Essays on the economics of climate change, international trade, and meat consumption

Leanne Cass

A thesis submitted for the degree of
Doctor of Philosophy

Department of Geography and Environment
London School of Economics and Political Science
July 2022

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 38,000 words.

Statement of co-authored work

I confirm that Chapter 4 was jointly co-authored with Dr Marion Dumas and I contributed 90% of this work.

Acknowledgements

I am indebted to my supervisors, Ben Groom and Eugenie Dugoua, whose support and guidance have been invaluable along the way to finding my footing as a researcher. In a similar vein, I would also like to thank Marion Dumas, who has been both a co-author and a mentor to me, and Alessandro Tavoni, who supported me during my first year at LSE.

I am also grateful to many people who kindly offered feedback, advice, and support along the way. In particular I would like to thank Giles Atkinson, Johannes Boehm, Frank Cowell, Roger Fouquet, Charles Palmer, Sefi Roth, Mario Larch, Katarina Lezova, Manuel Linsenmeier, Lutz Sager, Gregor Singer, and Frank Venmans. Email correspondence with Yoto Yotov has been especially helpful on my path to understanding structural gravity models, and I thank him for his guidance and intuitive answers to my technical questions. Many thanks to the examiners of this thesis, Rob Elliot and James Rising, for their time and care in reviewing this work, which yielded fruitful conversation and many helpful suggestions. I would also like to acknowledge the financial support I have received from LSE throughout my PhD, as well as the great opportunity I have had to be part of the community at the Grantham Research Institute on Climate Change.

On a personal note, I feel very lucky to have a large and supportive family - thank you to all of you. Most importantly, thank you Mum and Dad for being amazing role models and for supporting me unconditionally as I pursue my own passions and ambitions, even when they take me across an ocean. Thank you Lindsay, for listening and always knowing what to say to re-invigorate my self-belief. Thank you Nana, for always being my number one cheerleader in whatever I pursue. Thank you to my partner Jamie, who has stood patiently by me while I discovered for myself just how challenging a PhD can be.

Finally, I would like to acknowledge that the opportunity to pursue further education is a privilege that is unequally accessed, and the path that I walked to get here has been smoothed than by my White and socioeconomic privileges. Moreover, not long ago my position as a woman would have made this path largely unavailable to me. I am immensely grateful to have been able to pursue my intellectual interests this far and hope to contribute to widening this opportunity for others in the future.

Abstract

This thesis is composed of four environmental economics essays spanning the topics of climate change, international trade, and meat consumption. The first two chapters bring together approaches from the fields of international trade and climate econometrics. Using reduced-form empirical methods informed by theory, they investigate how international trade and weather shocks interact to affect economic outcomes. Chapter 1 estimates the impact of weather shocks on export sales relative to domestic market sales, finding that in the agricultural sector temperature shocks create additional barriers to international trade, as do precipitation shocks for the manufacturing sector in rainy countries. Chapter 2 examines the historical role of trade openness in the effect of temperature shocks on growth, finding some support for the hypothesis that connectedness to international markets may help to mitigate the impact of temperature shocks on economic growth. Chapter 3 continues the focus on international trade but shifts towards the topic of reduced meat consumption. Using insights from a structural gravity model for international trade, this paper explores how the impact of a tax on meat consumption in the EU can reach beyond borders and generate market signals via international trade mechanisms that undermine (to some degree) the aim of mitigating carbon emissions. Chapter 4 continues the theme of meat consumption but departs from the international trade perspective and instead focuses to the US market. Using granular data on households and stores, this paper interrogates popular stories about plant-based meat substitute products, finding that they remain a niche market and so far have shown limited potential for decarbonizing the food sector. This analysis highlights the potential need for policy intervention to spur demand growth and innovation in this product space. Overall, this thesis contributes to broadening our understanding of the challenges posed by climate change and climate policymaking.

Contents

List of Tables

List of Figures

Introduction

This thesis has two main themes: (i) The role of international trade in the economic impacts of climate change and climate policy, and (ii) potential pathways to mitigate greenhouse gas (GHG) emissions in the food and agriculture sector. I apply methods from the international trade literature to address gaps in our understanding of climate change economics, and focus on data-driven techniques to interrogate theoretical and popular understandings of these issues.

International trade plays an important role in modern economies in its own right, but trade also interacts with climate change, and understanding the role of trade in the effect of climate change and climate policy is an essential task in understanding how to deal with climate change. Copeland et al. (2022) point out that internationally traded goods account for over a quarter of global CO_2 emissions. In addition, climate change is a global environmental externality; all countries are affected and our success in tackling it depends on the combined efforts of many countries, so studying it in a global context is essential. Also, climate change and climate policy will have heterogeneous in the impacts on countries around the world, and the international trade network is a principal mediating factor in this spatial heterogeneity. This thesis explores several aspects in which international trade and climate change may interact: international trade may be particularly exposed to the impacts of climate change compared to other economic activity (as discussed in Chapter 1 of this thesis), but openness to international trade can also be a source of economic resilience to shocks such as climate change (Chapter 2). Finally, international trade linkages can play an important role in the effectiveness of policies to tackle climate change (Chapter 3). Despite the many ways in which international trade and climate change might interact, Dawson et al. (2020) find that trade-related topics have seen limited coverage in Intergovernmental Panel on Climate Change (IPCC) reports thus far; in a similar vein, Copeland et al. (2022) argue that the role of trade in environmental outcomes and policies is still contentious and a lot remains unknown in this field. This thesis contributes to addressing this gap. By bringing together the international trade and climate economics literatures we can better understand the economic impacts of climate change and how to cope with it, and importantly, we can draw on methodological insights from both these literatures to do so in a robust manner.

This thesis makes extensive use of gravity models of international trade, both empirically and theoretically. Gravity models have a long history in the international trade literature, and since the seminal contributions of Anderson and Wincoop (2003) and Eaton and Kortum (2002), which laid to rest the criticism that these models lack micro-foundations, these models have become the 'workhorses' of international trade analyses (Head and Mayer 2014). Structural gravity models combine empirical power with theoretical consistency. Gravity equations for trade are renowned for their empirical power in predicting bilateral trade flows, and meanwhile the structural gravity system is consistent with a variety of different micro-foundations and nests easily into general equilibrium models of the wider economy beyond trade (Yotov et al. 2016). Compared to traditional computable general equilibrium (CGE) models, structural gravity models are more transparent, more parsimonious, and are more empirically-grounded. The structural gravity system can be considered a mini CGE. While traditional CGE models can be opaque and usually involve a lot of parameters that are calibrated from outside studies, the structural gravity model is relatively simple, involving just a few parameters that are mostly calibrated or estimated directly from the data used in the study. Large, traditional CGE models offer greater complexity compared to structural gravity models, but this complexity can come at the expense of losing clarity on the extent to which model assumptions drive the results. While traditional CGE models make a valuable contribution to our understanding of general equilibrium issues related to international trade and climate economics, structural gravity models offer an important point of comparison in the difficult task of quantifying these effects.

Another prominent method in this thesis is the use of panel data econometrics, which I use both for estimating bilateral trade costs as well as quantifying the economic effects of climate change. The use of panel data econometrics to estimate bilateral trade costs is well-established in the structural gravity literature (Yotov et al. 2016). Likewise, the climate econometrics literature has seen a proliferation of studies employing panel data econometrics. Early empirical studies of climate change impacts using cross-sectional models were plagued by concerns of omitted variable bias. Linear panel models, such as Dell et al. (2012)'s empirical model, have been able to address some of these concerns by using fixed effects to control for time invariant unobserved heterogeneity between countries. However, the use of country fixed effects in empirical models that relate weather variables (such as temperature) to economic variables (such as income growth) means that identification of the effect of weather on economic outcomes comes from short-term deviations from mean weather within a given country. This identification strategy implies that the estimated effect is best interpreted as the economic impact of a short-term weather shock rather than the long-term impact of climate change. The validity of the linear panel model approach therefore rests heavily on the extent to which the (short-term) weather response differs from the (long-term) climate response. Arguments for why the climate response may differ from the weather response go in both directions, but the main concern is the potential that long-run adaptation causes the climate response to be less than the weather response.[1] Hsiang (2016) shows that if adaptation technology is continuous then the weather response and climate response are the same for a small change in climate, but his assumptions may be restrictive in many contexts.

Non-linear panel models help to address these concerns with linear panel models by allowing the marginal effect of weather to vary across the weather distribution. Kolstad and Moore (2020) explain that the marginal effect curve at high temperatures will be identified mainly from data from hot countries, and vice versa for the marginal effect at cold temperatures, so to the extent to which countries are already adapted to their local weather conditions, estimated effects from nonlinear panel models will be a mix of the long- and short-run responses. A recent contribution by Kahn et al. (2021) highlights a slightly different concern when using panel data models to estimate the economic impacts of climate change; they point out that temperature is non-stationary in most countries, which may lead to biased estimates in a panel data model of temperature and GDP growth. They use a panel ARDL model to estimate the effect of climate change on growth, with weather variables in deviations from their long-run moving averages. Future work could follow these developments to improve on the panel models used in this thesis. Overall, the extent to which panel data econometrics can fully account for the long-term and non-stationary nature of climate change is an important point of ongoing debate, but nevertheless the nonlinear panel models used in this thesis follow a large precedent in the literature, and are one of the most robust approaches available for the difficult task of empirically identifying the impact of climate change on economic outcomes.

Chapter 1 uses historical data on temperature, precipitation, and bilateral trade flows in a panel data model to investigate the sensitivity of export sales to weather shocks. Unlike previous papers studying the impact of climate and weather on trade,

[1] Arguments for why climate response may be larger than the weather response include the potential that economic systems can cope with one-off shocks but not repeated shocks, or similarly can cope with a marginal change in one weather variable such as temperature but not changes in multiple climate variables (Hallegatte et al. 2020). Moreover, some adaptation strategies may be available in the short run but not in the long run (for example, running down aquifers in a drought) (Auffhammer 2018).

I am able to disentangle impacts of weather shocks on the flow of trade from impacts on productivity. To do so, I combine a method from recent developments in the international trade literature with recent empirical approaches in climate change economics. More precisely, I use a model that controls for exporter-year fixed effects and thereby isolates the effect of weather shocks on export sales from productivity impacts on all sales. The results suggest that agricultural exports are more sensitive to temperature shocks compared to domestic market sales, particularly in hot countries. By comparison, manufacturing sector exports are relatively resilient to weather shocks, although increased precipitation in already rainy places seems to decrease exports relative to domestic sales. Economists usually conceptualize the economic damages of climate change as productivity impacts, but my results provide some evidence that weather and potentially climate can be an additional barrier to trade, so full economic damages of these shocks entail not only productivity impacts at the site of production, but also disruptions along the supply chain once goods leave the farm or factory. Given the close link between exports and economic development and welfare, these findings have economic significance. As well as the academic contribution to our understanding of the potential economic damages of weather shocks and climate change, this paper also contributes evidence for policy-making, highlighting the importance of aligning trade and climate change policy.

Chapter 2 investigates empirically the common hypothesis that openness to international trade could help countries to soften the economic impact of climate change. For example, fluctuations in productivity and comparative advantage due to climate change can be better adapted to if producers and consumers can access global markets. Some previous studies lend support for this hypothesis, but they have tended to rely heavily on structural approaches. In the spirit of the rapidly growing applied climate change economics literature, I take an empirical approach to investigating this question. Following recent developments in the international trade literature, I construct an instrument for trade openness in a manner consistent with international trade theory. I then use this instrument in an reduced-form model of GDP growth that follows approaches in the applied climate change economics literature. My results suggest that historically trade openness may help to mitigate the negative impact of temperature shocks on aggregate economic growth, which provides some empirical support to hypotheses that trade openness can help countries to adapt to climate change. However, results for the agriculture sector alone paint a more nuanced picture, suggesting that it is cold countries rather than hot countries that experience the beneficial effects of trade openness in the impact of temperature shocks on agricultural income growth. Given the particular vulnerability of the agriculture sector in hot countries to climate change, these results call into question the importance of trade openness as a measure to adapt to climate change. By testing the hypothesis that trade openness helps to moderate the impact of temperature on economic growth, this paper provides an important empirical counterpart to structural approaches to studying this topic. Moreover, this work contributes to both academic and policy debates on the interaction between international trade and climate change impacts and potential synergies between trade and climate policy.

In Chapter 3 I assess the global general equilibrium consequences of a tax on meat consumption in the European Union using a structural gravity framework, which has not yet been used to address this topic. I extend a standard structural gravity model of international trade to accommodate consumption taxes, and then I estimate and calibrate the parameters of this model using data on bilateral meat trade and income in the top 20 meat producers in the world. The results from simulating a 1% ad valorem tax on EU meat consumption suggest that incidence of the tax falls heavily on meat producers, particularly EU meat producers, while meat consumers around the world enjoy increased accessibility of meat and emerge as the relative 'winners' under the policy. Furthermore, the impact of the tax reverberates through the global network for meat trade and creates new bilateral trade opportunities. Global meat prices decrease and global meat sector income shifts away from the EU orbit. The results for the 1% tax suggest that global meat producer prices fall by 1.5% to 3.5% as a result of these global general equilibrium adjustments; under a 10% tax on EU meat consumption global meat prices decrease by as much as 25%. These estimates highlight the risk that reductions in meat consumption achieved by the tax within the EU may be offset by increases in meat consumption elsewhere. Given that the tax aims to tackle the global externality of climate change, these general equilibrium effects undermine the environmental aim of the policy to some extent. This paper contributes evidence to academic and policy discussions of how to tackle food-related GHG emissions, highlighting disadvantages of unilateral policy to tackle a global environmental externality.

