl and Koss 1984). Outliers falling above \$5,000 and below -\$5,000 are trimmed for exposition.

3.4.2 Place-tailored Weather Impacts on Farmland Prices


3.5 reports the MGWR parameter estimates of yearly temperature and precipitation obtained based on 2017 data and 3.6: $\sum_{k \in \Phi} \frac{\partial V_i}{\partial C_i} = \sum_{k \in \Phi} \{\beta_{1k}^{bw}(u_i, v_i)' + 2\beta_{2k}^{bw}(u_i, v_i)'C_{ik}\}$, where $k \in \Phi$ corresponds to the seasonal weather variables (either temperature or precipitation). Panel (a) indicates the marginal effect of yearly precipitation on farmland values. Not surprisingly, most counties benefit from an increase in precipitation. In contrast, most counties will experience losses from rising temperature as shown in panel (b). Our findings indicate significant heterogeneity in the impact of climate change in the U.S. In both maps, the overall trend with respect to climate change seems to be either positive (for precipitation) or negative (for temperature), but the direction of the impact changes for a small subset of counties. For example, the impact of more precipitation is positive for most of the states except in the counties close to the coast in Oregon and Washington. Similarly, more temperature damages most farmers in the U.S. except for some counties in the upper part of Midwest (South Minnesota and North Iowa). An explanation for this phenomenon could be that these counties have very low values of their estimated yearly effects that are

statistically equal to zero. While all our estimated seasonal effects are statistically different from zero, we are unable to test if the yearly summation are statistically different from zero i.e., no GWR method exists that allows to test linear combination of parameters.

In 3.6 and 3.7, we disaggregate yearly impacts into seasonal impacts: $\frac{\partial V_i}{\partial \mathbf{C}_i} = \beta_{1k}^{bw}(u_i, v_i)' + 2\beta_{2k}^{bw}(u_i, v_i)' \mathbf{C}_{ik}$. In reference to the previous heterogeneity observed in 3.5, 3.6 suggests that precipitation in Spring severely affects counties in the coastal areas of Washington and Oregon, forming an isolated cluster in the affected area. No seasonal temperature, however, seems to dominate the direction of the sign for the exceptional cluster of counties perceived in 3.7. Additionally, seasonal impacts observed in 3.6 and 3.7 have different levels of variance. A strip of counties between Northern California and Washington benefit the most from Winter precipitation, highlighting their status of irrigated counties. Precipitation in the Spring, however, damages these counties. For rain-fed counties East of the 100th more rain is associated with lower values of precipitation in the Spring affecting specially counties in the South and Midwest. All these marginal effects have been calculated and are available upon request.



(a) Marginal Effect: Yearly Precipitation



(b) Marginal Effect: Yearly Temperature

Figure 3.5: Marginal effect of yearly precipitation and temperature.
 Note: Figure shows the yearly marginal effect of precipitation and temperature, obtained with equation 3: $\sum_{k \in \Phi} \frac{\delta V_i}{\delta C_i} = \sum_{k \in \Phi} \beta_{1,k}^{bw}(u_i, v_i)' + 2\beta_{2,k}^{bw}(u_i, v_i)'C, k$, where $k \in \Phi$ indicates the variables associated with the weather variable i.e., either temperature or precipitation.




Figure 3.6: Seasonal precipitation marginal effect .

The Ricardian application of Dall’erba and Domínguez (2016) uncovers that the marginal impact of climate change is conditional on altitude. The results in 3.7 show a cluster of high-impacted counties in the Appalachian mountains, providing evidence that our MGWR Ricardian approach supports the conclusions derived in Dall’erba and Domínguez (2016). The cluster does not appear for precipitation, suggesting that the impact of precipitation is not altitude-dependent. In two separate studies, Schlenker, Hanemann, and Fisher (2005) and Schlenker and Roberts (2006) claim that the 100th meridian should serve as an appropriate split to cluster counties. The rationale is that counties East of the 100th meridian are rainfed and so more susceptible to precipitation variations. Panel (a) in 3.6 and panel (a) in 3.5