Chapter 4 continues on the theme of decarbonizing the food sector. Using rich data on retailer sales and household purchases, we interrogate popular stories about the market for plant-based substitutes for animal products. In particular we investigate whether the market for these products is growing in such as way as to provide meaningful hope for reducing food-related GHG emissions, and we explore whether demand is a driver of innovation in these products. We develop a system to identify these plant-based substitutes in the data and create customized food groups tailored to our research question. Despite the hype in the media and amongst the private sector about the investment opportunities in this market, we do not find evidence that a widespread shift towards plant-based substitutes is beginning to take off. Instead, these products are still a niche market, dominated across both time and space by young and highly-educated households. We also find that entries of new plant-based substitute products tend to be related to very localized spending growth and levels, and more so compared to other food groups. This finding is suggestive evidence that demand may be playing a role driving innovation in plant-based substitute products. We do not tackle the endogeneity between innovation and demand, but using US demographic trends we explore the potential that some of this demand growth is exogenous to product innovations. Overall, this paper suggests that private-sector investors and consumers on their own are not driving a green transition in the food sector, and like the energy and transport sectors, policy intervention is likely necessary to spur a widespread shift towards reduced meat consumption.

Chapter 1

Weather shocks and international trade

1.1 Introduction

As climate change progresses and weather conditions become more volatile as a result, evidence-based policy-making necessitates that we understand the potential economic impacts of these changes. A rapidly developing literature has uncovered empirical evidence linking weather shocks and potentially climate change with economic outcomes. These studies have focused heavily on the impact of weather shocks on productivity, delivering partial equilibrium estimates of the effect of a marginal increases in temperature and precipitation on the 'size of the output pie' while holding all else constant, including the 'division of the pie' between exports versus domestic market sales. Meanwhile, a vast literature highlights the importance of international trade for economic welfare. Not only can exports be an important channel for growth and development for the exporting country, but international trade also makes products cheaper and more accessible for consumers around the world by allowing producers to exploit their comparative advantage. That is, the 'division of the pie' between exports and domestic sales can have important implications for economic welfare, but we understand relatively little about how weather and climate change might impact the accessibility of international markets.

This study investigates empirically whether temperature and precipitation shocks impact barriers to international trade. More specifically, I quantify the difference in the impact of weather shocks on the value of exports relative to sales in producers' domestic markets. While some of the results I present provide insight into the effects of weather shocks on overall output and sales, the key aim of this study is to test for a difference in the effect on exports versus domestic sales. Unlike previous empirical studies of international trade and weather shocks, which do not disentangle impacts on underlying productivity from a potential particular sensitivity of exports to weather shocks, I use an empirical model that allows me to test specifically for a difference in the effect of weather shocks on exports versus sales in the domestic market. The results reveal the impact of weather shocks and climate change on barriers to international trade. Several potential mechanisms could explain why changes in temperature or precipitation might affect barriers to international trade, and competing hypotheses imply that the sign of the effect is ambiguous.

First of all, producers may prioritize their domestic markets (often called a 'home bias'), so if production is disrupted by weather shocks, they reduce export quantities more than domestic sales. This hypothesis is briefly mentioned in Jones and Olken (2010). Under this mechanism exports are more sensitive to weather shocks than domestic sales, so we would expect a negative sign on the effect of interest (increased temperatures are associated with a decrease in exports relative to domestic sales). Furthermore, supply chains for goods sold internationally may be more sensitive to shocks compared to supply chains for domestically-sold goods; for example, international supply chains may be more sensitive to perishability compared to domestic supply chains, and export-related transport infrastructure may be particularly sensitive weather shocks. For example, Becker et al. (2013) discuss how the vulnerability of seaports to climate change and extreme weather events could negatively impact international trade, and Chinowsky et al. (2019) find that increased temperatures lead to costly delays in the US rail network. These mechanisms suggest that weather shocks may exacerbate barriers to international trade, implying a negative coefficient on the parameter of interest: exports are more sensitive to weather shocks than domestic sales. Moreover, the outcome of interest in this study is the value of bilateral trade, and so the effect of interest could also occur through a price mechanism. More specifically, trade barriers may insulate producers from global competition and allow them to increase prices in their domestic markets in the face of a weather-induced negative production shock while prices on export markets remain steady. This mechanism also implies a negative sign on the effect of interest: the value of exports decreases relative to the value of domestic sales.

A competing hypothesis comes from a large literature on the propensity to export that tells us that firms that export are different from firms that do not (Atkin et al. (2017); Görg et al. (2012)). Given that increased propensity to export is associated with mainly positive firm traits (e.g. higher productivity), this literature suggests that firms that export might be more resilient to weather shocks than firms that do not export. In this case, we would expect a positive sign on the effect of interest, implying that the balance of trade shifts more towards exports relative to domestic sales, or in other words that exports increase relative to domestic sales.

Finally, long-term contractual obligations could be stronger for exports compared to domestic sales, leading to differential effects of weather shocks. If long-term contracts commit producers to export quantities, then a negative impact on production due to unexpected weather shocks may result in producers prioritizing these contracts over domestic sales (implying a positive sign on the coefficient of interest). On the other hand, if contracts tie producers to specific price levels while prices in the domestic market are more flexible, then a weather shock and resulting decrease in production may lead to a decrease in the value of exports relative to domestic sales (negative sign on coefficient of interest).

A case study from the Philippines helps to illustrate how weather shocks might impact exports differently from domestic sales. In 2019, unusually dry weather in the Philippines caused by the El Niño effect led to an oversupply of around 2 million kg of mangoes. According to local news reports, this excess supply was mainly absorbed by the domestic market. Local prices decreased by more than half and the government even set up a campaign to promote mango sales in Manila. Export quantities also increased, but media reports said that poorly developed supply chains lacking in standardization and regulation dampened increases in exports. In this case study, the weather shock led to an increase in domestic sale quantities relative to exports.[1]

To estimate of the differential effect of weather shocks on exports relative to domestic sales, I combine gravity model estimation techniques from the trade literature with developments from the climate econometrics literature. In some cases I also estimate the effect of weather shocks on the underlying levels of exports and domestic sales but in the most robust specifications I am unable to estimate this effect. Estimating the effect of weather shocks on trade is not straightforward given the potential biases inherent in empirical trade models. Models with international trade flows as the dependent variable are essentially cross-country comparisons and must inevitably deal with a myriad of potential confounding variables; how much two countries trade with each other is affected by complex array of factors many of which are difficult to measure and observe. Accordingly, a huge body of work in international trade has focused on the best techniques to mitigate potential omitted variable bias. A key development has been the use of importer and exporter fixed effects to properly control for 'multilateral resistance', which has now become part of best practice standards for empirical trade studies (Baldwin and Taglioni 2006).

[1] For media coverage of this event see `https://www.bbc.co.uk/news/world-asia-48581857` and `https://www.aseantoday.com/2019/07/too-many-mangoes-a-bumper-harvest-puts-the-spotlight-on-filipino-supply-chains/`

However, following these best practices means that country-specific variables such as weather shocks are absorbed into fixed effects.

To overcome these challenges I follow the innovations in Heid et al. (2017) and Beverelli et al. (2018) to control for the multilateral resistance parameters with a full set of importer-time and exporter-time fixed effects while still estimating the effect of country-specific variables (such as weather shocks) on exports relative to domestic sales. This approach includes domestic as well as international trade flows in the model and interacts the variables of interest (temperature and precipitation in this case) with a dummy indicator for international sales. Heid et al. (2017) apply this methodology to measure the effects of most favoured nation (MFN) tariffs and "Time to Export" on international relative to domestic trade. Beverelli et al. (2018) build on the methodology of Heid et al. (2017) to estimate the effect of institutional quality on exports relative to domestic sales. They find that poor institutions hinder exports and their GE simulation suggests that this effect translates into notable impacts on GDP. Ultimately, the approach developed in these papers provides a more robust basis for causal inference compared to methods used in previous papers exploring the relationship between temperature and trade.

Previous literature has demonstrated that weather and climate have notable economic impacts on a macroeconomic level. A rapidly expanding area of work uses historical weather data to estimate empirically the impact of weather and climate and economic outcomes. A particularly strong focus in this literature has been the effect of weather and climate on GDP. Seminal contributions include Dell et al. (2012), Burke et al. (2015), Kalkuhl and Wenz (2020), and Newell et al. (2021). Not only has this body of work contributed an empirical basis for the economic damages associated with climate change, but it has also made strides in developing an appropriate methodology for modelling the effects of weather on economic outcomes and linking weather effects with climate change impacts. Two key methodological developments have been the use of panel data techniques to deal with biases in cross-sectional analyses and functional forms that allow for non-linear effects of weather on economic outcomes. I follow these developments in the climate change economics literature, using a panel data setting and allowing for nonlinear effects of weather shocks; the main specification is a quadratic functional form for the effect of temperature on trade, but I also explore alternative functional forms based on higher-order polynomials and number of degree days in temperature bins.

A few papers have explored (ex-post) the relationship between weather shocks and trade and have found some evidence that increased temperatures are associated with a reduction in exports. An early contribution by Jones and Olken (2010) finds that increased temperatures are associated with reduced export growth in poor countries. The magnitude of their estimate is larger than the effect of temperature on GDP estimated in Dell et al. (2012), which they suggest may indicate that exports are more sensitive to temperature than GDP. More recently, Osberghaus (2019) reviews the literature on the effects of natural disasters and weather variation on international trade. He finds that most studies of the effect of temperature on trade find that increased temperatures reduces trade, with the agriculture sector particularly affected. The effect of precipitation on trade is ambiguous across the literature. These papers often focus on linear effects of temperature on trade, while the rapidly-developing climate change economics literature seems to have reached a consensus that nonlinear temperature effects are very important. Moreover, it's unclear if previous studies are just finding (through the lens of trade data) the effect of temperature on aggregate income, or if they are uncovering a particular sensitivity of international trade to weather shocks. Finally, these previous papers have often had to forgo a robust set of fixed effects to deal with potential confounding variables.

Dallmann (2019)'s contribution is the closest in this literature to this paper. She studies the effect of weather shocks on international trade using a gravity-like empirical model, finding that increased temperature in the exporting country tends to reduce bilateral trade. She suggest that this effect seems to be largely driven by the impact of temperature on production, but does not explore whether exports are more or less sensitive to these productivity impacts compared to domestic sales. Another recent paper that explores how international trade and climate can interact to have economic impacts is Dingel et al. (2019), which shows that climate change is likely to increase inequality between countries because it will increase the spatial correlation of productivities and therefore lead to higher gains from trade for rich countries compared to poor countries.

In short, we know from previous work that increased temperatures are associated with decreased productivity, and that this effect seems to translate into a decrease in international trade. This paper builds on this work by providing a clear answer on whether exports are particularly sensitive to weather shocks compared to overall income, while employing a strict set of controls to deal with potential omitted variable bias. The results imply that for the agricultural sector increases in temperature are associated with a shift in the balance of trade away from exports and towards domestic sales, particularly in already hot places. By comparison, manufacturing sector exports are relatively resilient to weather shocks, although increased precipitation in already rainy places seems to decrease exports relative to domestic sales. Moreover, employing a strict set of controls to deal with potential omitted variable bias has a notable impact on these estimates.

The remainder of the paper is organized as follows. The next section outlines the methodology used in this study, providing a brief theoretical background before describing the empirical model, and the following section describes the data. Then section 1.4 presents and discusses the empirical results, and section 1.5 uses a sufficient statistic approach to explore the implications of these results for economic welfare. Finally, section 1.6 concludes.

1.2 Empirical methodology

The following section outlines the methodology used to estimate a differential effect of weather shocks on exports relative to domestic sales. First I outline a standard theoretical basis for the empirical trade model and explain the challenges of estimating the effects of unilateral variables such as weather shocks. Then I present the estimating equations, and finally I explain how to interpret the coefficient estimates and how they relate to the marginal effects of interest.

1.2.1 Background

The structural gravity model, often dubbed the 'workhorse' of international trade analyses, can be derived from several different micro-foundations, all of which lead to the following standard expression for bilateral trade (Head and Mayer 2014):

$$X_{ij,t} = \frac{Y_{i,t}}{\Omega_{i,t}} \frac{E_{j,t}}{\Phi_{j,t}} \phi_{ij,t} \tag{1.1}$$

In this expression, $X_{ij,t}$ is the value of bilateral trade sold by exporter i to importer j in period t. $Y_{i,t}$ and $E_{j,t}$ are the value of the exporter i's total production and the value of importer j's total expenditure in period t, respectively. $\phi_{ij,t}$ is the bilateral accessibility of exporter i to importer j; this term includes the cost to transport goods from i to j as well as less quantifiable trade barriers such as cultural and institutional differences between i and j.

$\Omega_{i,t}$ and $\Phi_{j,t}$ are the importer and exporter multilateral resistance parameters in year t; they describe how well-integrated buyers and sellers in a given country are into the global trade network in a given year. $\Omega_{i,t}$ summarizes how well sellers in country i can access buyers around the world, and $\Phi_{j,t}$ summarizes how well consumers in country j can access products from around the world (Head and Mayer 2014). These parameters are essential components of the model, and not controlling for them properly has been dubbed the "gold medal mistake" of estimating structural gravity models (Baldwin and Taglioni 2006). Standard practice in a panel data setting is to control for these terms using importer-time and exporter-time fixed effects, and Head and Mayer (2014)'s Monte Carlo simulations demonstrate the superiority of this approach over other ways to control for the multilateral resistances. However, these importer-time and exporter-time fixed effects absorb all country-specific characteristics that are invariant across trade partners, preventing the researcher from estimating the effect of country-specific variables such as GDP, national policies, institutions, and weather. This challenge is the main difficulty in studying the effect of weather shocks on trade; we have a trade-off between including country-specific variables such as temperature and precipitation in the above model and using best practices for robust gravity model estimation.