provide evidence to support a split by the 100th meridian in terms of the impact of precipitation on farmland values. However, the marginal impact of temperature on farmland value does not display the same pattern. Indeed, panel (b) indicates the presence of substantial spatial heterogeneity in the marginal impact of temperature on both sides of the U.S. Irrigated counties West of the 100th meridian show more heterogeneity. Schlenker, Hanemann, and Fisher (2007) consider this phenomenon in their own Ricardian study of agriculture. Their study suggests that water availability is capitalized into land prices and concludes that farmers benefit from rain after the growing season (Fall and Winter). Our results support their conclusions. Panel (a) in 3.5 shows that most counties West of the 100th meridian benefit from precipitation, which replenishes aquifers. On the other hand, rain-fed counties east of the 100th meridian show more heterogeneity even with some losses, suggesting that too much precipitation may lead to floods. As we previously indicated, researchers often split counties a priori by the 100th meridian and include a dummy variable to generate two sets of slope coefficients. The results in 3.5 support their rationale without needing to undertake such an arbitrary separation.




Figure 3.7: Seasonal temperature marginal effect.

3.5 Conclusion

Economic estimates of the agricultural damages from climate change have largely relied on the Ricardian approach and on econometric techniques that either disregard or reduce spatial heterogeneity to a few groups. Given the size of the U.S. territory and the large variety of local climate conditions it offers, we believe that a much greater degree of spatial heterogeneity is present in the impacts climate change will have on US agriculture. This paper explores this question by relying on the recently developed MGWR estimator. Compared to global frameworks, this modeling ap-

proach allows us to uncover spatial heterogeneity in the effects of climate change on agricultural production without having to set a priori the number of clusters, their location and their membership; hence allowing us to generate place-specific marginal effects relying on a purely data-driven approach. Compared with global models, the MGWR forecasts exhibit a better capacity for predicting the large spatial variations in climate impact on farmland values and for accurately predicting future farmland values, as exemplified in our cross-validation exercise for the year 2017. By generating place-specific estimates, our approach offers policy makers and stakeholders a tool to address the effects of climate change on a case-by-case basis in each county and suggest mitigation and adaptation strategies that are place-tailored. For instance, our results suggest that precipitation increases have harmful effects for the Western part of the U.S., suggesting that droughts are a concern of farmers. On the other hand, our approach also captures the effect of water availability in the the value of lands for the Western counties.

Despite the novelty and strengths of our approach, we acknowledge that our approach does not include any form of interregional dependence that may place a large role in climate change mitigation. For instance, a recent contribution by Dall’Erba, Chen, and Nava (2021b) estimates that domestic trade is able to mitigate the negative impact of climate change on crop profit by a magnitude of \$14.5 billion. Fan, Fisher-Vanden, and Klaiber (2018) arrive to similar conclusions arguing that not only trade but also migration are mitigation tools against climate change. As such, further research shall not only offer county-level predictions, but it shall also account for domestic trade which has been largely ignored in the literature. Finally, this study also ignores the role of federal assistance and of state-level efforts to mitigate climate change impacts within their jurisdiction but not limited to farm subsidies (Deschenes and Greenstone, 2007; Dall’erba and Domínguez, 2016).

3.6 Chapter 3 References

- Cai, Chang and Sandy Dall’Erba (2021). “The impact of global warming on agriculture: A Ricardian analysis: Comment”. In: *Climate Change* 167(14).
- Dall’Erba, Sandy, Zhangliang Chen, and Noé J Nava (2021b). “US Interstate Trade Will Mitigate the Negative Impact of Climate Change on Crop Profit”. In: *American Journal of Agricultural Economics*.
- Dall’erba, Sandy and Francina Domiéguez (2016). “The impact of climate change on agriculture in the Southwestern United States: The Ricardian approach revisited”. In: *Spatial Economic Analysis* 11(1), pp. 46–66.
- Darwin, Roy (1999). “The impact of global warming on agriculture: A Ricardian analysis: Comment”. In: *American Economic Review* 89(4), pp. 1049–1052.
- Deschênes, Olivier and Michael Greenstone (2012). “The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Reply”. In: *American Economic Review* 102(7), pp. 3761–73.
- Fan, Qin, Karen Fisher-Vanden, and H Allen Klaiber (2018). “Climate change, migration, and regional economic impacts in the United States”. In: *Journal of the Association of Environmental and Resource Economists* 5(3), pp. 643–671.
- Fezzi, Carlo and Ian Bateman (2015). “The impact of climate change on agriculture: Nonlinear effects and aggregation bias in Ricardian models of farmland values”. In: *Journal of the Association of Environmental and Resource Economists* 2(1), pp. 57–92.
- Fisher, Anthony C, W Michael Hanemann, Michael J Roberts, and Wolfram Schlenker (2012). “The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment”. In: *American Economic Review* 102(7), pp. 3749–60.