Head and Mayer (2014) review possible approaches to estimating country-specific effects in gravity models; given the potential pitfalls of the approaches they consider, they recommend that researchers estimate several different specifications since none of them are an ideal solution. One common way that papers deal with this challenge is forgoing the importer-time and exporter-time fixed effects. For example, Dallmann (2019) estimates the effect of temperature and precipitation on international bilateral trade by not including importer and exporter fixed effects and instead relying on observable country-specific variables (such as GDP) and country-pair fixed effects to deal with potential endogeneity. The benefit of this approach is that the researcher is able to identify the direct effect of weather variables on bilateral trade flows. The key disadvantage of this approach is that it cannot control for unobservable potential confounding variables which vary at the importer-time or exporter-time level and affect bilateral trade and are correlated with the weather variables. For example, an exporter's overall connections to the global trading network ($\Omega_{i,t}$ in Equation 1.1, known as outward multilateral resistance in the gravity literature), is an important determinant of bilateral trade. If weather shocks affect one bilateral relationship, this effect will spill over into the exporter's other bilateral relationships via their multilateral resistance. Without exporter-year fixed effects to control for outward multilateral resistance, we cannot isolate the direct effect effect of weather shocks on trade from the effect of outward multilateral resistance. Finally, weather shocks are certainly correlated with underlying productivity in a given year, so without exporter-year fixed effects we cannot identify whether exports are particularly sensitive to weather shocks relative to overall sales.

This paper takes a novel approach to overcoming the challenges associated with estimating the effect of weather on trade. I follow the method developed in Heid et al. (2017) and Beverelli et al. (2018) to control for multilateral resistances with importer-time and exporter-time fixed effects and estimate the effect of temperature shocks on international *relative* to domestic trade. The cornerstone of this approach is to include domestic trade flows (i.e. $i = j$) in the model. Heid et al. (2017) show that this design enables the researcher to estimate the interaction between a dummy variable indicating international (versus domestic) trade and the country-specific variable of interest (e.g. temperature). For a proof that the parameter of interest is identifiable (and not collinear with any other model parameters) see the appendix of Heid et al. (2017). Importantly, this method cannot provide an estimate of the direct effects of temperature and precipitation on all sales (domestic and international), because they are absorbed by the fixed effects. However, this model does provide an estimate of the differential effect of weather shocks on exports compared to domestic sales. This estimate provides insight into whether exports may be more or less sensitive to weather shocks compared to domestic sales, an issue that hasn't been addressed by Dallmann (2019) or other previous literature. In other words, weather shocks may affect not simply how much is produced and sold overall, but also where these sales are made (domestic versus foreign markets).

1.2.2 Empirical model

To answer this question of whether temperature and precipitation differentially affect exports relative to domestic sales, I use an the empirical counterpart to the theoretical gravity model in equation 1.1. A common approach to forming an estimating equation from a multiplicative model such as equation 1.1 is to log-linearize the expression and use the OLS estimator. However, Silva and Tenreyro (2006) show that in the presence of heteroskedasticity (which is ubiquitous in trade data), the OLS estimator is biased when applied to a log-linear version of a multiplicative model. As a result, standard practice in the applied trade literature is to use the PPML estimator to estimate equation 1.1. The PPML estimator also has the advantage of being able to take account of zero trade flows, which are another prominent feature of trade data (Yotov et al. 2016).

I start with an empirical version of Equation 1.1 that is similar to the main specification in Dallmann (2019):

$$\begin{aligned} X_{ij,t} =& exp[h(T_{it}) + g(P_{it}) + +\rho_1 ln(GDP_{i,t}) + \rho_2 ln(GDP_{j,t}) \\ &+ \mu_{ij} + \alpha RTA_{ij,t} + YEAR_t] \times \varepsilon_{ij,t} \end{aligned} \tag{1.2}$$

$X_{ij,t}$ is the value of bilateral trade flows from exporter i to importer j in year t, and importantly this variable includes within-country sales - i.e. cases when $i = j$. The relationship between temperature and bilateral trade is given by:

$$h(T_{it}) = \theta_1 T_{it} + \theta_2 T_{it}^2 + INTL_{ij} \times (\theta_3 T_{it} + \theta_4 T_{it}^2)$$

T_{it} is annual mean temperature in country i in year t. Unlike previous papers, this model includes an interactive term with the INTL_{ij} dummy variable, which equals 1 if $i \neq j$; that is, it equals 1 when $X_{ij,t}$ represents international rather than domestic sales. θ_3 and θ_4 are the key coefficients of interest in this paper; they tell us if temperature shocks impact exports differently from domestic sales. A statistically significant estimate on this interactive term suggests that temperature affects not just how much is produced but also where that production tends to be sold.

Analogously, the relationship between precipitation and bilateral trade flows is given by $g(P_{it}) = \gamma_1 P_{it} + \gamma_2 P_{it}^2 + INTL_{ij} \times (\gamma_3 P_{it} + \gamma_4 P_{it}^2)$, where P_{it} is total annual precipitation in country i in year t. As with temperature, I allow for non-linear effects of precipitation on trade and I allow the effect of precipitation shocks to differ for exports compared to domestic sales.

The exporter-importer fixed effects in the equation above, μ_{ij}, control for time-invariant factors that affect the accessibility of import market j to exporter i. These controls absorb a myriad of factors that affect trade costs such as distance, geography, and cultural ties. Alongside these time-invariant drivers of trade costs, we would expect that changes in trade agreements over the sample period also affect trade costs, and I control for these effects with the $RTA_{ij,t}$ dummy variable, which indicates whether exporter i and importer j are part of a common regional trade agreement in year t. $YEAR_t$ is a year fixed effect, which controls for any global shocks in a given year such as a recession or the El Niño effect. $ln(GDP_{i,t})$ and $ln(GDP_{j,t})$ are the natural log of GDP in year t in the exporting and importing country respectively. These variables control for economic size, and are the empirical counterparts to $Y_{i,t}$ and $E_{j,t}$ in equation 1.1.

These choices for modelling the relationship between weather and trade flows follow developments in the climate economy literature. Following Dell et al. (2012), the use of panel data techniques to deal with the biases in cross-sectional analyses has become widespread in studies estimating the effects of weather and climate on economic outcomes. A panel specification with country fixed effects means that the model identifies the effects of weather shocks (deviations from countries' average weather) on economic outcomes; Kolstad and Moore (2020) explain that in a linear model these effects are short-run responses, and if adaptation opportunities are strong then extrapolating climate change effects from the effects of weather shocks is problematic. One way to deal with this issue to some extent is to introduce non-linearities into the effect of weather shocks on economic outcomes. Burke et al. (2015) make a seminal contribution demonstrating the importance of allowing for non-linearities in these relationships. Kolstad and Moore (2020) explain that allowing for non-linear effects means that the estimate is a mix of short- and long-run responses. The main specifications in this paper follow Burke et al. (2015) in using a quadratic functional form for $h(T_{it})$, but for robustness I also use specifications based on degree days. Compared to previous studies investigating the effect of weather on trade (which mainly use linear functional forms), this approach should help to address the challenge of connecting estimates of weather effects to climate change effects to some extent.

A key weakness in equation 1.2 is the lack of controls for multilateral resistance, $\Omega_{i,t}$ and $\Phi_{j,t}$. To address this issue, my preferred specification follows the approach in Beverelli et al. (2018), introducing exporter-year and importer-year fixed effects to control for these parameters:

$$\begin{aligned} X_{ij,t} = & exp[h(T_{it}) + g(P_{it}) + \pi_{i,t} + \chi_{j,t} + \mu_{ij} + \alpha FTA_{ij,t} \\ & + \eta INTL_{ij} \times YEAR_t] \times \varepsilon_{ij,t} \end{aligned} \tag{1.3}$$

In this specification, anything that varies at the exporter-year and importer-year level, such as GDP, is absorbed by the fixed effects. The direct effects of temperature and precipitation on trade are also absorbed into the exporter-year fixed effects, but the relative effects of weather on exports compared to domestic sales is identifiable. In other words, θ_1 and θ_2 as well as γ_1 and γ_2 are no longer identifiable but we can still obtain estimates for θ_3, θ_4, γ_3, and γ_4. The $INTL_{ij} \times YEAR_t$ dummy variables control for the average level of globalization in a given year across all countries, an innovation that Bergstrand et al. (2015) find plays an important role in reducing bias in empirical gravity models.

In the results presented below, I present results based on equations 1.2 and 1.3 to enable comparisons with previous literature and also illustrate the impact that the various sets of fixed effects have on the results. Equation 1.3 is the preferred specification throughout this paper because it controls most robustly for the myriad potential biases in trade models. Once again, this specification cannot deliver estimates of the effect of temperature and precipitation on total sales, but it can provide estimates of the effect of weather shocks on exports *relative* to domestic sales, which is the key parameter of interest in this study.

1.2.3 Interpreting the model estimates

To compute full marginal effects of temperature or precipitation we must take into account the quadratic and interactive terms. First, note that since the empirical specifications uses the PPML estimator, the coefficients are semi-elasticities: the proportional change in bilateral trade, $X_{ij,t}$, for a one unit change in the variable of interest. These semi-elasticities of bilateral trade with respect to temperature and precipitation are given by:

$$\beta_{Temp} = \theta_1 + 2\theta_2 T_{it} + INTL_{ij} \times (\theta_3 + 2\theta_4 T_{it})$$
$$\beta_{Precip} = \gamma_1 + 2\gamma_2 P_{it} + INTL_{ij} \times (\gamma_3 + 2\gamma_4 P_{it})$$

These parameters describe the $\beta \times 100\%$ change in bilateral trade associated with a 1 degree increase in annual mean temperature (ceteris parabis) and a 1 metre increase in total annual precipitation (ceteris parabis). For domestic sales ($\text{INTL}_{ij} = 0$), the semi-elasticities are simply $\beta_{Temp} = \theta_1 + 2\theta_2 T_{it}$ and $\beta_{Precip} = \gamma_1 + 2\gamma_2 P_{it}$, while for export sales the semi-elasticities include the interactive terms.

For the specification given by equation 1.2 we can identify estimates for the full semi-elasticities, but for the empirical model given by equation 1.3, we can only identify the interactive term, since θ_1 and θ_2 and γ_1 and γ_2 are absorbed into the fixed effects and not estimable (as explained above). These interactive terms, $INTL_{ij} \times (\theta_3 + 2\theta_4 T_{it})$ and $INTL_{ij} \times (\gamma_3 + 2\gamma_4 P_{it})$, are the difference (in percentage points) in the semi-elasticity for exports compared to domestic sales associated with 1 degree increase in temperature or a 1 metre increase in precipitation, ceteris parabis. They answer the central question of interest in this study: whether temperature and precipitation shocks have a differential effect on exports versus domestic sales. For ease of comparison across specification results, I report only the interactive terms of the estimated semi-elasticity estimates, regardless of whether or not the full elasticity is identified.

1.3 Data

The empirical model outlined above requires a cross-country panel dataset of bilateral trade flows, including domestic trade, plus data on regional trade agreements and data on weather in the exporting country. I outline the sources and construction of these variables below. The final dataset spans manufacturing and agriculture trade in 67 countries over 1991-2017; it is an unbalanced panel due to missing trade data for some years for some countries. Table 1.1 lists descriptive statistics for the model variables. See the appendix to this chapter for a list of countries included in the model.

International trade flows. Data on international bilateral trade flow comes from UN Comtrade for the manufacturing sector (United Nations 2021) and the FAO detailed trade matrix for the agriculture sector, which is part of the FAOSTAT database (Food and Agriculture Organization of the United Nations 2021). I mainly

Table 1.1: Descriptive statistics for main variables

	Mean	Std deviation	Min	Max
$\text{Trade}_{ij,t}$ (Billion USD)	5.0	125.52	0.0	14508.64
Manu. $\text{trade}_{ij,t}$ (Billion USD)	4.73	118.37	0.0	13647.25
Ag. $\text{trade}_{ij,t}$ (Billion USD)	0.31	8.55	0.0	901.78
$\text{Temperature}_{i,t}$ (°C)	15.6	5.7	3.48	28.14
$\text{Precipitation}_{i,t}$ (m)	9.35	5.99	0.12	35.33
$\text{RTA}_{ij,t}$	0.24	0.42	0.0	1.0
$\text{GDP}_{i,t}$ (Billion USD)	671.92	1856.9	0.71	19519.35
Low $\text{income}_{i,t}$	0.27	0.44	0.0	1.0
Weak $\text{institutions}_{i,t}$	0.49	0.5	0.0	1.0

use reported imports, which should be more reliable than reported exports, but I check for instances when a country reports no imports but a partner country reports exports and fill in missing values with these reported exports. The FAOSTAT trade matrix does not include observations for FAO items 328 and 2631 (groundnuts and cotton), so I use bilateral trade data from Comtrade for these items. As is common practice in the international trade literature, I assume that missing bilateral trade represents zero trade and therefore obtain a complete matrix of international bilateral trade flows by filling in missing values with zero.

Domestic trade flows. The main challenge in compiling a data for this study is obtaining observations of domestic trade flows, which are not readily available. Following the approach in Beverelli et al. (2018), I construct domestic trade as the difference between production and total exports: $X_{ii} = Y_i - \sum_{i \neq j} X_{ij}$. Crucially, since international trade flows are observed in gross values, I use gross values of production (not value-added). Also, I use aggregate exports to all countries reported in the data, not just the 67 countries in my sample. For manufacturing I primarily use the UNIDO database for gross production data (UNIDO 2020), but where available I fill in missing values with data from the CEPII TradeProd dataset (Sousa et al. 2012). For agriculture, I use FAOSTAT's value of gross production data series. I use data starting from 1991 for all sources, since this is the earliest available year for the FAOSTAT value of gross production data. As explained above, observations of domestic trade are a cornerstone of the methodology used in this paper; however, the limited availability of data on domestic trade is the most significant data limitation faced by this paper and defines the coverage of countries and years in the sample.[2]

[2]For further details on how I construct domestic trade flows and select the sample of countries, see the appendix of this chapter.