- Fotheringham, A Stewart, Ziqi Li, and Levi John Wolf (2021). “Scale, context, and heterogeneity: A Spatial Analytical Perspective on the 2016 US Presidential Election”. In: *Annals of the American Association of Geographers*, pp. 1–20.
- Fotheringham, A Stewart and David WS Wong (1991). “The modifiable areal unit problem in multivariate statistical analysis”. In: *Environment and Planning A* 23(7), pp. 1025–1044.
- Fotheringham, A Stewart, Wenbai Yang, and Wei Kang (2017). “Multiscale geographically weighted regression (MGWR)”. In: *Annals of the American Association of Geographers* 107(6), pp. 1247–1265.
- Gbetibouo, Glwadys Aymone and Rashid M Hassan (2005). “Measuring the economic impact of climate change on major South African field crops: A Ricardian approach”. In: *Global and Planetary Change* 47(2-4), pp. 143–152.
- Gelfand, Alan E, Hyon-Jung Kim, CF Sirmans, and Sudipto Banerjee (2003). “Spatial modeling with spatially varying coefficient processes”. In: *Journal of the American Statistical Association* 98(462), pp. 387–396.
- Hsiang, Solomon (2016). “Climate Econometrics”. In: *Annual Review of Resource Economics* 8, pp. 43–75.
- Karl, Thomas and Walter James Koss (1984). “Regional and national monthly, seasonal, and annual temperature weighted by area, 1895-1983”. In.
- Liu, Hui, Xiubin Li, Guenther Fischer, and Laixiang Sun (2004). “Study on the impacts of climate change on China’s agriculture”. In: *Climatic Change* 65(1), pp. 125–148.
- Mendelsohn, Robert and Ariel Dinar (2003). “Climate, water, and agriculture”. In: *Land Economics* 79(3), pp. 328–341.
- Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw (1994). “The impact of global warming on agriculture: A Ricardian analysis”. In: *American Economic Review*, pp. 753–771.

- Mendelsohn, Robert O and Emanuele Massetti (2017). “The use of cross-sectional analysis to measure climate impacts on agriculture: Theory and evidence”. In: *Review of Environmental Economics and Policy* 11(2), pp. 280–298.
- Páez, Antonio, Fei Long, and Steven Farber (2008). “Moving window approaches for hedonic price estimation: An empirical comparison of modelling techniques”. In: *Urban Studies* 45(8), pp. 1565–1581.
- Plantinga, Andrew J, Ruben N Lubowski, and Robert N Stavins (2002). “The effects of potential land development on agricultural land prices”. In: *Journal of Urban Economics* 52(3), pp. 561–581.
- PRISM Climate Group, Oregon State University (2021). *PRISM Gridded Climate Data*. data retrieved from PRISM website, <http://prism.oregonstate.edu>.
- Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher (2005). “Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach”. In: *American Economic Review* 95(1), pp. 395–406.
- Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher (2007). “Water availability, degree days, and the potential impact of climate change on irrigated agriculture in California”. In: *Climatic Change* 81(1), pp. 19–38.
- Schlenker, Wolfram and Michael J Roberts (2006). “Nonlinear effects of weather on corn yields”. In: *Review of Agricultural Economics* 28(3), pp. 391–398.
- Shiu, Yi-Shiang and Yung-Chung Chuang (2019). “Yield estimation of paddy rice based on satellite imagery: Comparison of global and local regression models”. In: *Remote Sensing* 11(2), p. 111.
- Soler, Ismael P and German Gemar (2018). “Hedonic price models with geographically weighted regression: An application to hospitality”. In: *Journal of Destination Marketing & Management* 9, pp. 126–137.
- Timmins, Christopher (2006). “Endogenous land use and the Ricardian valuation of climate change”. In: *Environmental and Resource Economics* 33(1), pp. 119–142.