Figure 1.1: Average annual temperature in sample countries, 1991-2017

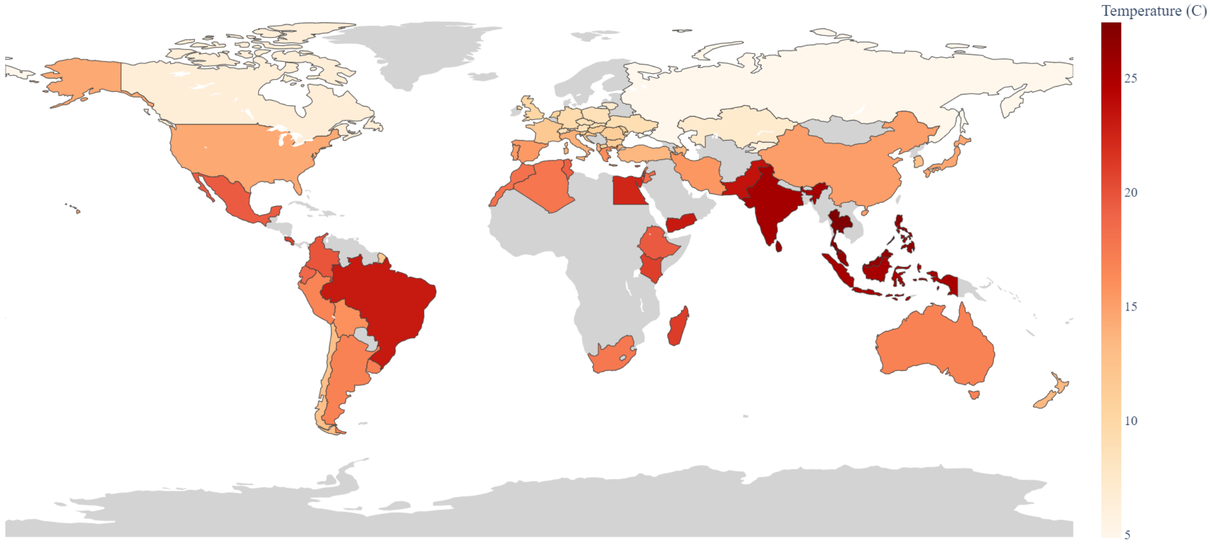

Regional trade agreements. I use the RTA dummy variable from the CEPII gravity database (Head and Mayer 2014). This variable indicates whether or not two countries have a regional trade agreement in a given year. For domestic trade observations, I set this this dummy equal to zero. For some specifications I also use GDP from this database, and for robustness checks I identify tropical countries using the latitude of capital cities variable from this database.

Weather. The country-level annual mean temperature variable is constructed from the ERA5 data for hourly grid-level temperature in Kelvin (Hersbach et al. 2020). The country-level total annual precipitation variable is from the University of Delaware Terrestrial Precipitation: 1900-2017 Gridded Monthly Time Series (V 5.01). (Willmott and Matsuura 2018). All transformations of the weather data (e.g. calculation of annual mean and square of annual mean) are done at the grid level before spatial aggregation. Country-level observations are constructed as population-weighted averages across grid cells, using the Gridded Population of the World v4 dataset for the year 2000 (Center for International Earth Science Information Network - CIESIN - Columbia University 2018). Figures 1.1 and 1.2 illustrate the average of these variables across the sample period.

Exporter characteristics. I explore heterogeneity in the effects of weather on trade according to a couple exporter characteristics: income and institutions. The Low income$_{i,t}$ dummy variable indicates whether country i was classified by the World Bank as 'Low income' or 'Lower middle income' in a given year (The World Bank 2022). The Weak institutions$_{i,t}$ dummy variable indicates whether country

Figure 1.2: Average annual precipitation in sample countries, 1991-2017

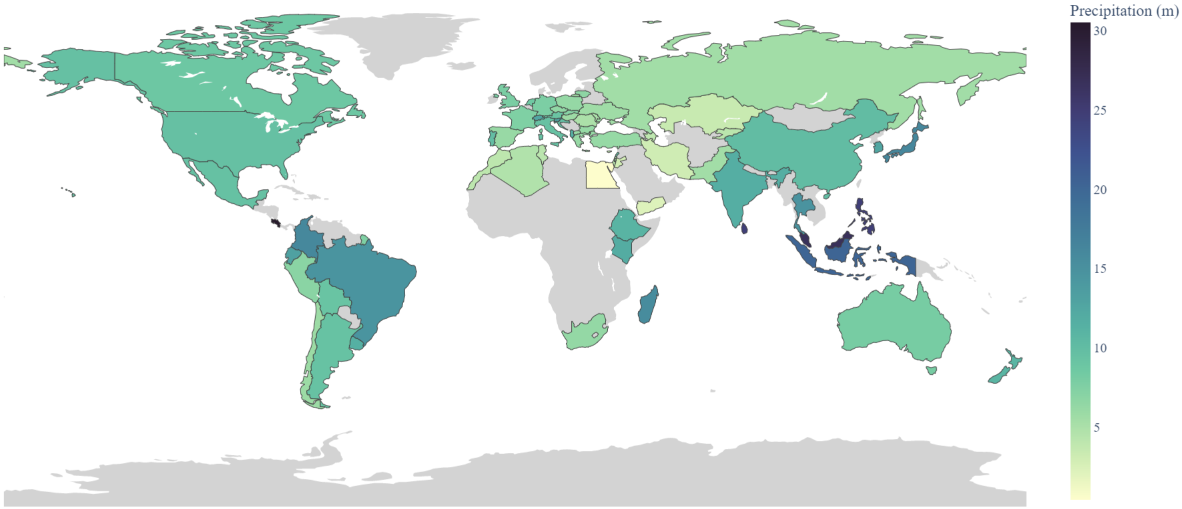

i is below the median observed value in year t for an institutional quality index. The institutional quality index is constructed as an unweighted average of the six variables in the World Governance Indicators dataset: control of corruption, government effectiveness, political stability and absence of violence, rule of law, regulatory quality, and voice and accountability (The World Bank 2020).

1.4 Results

The following discussion of results is organized as follows. First, I present estimation results for the effect of weather shocks on aggregate trade (manufacturing and agriculture combined). Next, I break the analysis down to the sector level and show results for agriculture and manufacturing separately. Then I explore heterogeneity in the effect of weather shocks on trade according to exporter characteristics such as income and institutional quality. Finally, I explore alternative functional forms for the weather-trade relationships.

1.4.1 Results for aggregate trade

Table 1.2 shows coefficient estimates for the gravity model described in Section 1.2.2, with marginal effect estimates for each specification in the bottom panel of the table. In all columns the dependent variable is the total nominal value of bilateral trade across the manufacturing and agriculture sectors. The marginal effects shown in the bottom panel are the estimated *difference* in the marginal effect of temperature or precipitation on exports relative to domestic sales. This estimate corresponds to

Table 1.2: Effects of weather shocks on aggregate trade

	(1)	(2)	(3)
$\text{Temp}_{i,t}$	1.2090	-0.5519	
	(1.1285)	(1.4879)	
$\text{Temp}^2_{i,t}$	-0.0022	0.0010	
	(0.0020)	(0.0026)	
$\text{INTL}_{ij}\times$ $\text{Temp}_{i,t}$	-0.3639	1.5718	0.4760
	(1.5696)	(2.0310)	(1.7891)
$\text{INTL}_{ij}\times$ $\text{Temp}^2_{i,t}$	0.0009	-0.0028	-0.0008
	(0.0028)	(0.0036)	(0.0031)
$\text{Precip}_{i,t}$	0.0010	-0.0115	
	(0.0137)	(0.0124)	
$\text{Precip}^2_{i,t}$	0.0001	0.0007*	
	(0.0004)	(0.0004)	
$\text{INTL}_{ij}\times$ $\text{Precip}_{i,t}$	0.0048	0.0157	0.0135
	(0.0184)	(0.0134)	(0.0117)
$\text{INTL}_{ij}\times$ $\text{Precip}^2_{i,t}$	-0.0005	-0.0011***	-0.0009**
	(0.0005)	(0.0004)	(0.0004)
$\text{RTA}_{ij,t}$	0.2613***	0.1640***	0.1738***
	(0.0509)	(0.0340)	(0.0398)
$\ln(\text{GDP}_{it})$	0.5804***	0.5644***	
	(0.0712)	(0.0551)	
$\ln(\text{GDP}_{jt})$	0.6717***		
	(0.0502)		
Observations	109092	109092	109092
Difference in marginal effect on exports relative to domestic sales:			
Temp at 15°C	0.1314***	-0.0138	-0.0073
	(0.0230)	(0.0300)	(0.0269)
Temp at 25°C	0.1486**	-0.0688	-0.0241
	(0.0755)	(0.0969)	(0.0859)
Precip at 8 m	-0.0024	-0.0023	-0.0016
	(0.0121)	(0.0087)	(0.0085)
Precip at 20 m	-0.0133	-0.0292***	-0.0242**
	(0.0090)	(0.0096)	(0.0105)
Year FE	✓		
INTL-Year dummies		✓	✓
Importer-Year FE		✓	✓
Exporter FE		✓	
Exporter-Year FE			✓

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

Notes: Standard errors (in parentheses) are clustered by exporter-importer pairs. All specifications control for exporter-importer FE.

the blue lines in the plots in Figures 1.3 and 1.4, which show the marginal effect estimates across levels of temperature and precipitation observed in the sample. The red lines in these figures illustrate the estimated marginal effect on domestic sales, and the green lines represent the estimated marginal effect on export sales.

Figure 1.3: Estimated marginal effect of temperature on aggregate trade

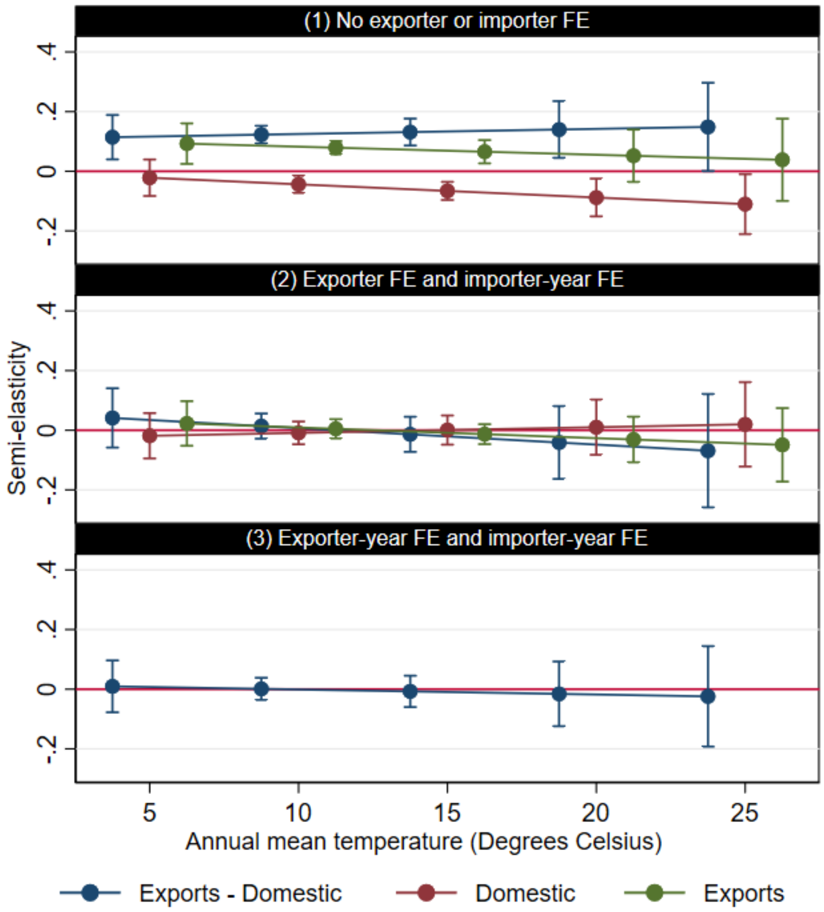

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.2. 'Exports - Domestic' (the blue line) denotes the difference in the semi-elasticity for exports versus domestic sales - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

Moving from left to right across Table 1.2, each column includes a progressively more strict set of controls for potential confounding variables. In column (1), I start with a specification similar to the main specification in Dallmann (2019), corresponding to Equation 1.2 above. Exporter-importer fixed effects absorb time invariant bilateral variables such as distance and sharing a border as well as any unobservable time invariant factors that affect bilateral accessibility. The $RTA_{ij,t}$ dummy controls for variation over the sample period in trade agreements. Although the results suggest that the effects of temperature and precipitation on trade are statistically insignif-

Figure 1.4: Estimated marginal effect of precipitation on aggregate trade

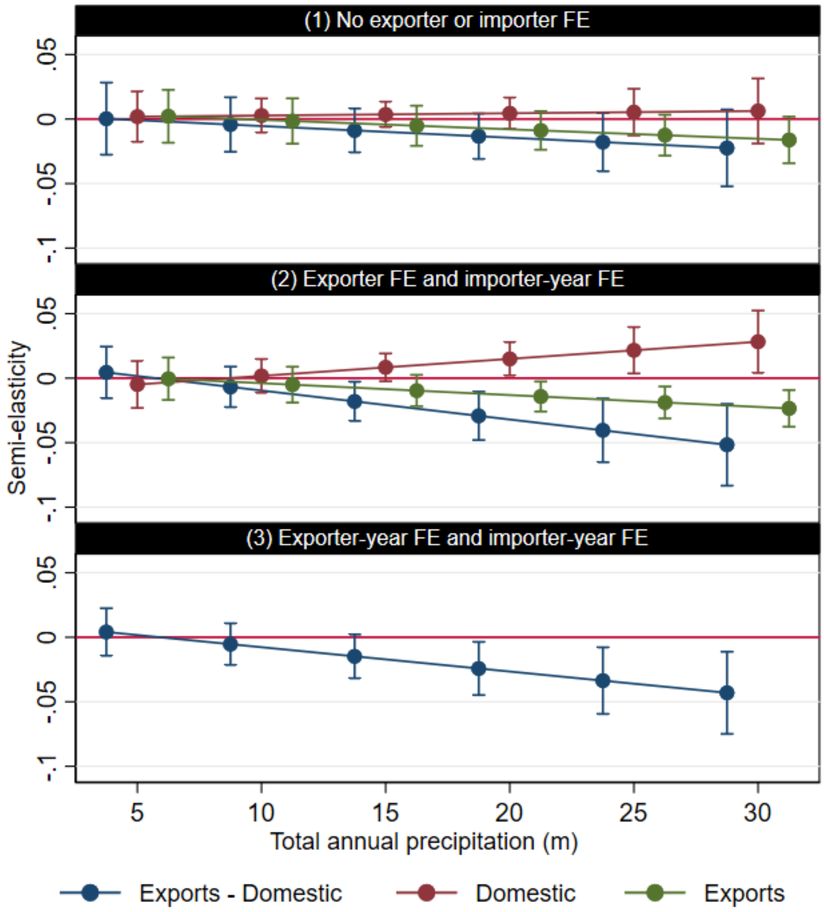

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.2. 'Exports - Domestic' (the blue line) denotes the difference in the semi-elasticity for exports versus domestic trade - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

icant on average across the sample, at 15°C (approximately the sample median of the temperature variable) the results indicate that the elasticity of exports to a temperature shock on exports is 13.1 percentage points higher than the elasticity of domestic sales. This results implies that at an annual mean temperature of 15°C, domestic sales are more sensitive to a temperature shock than exports. However, we should keep in mind that this specification controls for exporter and importer GDP, but otherwise does not control for the multilateral resistance faced by the exporter and importer. As explained above, this specification commits the 'gold medal mistake' of not properly controlling for exporter and importer multilateral resistances.