- Wang, Jinxia, Robert Mendelsohn, Ariel Dinar, Jikun Huang, Scott Rozelle, and Lijuan Zhang (2009). “The impact of climate change on China’s agriculture”. In: *Agricultural Economics* 40(3), pp. 323–337.
- Wolf, Levi John, Taylor M Oshan, and A Stewart Fotheringham (2018). “Single and multiscale models of process spatial heterogeneity”. In: *Geographical Analysis* 50(3), pp. 223–246.
- Yu, Hanchen, A Stewart Fotheringham, Ziqi Li, Taylor Oshan, Wei Kang, and Levi John Wolf (2020). “Inference in multiscale geographically weighted regression”. In: *Geographical Analysis* 52(1), pp. 87–106.
- Yu, Hanchen, A Stewart Fotheringham, Ziqi Li, Taylor Oshan, and Levi John Wolf (2020). “On the measurement of bias in geographically weighted regression models”. In: *Spatial Statistics* 38, p. 100453.
- Zhang, Honglei, Jie Zhang, Shaojing Lu, Shaowen Cheng, and Jinhe Zhang (2011). “Modeling hotel room price with geographically weighted regression”. In: *International Journal of Hospitality Management* 30(4), pp. 1036–1043.

APPENDIX A

SUPPLEMENTARY MATERIAL FOR CHAPTER 1

A.1 Households exhibit rank 3 Engel curve behavior




Figure A.1: Engel curves of beverage budget shares on log beverage expenditure.
Note: Graphs fit non-parametric regressions of the log of beverage expenditure (horizontal axes) on each budget share (vertical axes).

A.2 Item aggregation from ENIGH 2018

Table A.1: Summary of statistics for each item in the composite goods

ENIGH code	Item description	Mean(quantity)	Mean(expenditure)
SSB			
A220	Soda drinks	1.90	23.74
Juice			
A217	Water with fruits	1.35	20.10
A218	Processed juice	1.19	18.93
A221	Energy drinks	1.04	29.53
Milk			
A075	Pasteurized milk	1.64	24.60
A076	Condensed milk	0.77	23.81
A077	Evaporated milk	0.60	20.71
A080	Raw milk	2.10	19.55
A081	Non-cow milk	1.77	51.09
Other			
A214	Bottle water	22.35	23.49
A216	Mineral water	1.56	18.85

Note: ENIGH codes for all items in our aggregations can be used to trace back composite goods to the ENIGH documentation to facilitate replication. Item description is translated from ENIGH documentation into English. Quantity and expenditure are reported by surveyed households to INEGI. Averages are per time the item is purchased. Monetary values are measured in 2018 Mexican pesos.

A.3 Marshallian elasticities across methodologies

Table A.2: Marshallian elasticities

With respect to		Price				Zero purchases
		SSB	Juice	Milk	Other	
Quantity SSB	ND	-0.951	-0.005	-0.005	0.002	29%
	Uncensored	-0.652	0.009	-0.406	0.162	
	Endogenous	-1.320	0.021	0.105	1.507	
Juice	ND	-0.039	-0.948	0.112	-0.049	76%
	Uncensored	-0.005	-0.442	-0.102	0.031	
	Endogenous	-0.033	-0.903	-0.027	0.196	
Milk	ND	0.022	0.075	-1.031	-0.029	34%
	Uncensored	-0.230	-0.400	-0.481	0.189	
	Endogenous	-0.009	-0.101	-1.015	1.683	
Other	ND	-0.334	-0.228	-0.276	-0.639	65%
	Uncensored	-0.013	-0.013	0.014	-0.734	
	Endogenous	0.272	0.299	0.101	1.242	

Note: Table shows Marshallian elasticities for our basket of four goods across three methodological specifications. ND employs our correction. Uncensored does not employ our correction, so it does not correct for censorship. Endogenous does not employ our correction and does not correct for unit-value endogeneity i.e., section 3.1.

A.4 Marshallian elasticities across years

Table A.3: Marshallian and budget elasticities

With respect to		Price				Budget
		SSB	Juice	Milk	Other	
Quantity						
SSB	2018	-0.951	-0.055	-0.005	0.002	1.009
	2016	-0.988	-0.053	0.082	-0.007	0.996
	2014	-0.833	-0.025	0.042	-0.032	0.847
Juice	2018	-0.039	-0.948	0.112	-0.049	0.923
	2016	-0.075	-0.949	0.108	-0.040	0.956
	2014	-0.072	-0.891	0.127	-0.027	0.863
Milk	2018	0.022	0.075	-1.031	-0.029	0.963
	2016	0.091	0.069	-1.142	-0.022	1.004
	2014	-0.143	-0.036	-1.088	0.020	1.247
Other	2018	-0.334	-0.228	-0.276	-0.639	1.478
	2016	-0.273	-0.188	-0.252	-0.678	1.391
	2014	-0.003	0.023	-0.209	-0.771	0.954

Note: Table shows Marshallian and budget elasticities for our basket of four goods across several years.