The estimates in columns (2) and (3) of Table 1.2 demonstrate that including a robust set of controls can have a notable impact on the coefficient estimates for the weather variables. Column (2) goes some way towards remedying the 'mistake' in column (1) by adding an importer-year fixed effects (FE) as well as time-invariant exporter FE. This specification reflects Table 22 in the appendix of Dallmann (2019). The importer-year FE control robustly for importer multilateral resistance, absorbing any factors affecting an importer's overall connection to the trading network in a given year. The exporter FE absorb any time invariant factors affecting an exporter's overall connection to consumer markets; these FE control imperfectly for exporter multilateral resistance but allow identification of the full effect of weather on trade (i.e. θ_1, θ_2, γ_1, and γ_2 are not absorbed into the FE). Column (2) also includes INTL-year dummies, as suggested by Beverelli et al. (2018)) to account for the effects of increasing globalization over the sample period. In contrast to the results in column (1), in column (2) the estimated marginal effects of temperature are insignificant across the full distribution of the temperature variable. Meanwhile, precipitation has a statistically significant marginal effect on trade at upper levels of the precipitation distribution, as shown in the second panel of Figure 1.4. At 20 metres total annual precipitation, the semi-elasticity of exports in response to an additional metre of precipitation is 2.9 percentage points lower than the semi-elasticity of domestic sales.

Finally, column (3) of Table 1.2 adds exporter-year FE. This specification corresponds to Equation 1.3. Note that the overall effect of weather is absorbed into these FE and only the interactive term with the INTL dummy can be estimated now. In other words, the effect of weather on exports *relative* to domestic sales is identified but we cannot identify the underlying level effects. However, this specification follows the consensus established in the literature to control for multilateral resistances with both importer-year and exporter-year FE. The results are similar to those in column (2). The impact of temperature on exports relative to domestic sales is statistically insignificant. Meanwhile, as illustrated in the bottom panel of Figure 1.4, in countries with already high levels of precipitation, an increase precipitation leads to a decrease in exports relative to domestic sales: the balance of trade shifts away from international markets and towards the domestic market.

A notable feature of the precipitation variable is that its distribution is quite wide, with low income tropical countries dominating the upper tail. With this feature of the data in mind, I do a leave-one-out analysis to test whether the significant marginal effect of precipitation at 20 metres annual precipitation is driven by an outlier country in the precipitation distribution. I compile a list of tropical and low income countries from within my sample countries according to the following characteristics: (i) the latitude of a country's capital is within the Tropic of Capricorn and Tropic of Cancer[3] and (ii) the country is classified as 'Low income' or 'Lower

[3] I obtain this data from the CEPII gravity database (Head and Mayer 2014).

Figure 1.5: Robustness of precipitation effect to leaving out tropical low income countries

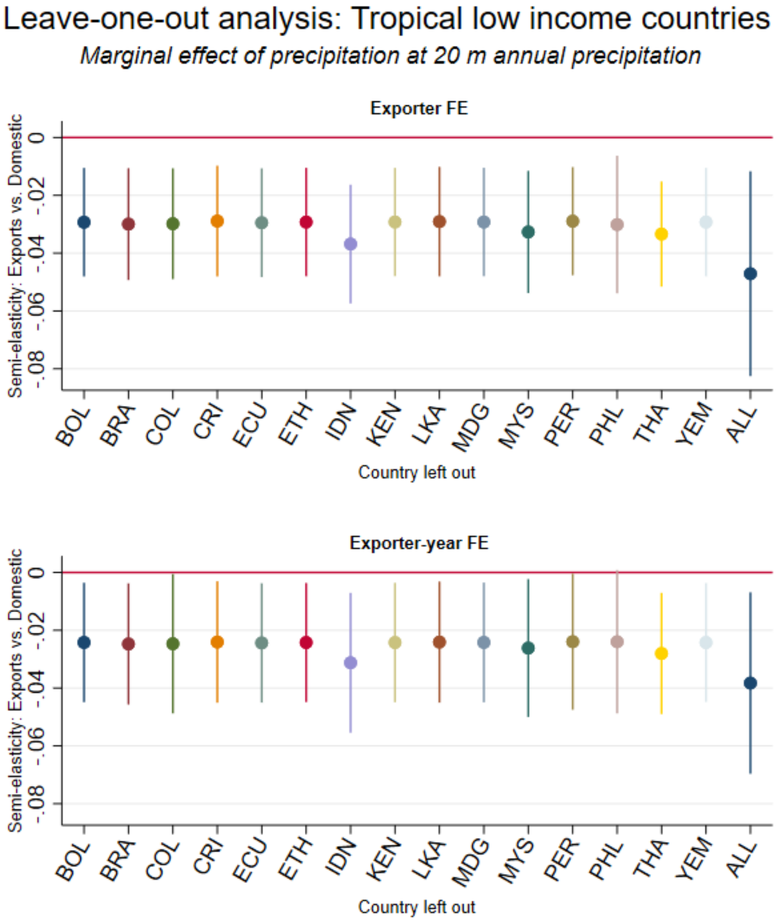

Notes: Countries labelled according to their 3-digit ISO codes; see section 1.A, the appendix to this chapter, for the correspondence between ISO3 codes and country names. 'ALL' denotes all tropical low income countries. Error bars represent 95% confidence intervals.

middle income' by the World Bank at any point during the sample period (The World Bank 2022). These criteria result in a list of 15 countries, which includes the countries at the very top of the distribution of the precipitation variable. Next, I iteratively re-estimate the model, each time leaving out one of the tropical low income countries, and then finally leaving out all of these countries. I do this procedure for both the model in column (2) of Table 1.2, which controls for exporter fixed effects, as well as the model in column (3), which controls for exporter-year fixed effects.

As Figure 1.5 illustrates, the estimated difference in the marginal effect of precipitation on exports relative to domestic sales at 20 m annual precipitation is fairly stable across each of these iterations, suggesting that the result is not driven by a single outlier country. The main potential exception occurs when the Philippines is left out of the model with exporter-year fixed effects; in this case the marginal effect of a precipitation shock at 20 metres annual precipitation does not have a significantly different effect on exports relative to domestic sales. This result suggests that the Philippines may be an important driver of the finding that exports are particularly sensitive to precipitation shocks in rainy countries. Nevertheless, when all tropical low income countries are left out of the analysis, the estimated marginal effect at 20 metres remains statistically significant (although standard errors widen, which is not surprising given that removing these countries eliminates all observations of annual precipitation above 20 meters and so the estimate is an out-of-sample prediction). Overall, the leave-one-out analysis suggests that the result that exports are relatively sensitive to precipitation shocks is fairly robust across samples and does not simply reflect a particular sensitivity of tropical low income countries to weather shocks. Moreover, the magnitude of the marginal effect of precipitation at 20 metres may be slightly larger when all tropical and low income countries are left out of the sample compared to when they are included, suggesting that exports from tropical low income countries may even be less sensitive to precipitation shocks compared to exports from other countries, perhaps because they are better adapted.

Overall, these estimates do not identify a significant effect of temperature shocks on aggregate trade. In rainy places, exports seem to be more sensitive than domestic sales to additional precipitation, but otherwise these results do not indicate a particular sensitivity of exports to weather shocks. Moreover, these results demonstrate the importance of including a robust set of controls for multilateral resistances and globalization effects when estimating the effects of weather shocks on trade. Contrary to previous studies that have found a significant negative impact of temperature on trade, once a robust set of fixed effects are included, as is suggested by the gravity literature, this study finds that the effects of temperature on aggregate trade are statistically insignificant. Of course, including a demanding set of fixed effects reduces concerns of omitted variable bias, but it also reduces the identifying variation in the model, so these results could potentially reflect a lack of statistical power rather than true zero effects. Moreover, these results for aggregate trade may hide sector-specific effects of weather shocks on exports relative to domestic sales. Accordingly, the next section investigates these relationships separately for the agriculture and manufacturing sectors.

1.4.2 Results by sector

Previous studies, such as Dallmann (2019), have found that the sensitivity of international trade to weather shocks varies by sector, and Osberghaus (2019) notes that several studies on this topic have often found that agricultural trade is particularly affected by temperature shocks. Following this precedent, Table 1.3 presents results

Table 1.3: Separating manufacturing and agricultural trade

	(1) Manu	(2) Manu	(3) Ag	(4) Ag
$\text{Temp}_{i,t}$	-0.6425		-2.7630	
	(1.4945)		(1.7394)	
$\text{Temp}^2_{i,t}$	0.0011		0.0049	
	(0.0026)		(0.0031)	
$\text{INTL}_{ij}\times$ $\text{Temp}_{i,t}$	1.7192	0.5288	4.6851	4.8263**
	(2.0748)	(1.8177)	(2.9329)	(2.2375)
$\text{INTL}_{ij}\times$ $\text{Temp}^2_{i,t}$	-0.0030	-0.0009	-0.0083	-0.0085**
	(0.0036)	(0.0032)	(0.0051)	(0.0039)
$\text{Precip}_{i,t}$	-0.0098		-0.0042	
	(0.0136)		(0.0131)	
$\text{Precip}^2_{i,t}$	0.0006		0.0003	
	(0.0004)		(0.0004)	
$\text{INTL}_{ij}\times$ $\text{Precip}_{i,t}$	0.0135	0.0104	-0.0038	0.0017
	(0.0143)	(0.0118)	(0.0176)	(0.0193)
$\text{INTL}_{ij}\times$ $\text{Precip}^2_{i,t}$	-0.0011**	-0.0009**	-0.0001	-0.0003
	(0.0004)	(0.0004)	(0.0004)	(0.0006)
$\text{RTA}_{ij,t}$	0.1635***	0.1779***	0.0160	0.0275
	(0.0362)	(0.0421)	(0.0904)	(0.0607)
$\ln(\text{GDP}_{it})$	0.6030***		-0.0617	
	(0.0586)		(0.0939)	
Observations	109092	109092	109092	109092
Difference in marginal effect on exports relative to domestic sales:				
Temp at 15°C	-0.0140	-0.0037	-0.0813**	-0.0849*
	(0.0323)	(0.0290)	(0.0399)	(0.0437)
Temp at 25°C	-0.0741	-0.0222	-0.2467**	-0.2553***
	(0.1007)	(0.0891)	(0.1163)	(0.0855)
Precip at 8 m	-0.0042	-0.0037	-0.0058	-0.0031
	(0.0094)	(0.0095)	(0.0126)	(0.0140)
Precip at 20 m	-0.0308***	-0.0249*	-0.0089	-0.0103
	(0.0103)	(0.0145)	(0.0106)	(0.0180)
Exporter FE	✓		✓	
Exporter-Year FE		✓		✓

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

Notes: Standard errors (in parentheses) are clustered by exporter-importer pairs. All specifications control for exporter-importer FE, importer-year FE and INTL-year dummies.

from estimating the model separately for the manufacturing and agricultural sectors. In columns (1) and (2) the dependent variable is the nominal value of bilateral trade in manufacturing goods and in columns (3) and (4) the dependent variable is the nominal value of bilateral trade in agricultural goods. For each sector, I estimate specifications with the set of controls corresponding to those used in columns (2) and (3) of Table 1.2. As with Table 1.2, the bottom panel of Table 1.3 shows the estimated difference in the semi-elasticity for exports versus domestic sales.

The results for manufacturing trade (columns (1) and (2) of Table 1.3) paint a similar picture as the results for aggregate trade in Table 1.2, which is unsurprising given the large size of manufacturing relative to the agriculture sector. Overall, these results do not identify a significant effect of temperature shocks on manufacturing exports relative to domestic sales. In rainy places exports seem to be more sensitive to an increase in precipitation relative to domestic sales, but the standard errors on these estimates are large.[4] Of course, within the manufacturing sector is a wide array of industries, and underneath these results could be significant effects for particular sub-sectors. Dallmann (2019)'s results suggest that trade in some manufacturing sub-sectors is more sensitive than others to weather shocks. Future work could further investigate potential heterogeneity within the manufacturing sector.

The results for agricultural trade (columns (3) and (4) of Table 1.3) confirm previous findings in the literature that this sector is particularly sensitive to temperature shocks. As shown in Figures 1.6, in relatively hot places, a temperature shock leads to decreases in exports relative to domestic sales. These results suggest that previous studies of the effect of weather on agricultural trade are not simply identifying - through the lens of trade - the effect of temperature shocks on underlying production. Export sales seem to be particularly sensitive to weather shocks, confirming the suggestive evidence of this effect in Jones and Olken (2010)'s results. At 15°C, the semi-elasticity of agricultural export sales with respect to a 1°increase in annual mean temperature is about 8 percentage points lower than the semi-elasticity for domestic sales. At 25°C, this gap widens to about 25 percentage points. Meanwhile, as illustrated by Figure 1.7, the marginal effect of precipitation on exports relative to domestic sales is statistically insignificant across the distribution of the total annual precipitation variable. Together with the results for the manufacturing sector, this result suggests that the significant effect of precipitation shocks on aggregate trade found in the previous section is driven solely by the manufacturing sector.

The estimates for the differential marginal effect of temperature on agricultural ex-

[4] See the appendix of this chapter for figures illustrating the estimated semi-elasticities for the manufacturing sector.

Figure 1.6: Estimated marginal effect of temperature on agriculture trade

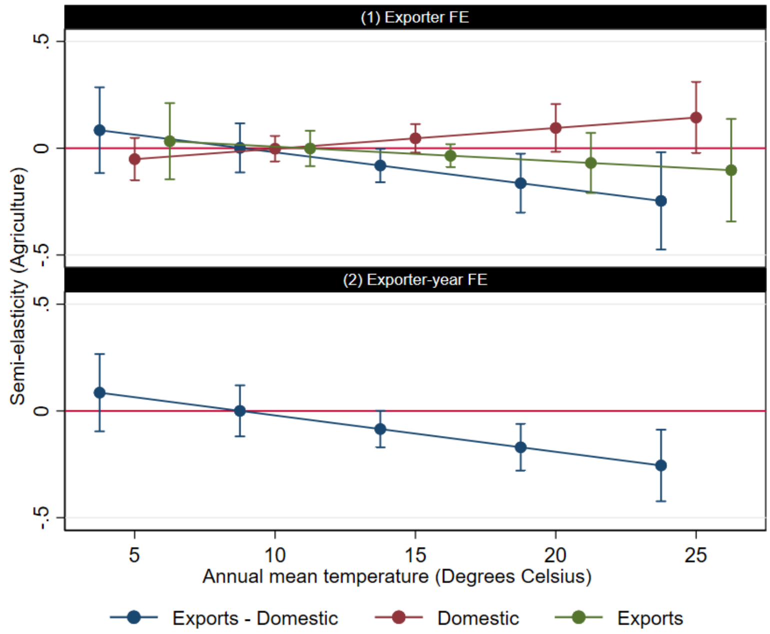

Notes: Marginal effect estimates shown here correspond to the estimates in columns (3) and (4) of Table 1.3. 'Exports - Domestic' (the blue line) denotes the difference in the semi-elasticity for exports versus domestic trade - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

Figure 1.7: Estimated marginal effect of precipitation on agriculture trade

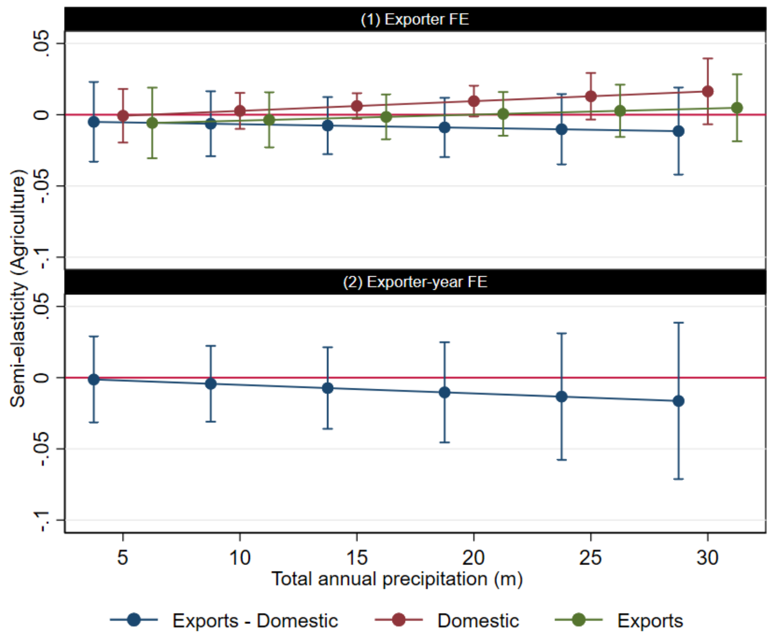

Notes: Marginal effect estimates shown here correspond to the estimates in columns (3) and (4) of Table 1.3. 'Exports - Domestic' (the blue line) denotes the difference in the semi-elasticity for exports versus domestic trade - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

ports compared to domestic sales are quite large, particularly at the upper end of the temperature distribution. Importantly, the estimated marginal effect of 0.25 at 25°C does not necessarily imply a 25% decrease in exports, but a 25 percentage point *difference* in the marginal effect on domestic sales versus exports. Exports likely do not decrease by as much as 25%, but instead domestic sales increase and exports decrease and the combined effect is a 25 percentage point gap between these two marginal effects. Indeed, the results in column (3) of Table 1.3 imply that the value of domestic sales may increase in response a temperature shock while the value of exports decreases, implying that the marginal effect of a temperature shock at 25°C may not be as large as a 25% decrease in exports. Moreover, the dependent variable is the nominal value of trade, so mechanisms for this relative sensitivity of exports could be channeled through both quantities and prices. That is, the decrease in exports relative to domestic sales could reflect both a decrease in underlying quantities exported relative to domestic quantities sold, as well as a decrease in the relative price of exports compared to domestically-sold goods. If both a price and quantity mechanism are at work here, then the two effects would reinforce each other and exacerbate the magnitude of the estimated total effect.

In terms of potential mechanisms channeled through quantities sold, as hypothesized in the Introduction of this chapter, producers may have a 'home bias' and prioritise their domestic markets, so when a weather shock negatively impacts production, they decrease exports more than domestic sales quantities. Another possibility is that international supply chains are more sensitive than within-country supply chains. In effect, weather shocks may act as an additional barrier to trade that increases the difficulty of getting goods to export markets. Meanwhile, potential mechanisms via a decrease in the relative price of exports compared to the price of goods sold domestically imply that international trade barriers insulate producers from global competition and permit price fluctuations in domestic markets that do not occur in export markets. The results here are consistent with a potential mechanism in which agricultural producers are able to increase prices in their domestic market following a contraction in supply due to a temperature shock, but export prices remain steady, so the value of exports decreases relative to the value of domestic sales. A similar mechanism may occur through long-term export contracts, if these contracts tie producers to price levels before production quantities are realised. Overall, the large magnitude of the estimate for relative effect of temperature shocks on exports versus domestic sales suggests that several mechanisms may be at play here.

Another important point regarding the magnitude of these estimates relates to the body of work studying the "border puzzle" in empirical models of bilateral trade.

This literature has often found that international borders have a very large negative impact on the flow of trade. Early estimates suggested that international borders decreased the flow of trade by thousands of percent compared to domestic trade (Bergstrand et al. 2015). Modern improvements in the estimation of gravity equations have tempered these estimates; nevertheless, in a recent analysis by Yotov et al. (2016) the estimated partial equilibrium effect of the $INTL_{ij}$ border dummy variable on bilateral trade is -91.6%. The effect of interest in this paper - the interactive term $INTL_{ij} \times h(T_{it})$ - can be interpreted as the effect of temperature shocks on the international border effect. From this perspective, the magnitudes of the estimates in this paper are not unprecedented, since we know from previous empirical literature that international borders have a large negative impact on trade. The results in this paper imply that these border effects in the agriculture sector fluctuate quite a bit with temperature shocks. Overall, the large magnitude of the effect of temperatures shocks on agricultural exports relative to domestic sales may be reasonable when one considers previous estimates of border effects as well as the myriad of potential mechanisms that this estimate may capture, but continued research to better understand the potential impacts of weather and climate change on international trade would be beneficial to better understand the extent to which these magnitudes are realistic.

Regardless of the exact mechanisms driving the relative sensitivity of agricultural exports to weather shocks, we can expect that this effect may have implications for economic welfare. From a utilitarian perspective on social welfare, which underlies neoclassical models of international trade, the welfare impact of these results is the same regardless of whether the mechanism occurs through prices or quantities. Either way, the weather shock decreases international bilateral accessibility; that is, openness to international trade decreases and goods become less accessible (either in price or quantity terms) to consumers around the world. On aggregate, gains to consumers from trade openness outweigh potential losses to producers that cannot remain competitive, and so a weather shock that decreases trade openness leads to a decrease in aggregate economic welfare. On the other hand, from a distributional perspective, the implications of whether this effect occurs through a price or quantity mechanism may matter for welfare. The estimated negative effect of temperature shocks on agricultural exports relative to domestic sales is consistent with both (i) export quantities sold decrease relative to domestic quantities sold, and (ii) the price of export sales decreases relative to domestic prices. Mechanism (i) may be beneficial for domestic consumers, and in particular could help to maintain domestic food security in the event of a temperature shock. However, mechanism (ii) could undermine domestic food security, but could help to shield domestic producers from the negative productivity impacts of temperature shocks. In summary, from the perspective of Pareto efficiency, regardless of the mechanism underlying these estimated effects the welfare implication is that temperature shocks decrease welfare via a decrease in trade openness. However, from a distributional perspective, mechanisms via prices may be beneficial for local producers but harmful for local consumers.

1.4.3 Heterogeneity by exporter characteristics

Table 1.4 explores these results for bilateral agricultural trade based on characteristics of the exporting country. Column (1) interacts the temperature-trade function with a dummy variable indicating whether the exporter is a low income country, and column (2) interacts this function with a dummy variable indicating whether the exporter has weak institutions (see section 1.3 for a detailed description of these variables). The set of controls used in both specifications corresponds to those in columns (2) and (4) of Table 1.3: exporter-importer FE, importer-year and exporter-year FE, and INTL-year dummies. The bottom panel of Table 1.4 shows the estimated difference in the marginal effects for exports versus domestic sales according to these exporter characteristics, and Figures 1.8 and 1.9 illustrate these estimates across the distributions of the temperature and precipitation variables.

Table 1.4: Exploring heterogeneity by country characteristics

	(1)	(2)
$\text{INTL}_{ij} \times \text{Temp}_{i,t}$	3.5346^{*}	5.2463^{***}
	(1.9685)	(1.7744)
$\text{INTL}_{ij} \times \text{Temp}^2_{i,t}$	-0.0063^{*}	-0.0093^{***}
	(0.0034)	(0.0031)
$\text{Low income}_{i,t} \times \text{INTL}_{ij} \times \text{Temp}_{i,t}$	-0.0751^{***}	
	(0.0186)	
$\text{Low income}_{i,t} \times \text{INTL}_{ij} \times \text{Temp}^2_{i,t}$	0.0003^{***}	
	(0.0001)	
$\text{Weak institutions}_{i,t} \times \text{INTL}_{ij} \times \text{Temp}_{i,t}$		-0.0418^{**}
		(0.0203)
$\text{Weak institutions}_{i,t} \times \text{INTL}_{ij} \times \text{Temp}^2_{i,t}$		0.0002^{**}
		(0.0001)
$\text{INTL}_{ij} \times \text{Precip}_{i,t}$	0.0503^{**}	0.0605

	(0.0253)	(0.0376)
$\text{INTL}_{ij} \times \text{Precip}^2_{i,t}$	-0.0028***	-0.0021*
	(0.0009)	(0.0012)
$\text{Low income}_{i,t} \times \text{INTL}_{ij} \times \text{Precip}_{i,t}$	-0.1110***	
	(0.0414)	
$\text{Low income}_{i,t} \times \text{INTL}_{ij} \times \text{Precip}^2_{i,t}$	0.0041***	
	(0.0011)	
$\text{Weak institutions}_{i,t} \times \text{INTL}_{ij} \times \text{Precip}_{i,t}$		-0.1149**
		(0.0544)
$\text{Weak institutions}_{i,t} \times \text{INTL}_{ij} \times \text{Precip}^2_{i,t}$		0.0031**
		(0.0015)
$\text{RTA}_{ij,t}$	0.0334	-0.0771
	(0.0591)	(0.0487)
Observations	109092	78584
Difference in marginal effect on exports relative to domestic sales:		
High income		
Temp at 15°C	-0.0835***	
	(0.0275)	
Temp at 25°C	-0.2091***	
	(0.0705)	
Precip at 8 m	0.0062	
	(0.0148)	
Precip at 20 m	-0.0601***	
	(0.0200)	
Low income		
Temp at 15°C	-0.0081	
	(0.0305)	
Temp at 25°C	-0.1284*	
	(0.0675)	
Precip at 8 m	-0.0400	
	(0.0256)	
Precip at 20 m	-0.0090	
	(0.0146)	
Strong institutions		
Temp at 15°C		-0.1215***
		(0.0405)

Temp at 25°C	-0.3078***
	(0.0762)
Precip at 8 m	0.0261
	(0.0216)
Precip at 20 m	-0.0254
	(0.0225)
Weak institutions	
Temp at 15°C	-0.0760**
	(0.0386)
Temp at 25°C	-0.2593***
	(0.0621)
Precip at 8 m	-0.0387*
	(0.0224)
Precip at 20 m	-0.0152
	(0.0121)

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

Notes: The dependent variable in both specifications is the value of agricultural trade. Standard errors (in parentheses) are clustered by exporter-importer pairs. All specifications control for exporter-importer FE, importer-year FE, exporter-year FE, and INTL-year dummies.