APPENDIX B

SUPPLEMENTARY MATERIAL FOR CHAPTER 3

B.1 Inference in the MGWR

Contrary to the basic-GWR, MGWR is defined as a Generalized Additive Model (GAM) whose values are estimated by a back-fitting algorithm (Fotheringham, Yang, and Kang, 2017). In this case inference is a limitation since there is no single hat matrix mapping the estimated values $\hat{y}_i$ onto y_i . Yu, Fotheringham, Li, Oshan, Kang, et al. (2020) reframe the basic-GWR as a GAM, and demonstrate that the hat matrix for a MGWR model can be derived within the back-fitting estimation algorithm. We display their derivation, by first showing that the hat matrix for a basic-GWR can be written as an additive term. Then, we demonstrate how a hat matrix can be estimated and used for the computation of local standard errors.

The hat matrix for the basic-GWR $\mathbf{S}$, is defined in B.1 with each row given by $\mathbf{s}_i = \mathbf{X}_i(\mathbf{X}'\mathbf{W}(u_i, v_i)\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}(u_i, v_i)$. Here, $\mathbf{X}_i$ is the i^{th} row in the matrix of all predictors $\mathbf{X}$.

$$\mathbf{S} = \begin{pmatrix} \mathbf{X}_1(\mathbf{X}'\mathbf{W}(u_1, v_1)\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}(u_1, v_1) \\ \dots \\ \mathbf{X}_n(\mathbf{X}'\mathbf{W}(u_n, v_n)\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}(u_n, v_n) \end{pmatrix} \quad (\text{B.1})$$

Let us define the additive hat matrices by $\mathbf{R}_k$ for the set of local parameters β_k associated with the k^{th} explanatory variable, such that it has the following two properties regarding the dependent variable $\mathbf{y}$:

$$\hat{\mathbf{f}}_k = \hat{\mathbf{R}}_k \mathbf{y} \quad (\text{B.2})$$

$$\hat{\mathbf{y}} = \mathbf{S} \mathbf{y} = \Sigma_k \mathbf{R}_k \mathbf{y} \quad (\text{B.3})$$

B.2 states that $\mathbf{R}_k$ projects $\mathbf{y}$ onto each fitted additive term $\hat{\mathbf{f}}_k$, and B.3 implies $\Sigma_k \mathbf{R}_k = \mathbf{S}$. Next, we express each fitted additive term. Each comes from within each iteration of the back-fitting algorithm described in the main text, as a column vector:

$$\hat{\mathbf{f}}_k = \begin{pmatrix} x_{1k} \hat{\beta}_{1k} \\ \dots \\ x_{nk} \hat{\beta}_{nk} \end{pmatrix} \quad (\text{B.4})$$

Each estimated local parameter can be written as $\hat{\beta}_{ik} = \mathbf{e}_k \hat{\boldsymbol{\beta}}_i$, where $\mathbf{e}_k$ is the k^{th} row of an identity matrix whose dimension is the number of regressors plus 1. Thus, we can re-write $\hat{\beta}_{ik}$ as $\hat{\beta}_{ik} = \mathbf{e}_k (\mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{X})^{-1} \mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{y}$, and B.4 as:

$$\hat{\mathbf{f}}_k = \begin{pmatrix} x_{1k} \mathbf{e}_k (\mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{X})^{-1} \mathbf{X}' \mathbf{W}(u_i, v_i) \\ \dots \\ x_{nk} \mathbf{e}_k (\mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{X})^{-1} \mathbf{X}' \mathbf{W}(u_i, v_i) \end{pmatrix} \mathbf{y} \quad (\text{B.5})$$

Therefore, if the $\mathbf{R}_k$ is a hat matrix that projects $\mathbf{y}$ onto each fitted additive term $\hat{\mathbf{f}}_k$ (B.2), B.5 allows us to write each additive hat matrix as:

$$\hat{\mathbf{R}}_k = \begin{pmatrix} x_{1k} \mathbf{e}_k (\mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{X})^{-1} \mathbf{X}' \mathbf{W}(u_i, v_i) \\ \dots \\ x_{nk} \mathbf{e}_k (\mathbf{X}' \mathbf{W}(u_i, v_i) \mathbf{X})^{-1} \mathbf{X}' \mathbf{W}(u_i, v_i) \end{pmatrix} \quad (\text{B.6})$$