As shown in Column (1) of Table 1.4, temperature shocks lead to a shift in the balance of trade away from exports and towards domestic market sales for both high and low income countries. However, this effect may be slightly stronger and emerge at lower temperatures for high income compared to low income countries. At 25°C, the estimated marginal effect of a temperature shock on exports in high income countries is 20.9 percentage points lower than the effect on domestic sales; in low income countries this gap is 12.8 percentage points. The estimates in column (1) show even stronger heterogeneity between high and low income countries for the effect of precipitation shocks on exports relative to domestic sales: a statistically significant difference in the marginal effect on exports relative to domestic sales only emerges for high income countries. This result implies that in high income countries that are already fairly rainy, exports are more sensitive to an increase in precipitation than domestic sales, while in low income countries an increase in precipitation has a similar impact on exports as on domestic sales.

Several possible factors may explain this heterogeneity in the estimates for high versus low income exporters. Perhaps high versus low income countries tend to specialize in different agricultural sub-sectors, which leads to differing vulnerability of

Figure 1.8: Heterogeneity in the marginal effect of temperature on agricultural trade

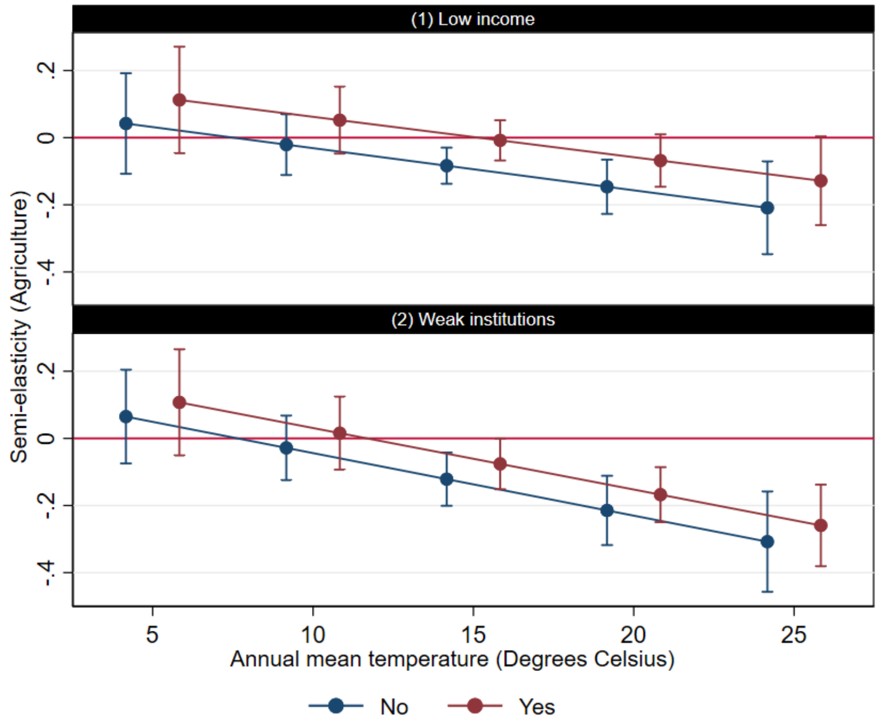

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.4. The blue line, 'No', denotes the difference in the marginal effect for exports versus domestic sales for countries without the exporter characteristic in the plot title, and the red line, 'Yes', shows this effect for countries with this characteristic.

trade balances to weather shocks. Given that low income countries tend to be hotter (and to some extent wetter) than high income countries in the data, the weaker effect in these places may indicate some long term adaptation to typical weather conditions. Moreover, the fact that statistically significant effects emerge at lower levels of temperature and precipitation for high versus low income countries may indicate a lack of statistical power to identify the marginal effects outside the usual temperature ranges for these groups. Finally, as shown in Table 1.1, less than a third of the sample of country-year observations are in the low income category, so this relatively low coverage of low income countries in the sample data may partly explain the lack of precise estimates for this group. Unfortunately this limitation is inherent in the data requirements of this research design, because data on the gross value of production (which is necessary to construct observations of domestic sales) is less available for low income compared to high income countries.

Column (2) of Table 1.4 allows the effect of weather shocks on exports relative to domestic sales to vary based on whether the exporting country has strong versus

Figure 1.9: Heterogeneity in the marginal effect of precipitation on agricultural trade

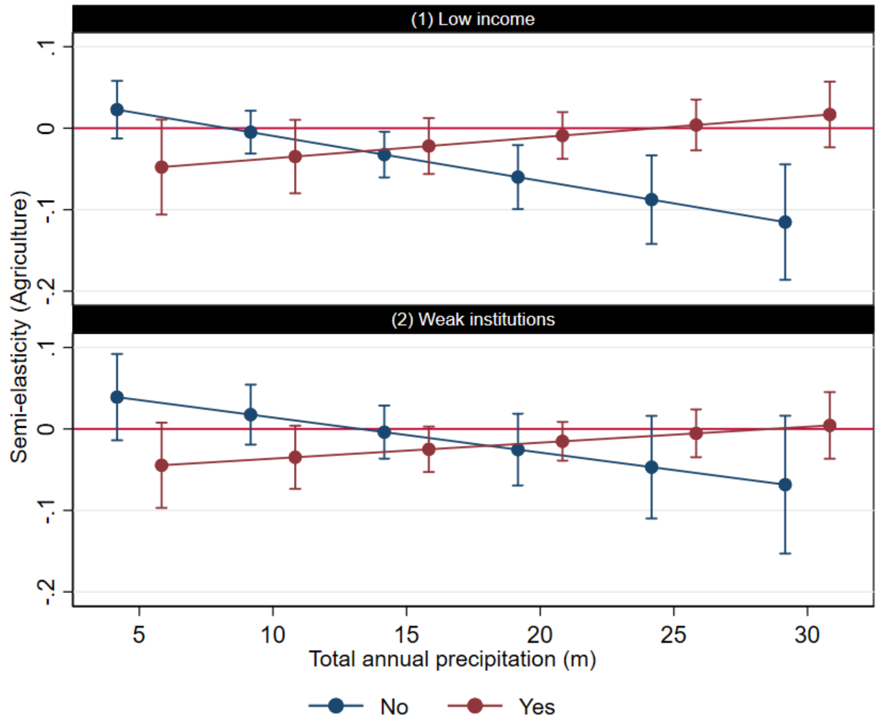

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.4. The blue line, 'No', denotes the difference in the marginal effect for exports versus domestic trade for countries without the exporter characteristic in the plot title, and the red line, 'Yes', shows this effect for countries with this characteristic.

weak institutional quality relative to the rest of the countries in the sample. The results imply that at a marginal increase in temperature negatively impacts exports relative to domestic sales regardless of the institutional quality of the exporting country. This effect may be slightly stronger in countries with strong institutions, but the standard errors on these estimates are large so this heterogeneity is not precisely identified.

On the other hand, institutions may be a source of heterogeneity in the effect of precipitation on trade balances. At 8 metres total annual precipitation, countries with weak institutions see a marginal effect of precipitation on exports that is 3.9 percentage points lower than the marginal effect on domestic sales, while countries with strong institutions see no statistically significant difference in the marginal effect of precipitation on exports relative to domestic sales. Beverelli et al. (2018) find that weak institutions hinder exports relative to domestic sales. The results here suggest that precipitation shocks may exacerbate this negative impact of weak institutions on exports relative to domestic sales.

Overall, strong institutions may be an important mitigating factor in the vulnerability of exports to precipitation shocks, but otherwise low income levels and weak institutions do not seem to increase the vulnerability of countries' trade balances to weather shocks. Nevertheless, some heterogeneity based on income level and institutional quality of an exporting country may exist in the relative marginal effect curve. As mentioned above, this heterogeneity may reflect systematic differences in specialization in agricultural sub-sectors and in long-term adaptation. Future work on this topic should aim to understand this heterogeneity more thoroughly, though data limitations may present a challenge here.

1.4.4 Alternative functional forms for temperature

As a robustness check for the quadratic functional form assumption for the weather and trade relationships, Table 1.5 shows coefficient estimates for specifications based on alternative functional forms for the effects of temperature and precipitation on exports relative to domestic sales. Columns (1) and (2) introduce third-order and fourth-order polynomial terms, respectively, and columns (3) to (5) use flexible functional forms based on number of days in a year in a given temperature range. The dependent variable in all columns is the value of bilateral trade in agricultural products. All columns use the same set of controls as in column (5) of Table 1.3: exporter-importer FE, importer-year FE, exporter-year FE, and INTL-year dummies.

Table 1.5: Alternative functional forms

	(1)	(2)	(3)	(4)	(5)
$\text{INTL}_{ij}\times$ $\text{Temp}_{i,t}$	9.5303				
	(73.2664)				
$\text{INTL}_{ij}\times$ $\text{Temp}^2_{i,t}$	-0.0249	0.0247			
	(0.2538)	(0.1302)			
$\text{INTL}_{ij}\times$ $\text{Temp}^3_{i,t}$	0.0000	-0.0001			
	(0.0003)	(0.0006)			
$\text{INTL}_{ij}\times$ $\text{Temp}^4_{i,t}$		0.0000			
		(0.0000)			
$\text{INTL}_{ij}\times$ $\text{Precip}_{i,t}$	0.0479	0.0437	-0.0035	-0.0018	0.0154
	(0.0478)	(0.0608)	(0.0211)	(0.0192)	(0.0198)
$\text{INTL}_{ij}\times$ $\text{Precip}^2_{i,t}$	-0.0036	-0.0032	-0.0000	-0.0001	-0.0005

	(0.0030)	(0.0043)	(0.0005)	(0.0006)	(0.0005)
INTL$_{ij}\times$ Precip$^3_{i,t}$	0.0001	0.0000			
	(0.0000)	(0.0001)			
INTL$_{ij}\times$ Precip$^4_{i,t}$		0.0000			
		(0.0000)			
$D_{it,b}\times$ INTL$_{ij}, b \in (-\infty, -5]$°C			0.0043		
			(0.0046)		
$D_{it,b}\times$ INTL$_{ij}, b \in (-5, 0]$°C			-0.0101**		
			(0.0047)		
$D_{it,b}\times$ INTL$_{ij}, b \in (0, 5]$°C			0.0025		
			(0.0025)		
$D_{it,b}\times$ INTL$_{ij}, b \in (5, 10]$°C			0.0020		
			(0.0032)		
$D_{it,b}\times$ INTL$_{ij}, b \in (15, 20]$°C			-0.0057**		
			(0.0026)		
$D_{it,b}\times$ INTL$_{ij}, b \in (20, 25]$°C			0.0007		
			(0.0028)		
$D_{it,b}\times$ INTL$_{ij}, b \in (25, 30]$°C			-0.0059**		
			(0.0029)		
$D_{it,b}\times$ INTL$_{ij}, b \in (30, \infty)$°C			-0.0005		
			(0.0023)		
$D_{it,b}\times$ INTL$_{ij}, b \in (15, 25]$°C				-0.0025	
				(0.0021)	
$D_{it,b}\times$ INTL$_{ij}, b \in (25, \infty]$°C				-0.0052**	
				(0.0025)	
$D_{it,b}\times$ INTL$_{ij}, b \in (20, \infty]$°C					0.0014
					(0.0021)
RTA$_{ij,t}$	0.0274	0.0276	0.0161	0.0194	0.0212
	(0.0610)	(0.0610)	(0.0593)	(0.0608)	(0.0616)
Observations	109092	109092	109092	109092	109092

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

Notes: The dependent variable in all specifications is the value of agricultural trade. Standard errors (in parentheses) are clustered by exporter-importer pairs. All specifications control for exporter-importer FE, importer-year FE, exporter-year FE, and INTL-year dummies.

In column (3), the effect of temperature on exports relative to domestic sales is given by $h(D_{it}) = INTL_{ij} \times \sum_{b=0}^{B} \theta_b D_{it,b}$. The temperature distribution is divided into B 5 °C bins, and $D_{it,b}$ is the number of days during year t in which the mean temperature in country i falls into bin b. The coefficient estimate $\hat{\theta}_b$ indicates the marginal effect (relative to the reference bin of $(10, 15]$°C) of an additional day in bin b on the semi-elasticity of exports relative to domestic sales. The results are roughly consistent with a quadratic functional form for the effect of temperature, with stronger effects of temperature shocks on exports relative to domestic sales occurring at more extreme temperatures. This specification greatly reduces functional form assumptions compared to the polynomial specifications, but it has the disadvantage of requiring more statistical power to estimate all of the coefficients, and the prevalence of statistically insignificant results in column (3) may reflect a lack of power. I address this issue by estimating specifications with fewer bins in columns (4) and (5). Given that the results above suggest that the effect is occurs mainly in hot places, I focus on bins in the upper ranges of the temperature distribution. The results in these columns confirm that high temperature levels, particularly above 25°C, are a key driver of the relative sensitivity of exports to temperature shocks.