Now, we show how the additive hat matrix, and the MGWR hat matrix, can be estimated within the back-fitting algorithm. First, let $\mathbf{A}_j$ be the hat matrix of the partial model such that $\hat{\mathbf{f}}_k^* = \mathbf{A}_k (\hat{\mathbf{f}}_k + \hat{\boldsymbol{\epsilon}})$. Within the back-fitting algorithm, $\hat{\mathbf{f}}_k^*$ is the updated fitted additive term from the previous iteration fitted additive term $\hat{\mathbf{f}}_k$. In each iteration, the back-fitting algorithm calculates the fitted residual

$\hat{\epsilon} = \mathbf{y} - \Sigma_k \hat{\mathbf{f}}_k$; but B.3 implies that $\mathbf{S} = \Sigma_k \mathbf{R}_k$, so we can write the following equality:

$$\hat{\mathbf{f}}_k^* = \mathbf{A}_k(\hat{\mathbf{f}}_k + \hat{\epsilon}) = \mathbf{A}_k(\mathbf{R}_k \mathbf{y} + \mathbf{y} - \mathbf{S} \mathbf{y}) = (\mathbf{A}_k \mathbf{R}_k + \mathbf{A}_k - \mathbf{A}_k \mathbf{S}) \mathbf{y} \quad (\text{B.7})$$

The updated additive hat matrix is $\mathbf{R}_k^* = \mathbf{A} \mathbf{R}_k + \mathbf{A}_k - \mathbf{A}_k \mathbf{S}$ and the updated MGWR hat matrix is $\mathbf{S}^* = \mathbf{S} - \mathbf{R}_k + \mathbf{R}_k^*$. With $\mathbf{R}_k$ and $\mathbf{S}$, local standard errors can be computed. First, let us re-write B.4 as $\hat{\mathbf{f}}_k = \text{diag}(\hat{X}_k) \hat{\boldsymbol{\beta}}_k$, where $\text{diag}(\hat{X}_k)$ is shown in B.8. Then, we can re-write $\hat{\boldsymbol{\beta}}_k = [\text{diag}(\mathbf{X}_k)]^{-1} \mathbf{R}_k \mathbf{y} = \mathbf{C} \mathbf{y}$, where $\mathbf{C} = [\text{diag}(\mathbf{X}_k)]^{-1} \mathbf{R}_k$.

$$\text{diag}(\mathbf{X}_k) = \begin{pmatrix} x_{1k} & 0 & 0 & 0 \\ 0 & x_{2k} & 0 & 0 \\ 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & x_{nk} \end{pmatrix} \quad (\text{B.8})$$

Finally, we can write the variances of the local parameters $\hat{\boldsymbol{\beta}}_k$ as in B.9, where the normalized residual sum of squares from MGWR is defined by B.10 and the trace of $\mathbf{S}$ is defined by $\tau_1 = \text{trace}(\mathbf{S})$. Therefore, we can compute local standard errors for each parameter with: $SE(\hat{\boldsymbol{\beta}}_k) = \sqrt{\text{var}(\hat{\boldsymbol{\beta}}_k)}$.

$$\text{var}(\hat{\boldsymbol{\beta}}_k) = \text{diag}(\mathbf{C} \mathbf{C}' \hat{\sigma}^2) \quad (\text{B.9})$$

$$\hat{\sigma}^2 = \frac{\sum_N (y_i - \hat{y}_i)^2}{n - \tau_1} \quad (\text{B.10})$$

The local standard errors for all local parameters estimated with the MGWR are defined by $SE(\hat{\boldsymbol{\beta}}) = [SE(\hat{\boldsymbol{\beta}}_1), SE(\hat{\boldsymbol{\beta}}_2), \dots, SE(\hat{\boldsymbol{\beta}}_k)]$, and the pseudo-t test for each local parameter is given by:

$$\frac{\hat{\beta}_{ik}}{SE(\hat{\beta}_{ik})} \sim t_{n-\tau_1} \quad (\text{B.11})$$