1.5 Counterfactual welfare simulation

As noted in Section 1.2, the theoretical expression for bilateral trade flows given by Equation 1.1 can be derived from a variety of micro-foundations. Arkolakis et al. (2012)'s seminal contribution demonstrates that this wide class of quantitative trade models share a common sufficient statistic for changes in welfare:

$$\Delta W_i = \left(\frac{\lambda_{ii}^{CFL}}{\lambda_{ii}^{BLN}}\right)^{\frac{1}{1-\sigma}}$$

ΔW_i is the change in welfare for country i moving from the baseline (BLN) to counterfactual (CFL) scenario. $\lambda_{ii} = X_{ii}/E_i$ is country i's expenditure on domestically-produced goods as share of total expenditure. $\sigma > 1$ is the trade elasticity of substitution, and $1 - \sigma$ is the elasticity of trade flows with respect to trade costs. In many demand-side derivations of the structural gravity model, a key assumption is that consumers prefer variety and each country produces a differentiated variety of the traded good, and the σ parameter, often called the Armington elasticity, is the representative consumer's elasticity of substitution between different varieties (Head and Mayer 2014). In supply-side derivations of the structural gravity model such as the Eaton and Kortum (2002) model, $1 - \sigma$ reflects the degree of variation in the productivity of firms around the world. Regardless of the micro-foundations that underpin its interpretation, σ is a parameter from a static trade model and is distinct from the elasticity of marginal utility in the climate economics literature, which is sometimes denoted as σ, and which reflects social preferences for smooth-

Figure 1.10: Returning to 1980s weather: Estimated impacts on welfare via changes in trade flows, $\sigma = 6$

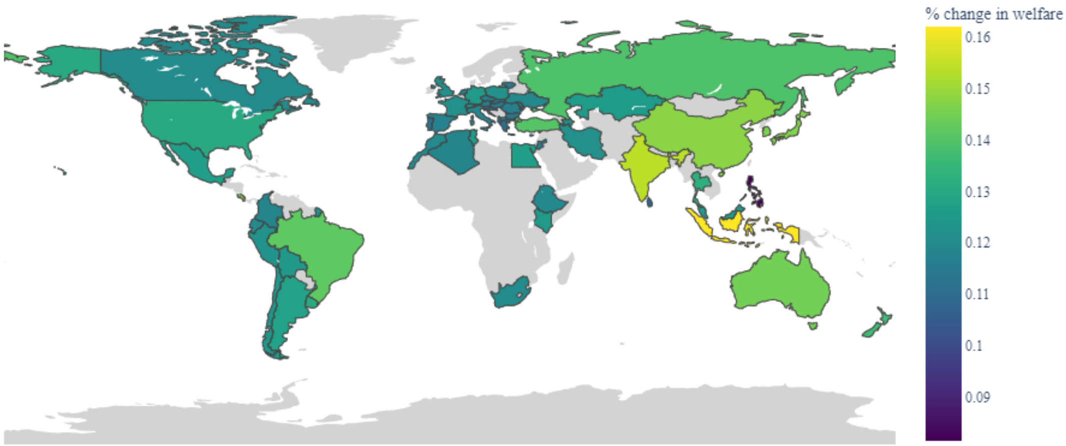

ing consumption levels across generations (Ramsey 1928; Drupp et al. 2018). With this in mind, the above sufficient statistic for welfare does not take into account the dynamics of decision making, in particular preferences for social discounting, and the results presented here should be interpreted with this caveat in mind.

To assess the potential welfare implications of the differential impact of temperature on exports versus domestic market sales, I use this sufficient statistic alongside the coefficient estimates in column (3) of Table 1.2. First, I plug these coefficient estimates into equation 1.3 to predict bilateral trade ($\hat{X}_{ij}$) in a baseline and counterfactual scenario: the baseline scenario uses the average of the temperature and precipitation variables observed in each country from 2008 to 2018, and the counterfactual scenario takes these averages across 1980 to 1989. The median change in temperature observed in sample countries between these two decades is 0.92 degrees, and the median precipitation change is 0.23 metres. All other variables are constant across the baseline and counterfactual scenarios, and I use observed trade agreements as well as fixed effects estimates for 2017.

Next, I compute predicted total expenditure by summing up predicted bilateral trade across all exporters, $\hat{E}_{j,BLN} = \sum_i \hat{X}_{ij,BLN}$, and then I calculate the predicted share of expenditure on domestic goods: $\hat{\lambda}_{ii,BLN} = \hat{X}_{ii,BLN}/\hat{E}_{i,BLN}$. I do these calculations for both the baseline and counterfactual scenarios. Finally, I plug $\hat{\lambda}_{ii,BLN}$ and $\hat{\lambda}_{ii,CFL}$ into the equation above to compute $\Delta\hat{W}_i$ for each country. I use a value of 6 for σ, which is Head and Mayer (2014)'s preferred estimate after reviewing the literature on this parameter. As a robustness check, I also do this analysis with a value of 4 for σ, which is close to the median of the full sample of estimates in included in Head and Mayer (2014)'s meta-analysis of estimates of the trade elasticity.[5]

This simulation assesses the welfare impacts (via changes in trade flows) of returning to weather conditions of the 1980s but keeping international trade costs, policies, institutions, and all other bilateral and country-specific factors at 2017 conditions. Figure 1.10 illustrates the results of this simulation. Returning to weather conditions of the 1980s leads to a 0.08% to 0.16% increase in welfare across the sample countries. Note that this impact of weather changes on welfare reflects only changes due adjustment in trade flows; other impacts of changes in temperature and precipitation across this time period, such as health and productivity impacts, are not reflected in this simulation. These decreases in welfare due to temperature increases of around 1°since the 1980s are not huge, but given the underlying size of the economies in the sample, they are not negligible either. Comparing Figure 1.10 with Figure 1.13 (in the appendix to this chapter) reveals that the calibration of the σ parameter affects the magnitudes of these results. The calibration of σ at the alternative value of 4 reflects an assumption that trade is less elastic to trade costs, so trade becomes more important to welfare compared to the calibration in Figure 1.10. Compared to the main results that assume $\sigma = 6$, in the alternative results for $\sigma = 4$ the magnitudes of the estimated trade-related welfare impacts of climate change increase, ranging from 0.13% to 0.27%. Nevertheless, the heterogeneity in these trade-related welfare impacts across countries is quite similar regardless of the calibration of σ. In general, countries in the north western quadrant of the map in Figure 1.10 tend to have seen relatively small trade-related impacts of climate change since the 1980s, while countries in the south eastern quadrant have seen relatively large impacts.

These welfare impacts of weather changes via changes in exports relative to domestic trade help to illustrate the potential economic significance of the empirical estimates presented above. Nevertheless, this simulation should be interpreted with a couple of important caveats in mind. First, the coefficient estimates from the empirical model are identified from weather shocks - that is, deviations from the usual temperature and precipitation experienced in a given country. The response of trade flows to a temperature shock may be different than the response to changes in temperature over 30 years because a gradual long-term change offers opportunities for exporters to learn and adapt. Accordingly, these simulated welfare impacts of changes in weather may overstate the impact of climate change. Finally, this sufficient statistic for welfare is derived from a stylized static framework for international trade which

[5]See Figure 1.13 in the appendix of this chapter for an illustration of the results with this alternative value for σ.

leads to a continuous monotonic relationship between changes in trade openness and changes in economic welfare; a more complex social welfare function that allows for dynamics, for example, or takes more account of distributional impacts within a country, may be more appropriate for a full assessment of the welfare implications of the results in this study.

1.6 Conclusion

This paper uses an approach that brings together developments from the international trade and climate change economics literatures to investigate the differential impact of weather shocks on exports relative to domestic sales. In contrast to previous empirical papers that study the impact of weather shocks on international trade, I include domestic trade flows in my model and control robustly for multilateral resistance using exporter-year and importer-year fixed effects. First I demonstrate that including this robust set of controls leads to somewhat different results compared to other papers that study this topic; in particular, the results suggest that manufacturing trade is not affected by temperature shocks. Manufacturing exports in very rainy places decrease relative to domestic sales in response to an increase in precipitation. Meanwhile, in the agricultural sector increased temperatures leads to a decrease in exports relative to domestic sales, mostly in already hot countries. I do not find strong evidence of heterogeneity in this effect based on income levels or institutional quality. Finally, a simple sufficient statistics analysis to assess the welfare implications of these results suggests that returning to temperature and precipitation levels of the 1980s but otherwise keeping the global trading network constant at recent conditions would increase welfare by 0.08% to 0.16% in sample countries due to changes in exports relative to domestic sales.

This paper contributes to our understanding of how to conceptualize the economic impacts of climate and weather. Climate change economists often model the economic damages associated with increased temperatures as part of the production function, which implies that these damages are productivity impacts. This paper brings some empirical insight into how this assumption might be a simplification. More precisely, the economic damages of weather shocks likely do not stop when agricultural goods leave the farm, but continue to have impacts along the supply chain. The results of this paper suggest that weather shocks are an additional barrier to international trade, or perhaps exacerbate existing barriers to international trade.

Understanding the mechanisms underlying this effect as well as its economic significance are two key areas for future research on this topic. Several potential underlying mechanisms are consistent with these results, including that producers have a ‘home bias’, and that weather shocks create an additional barrier to international trade, perhaps through difficulties in transporting goods internationally or by increasing the gap between domestic and export prices. Future work could break the analysis down into more granular sub-sectors to try to get a better idea of what is driving these results. Moreover, given that these effects are identified from weather shocks, the economic significance of these results for long-term climate change is unclear. Understanding the mechanisms behind the effect will help to understand the extent to which exporters can adapt to these impacts in the long term. Alternative models and methods may also be useful for understanding how climate change may have economically significant impacts via impacts on trade flows.

Finally, some policy takeaways arise from this paper. The results confirm findings in many other papers that the agricultural sector is particularly sensitive to weather shocks, and so climate and trade policy should take into account these sector-specific vulnerabilities. In particular, the results stress the importance of policy alignment. Climate and trade interact with each other in their effects on economic welfare, and so climate and trade policy should not exist in silos but instead take into account these interactions. For example, policy initiatives to support trade openness and export-driven growth could benefit from including climate change adaptation measures.

1.A Appendix to Chapter 1

1.A.1 List of countries in the model

ISO3 Codes in parentheses: Albania (ALB), Algeria (DZA), Argentina (ARG), Australia (AUS), Austria (AUT), Azerbaijan (AZE), Bolivia (Plurinational State of) (BOL), Brazil (BRA), Bulgaria (BGR), Canada (CAN), Chile (CHL), China (CHN), Colombia (COL), Costa Rica (CRI), Croatia (HRV), Cyprus (CYP), Czechia (CZE), Czechoslovakia (CZE), Ecuador (ECU), Egypt (EGY), Ethiopia (ETH), France (FRA), Germany (DEU), Greece (GRC), Hungary (HUN), India (IND), Indonesia (IDN), Iran (Islamic Republic of) (IRN), Israel (ISR), Italy (ITA), Japan (JPN), Jordan (JOR), Kazakhstan (KAZ), Kenya (KEN), Kyrgyzstan (KGZ), Lebanon (LBN), Lithuania (LTU), Madagascar (MDG), Malaysia (MYS), Mexico (MEX), Morocco (MAR), Netherlands (NLD), New Zealand (NZL), North Macedonia (MKD), Pakistan (PAK), Peru (PER), Philippines (PHL), Poland (POL), Portugal (PRT), Republic of Korea (KOR), Republic of Moldova (MDA), Romania (ROU), Russian Federation (RUS), Slovakia (SVK), Slovenia (SVN), South Africa (ZAF), Spain (ESP), Sri Lanka (LKA), State of Palestine (PSE), Switzerland (CHE), Thailand

(THA), Tunisia (TUN), Turkey (TUR), Ukraine (UKR), United Kingdom (GBR), United States of America (USA), Uruguay (URY), Yemen (YEM)

1.A.2 Additional notes on data cleaning

For the manufacturing sector, I calculate domestic trade at the country-year-sector level (I define country-specific sectors based on combinations of ISIC codes a given country reports production under over the sample period). Two potential issues lead to missing values for domestic trade: (i) production data is missing, and (ii) the production value is less than the value of exports and therefore domestic trade is negative, in which case I recode domestic trade to missing. I am unable to observe when missing production values indicate true missing values versus zero actual production. To avoid creating mechanical year-to-year variation in domestic trade based on what sectors are observed in a given country, after aggregating across ISIC codes to the country-year level, I recode domestic trade as missing if the number of sectors observed for a given country-year is more than 2 different than the mode number of sectors observed for that country.

I check for overlaps in the product coverage across FAOSTAT and Unido. I drop 254 such overlapping FAO items from the FAOSTAT data and instead use the observations in the Unido and Comtrade data and allocate these products to the manufacturing sector. I also drop FAO items that cannot be matched to an ISIC or HS code and items for which production data is unavailable despite the availability of trade data (e.g. live animals).

I drop countries from the sample if they are not in both UNIDO and FAO production data. To limit the prevalence of missing observations for domestic trade I drop addditional countries and settle on the sample of 67 countries listed above.

1.A.3 Additional figures

Figure 1.11: Estimated marginal effects of temperature on manufacturing trade

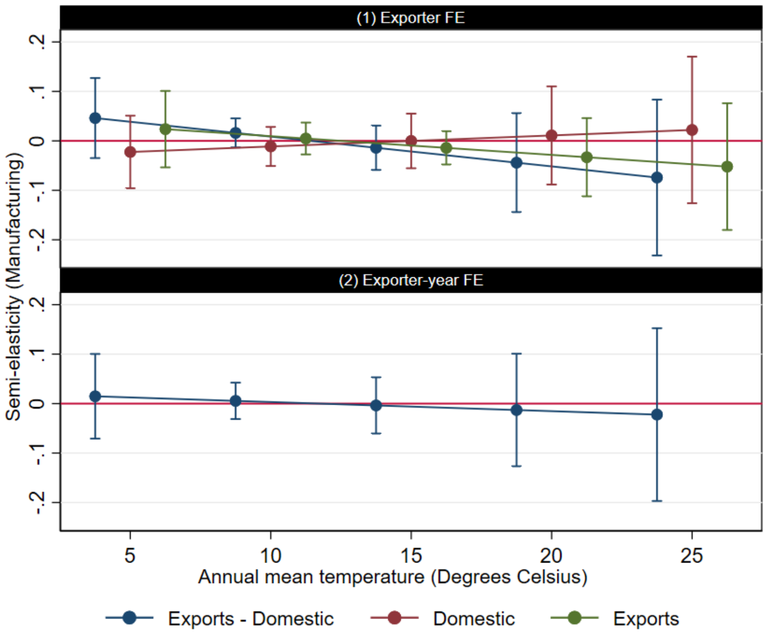

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.3 - i.e. the plot in first row, titled (1), corresponds to estimates from column (1) in Table 1.3. 'Exports - Domestic' (the blue line) denotes the difference in the marginal effect for exports versus domestic trade - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

Figure 1.12: Estimated marginal effect of precipitation on manufacturing trade

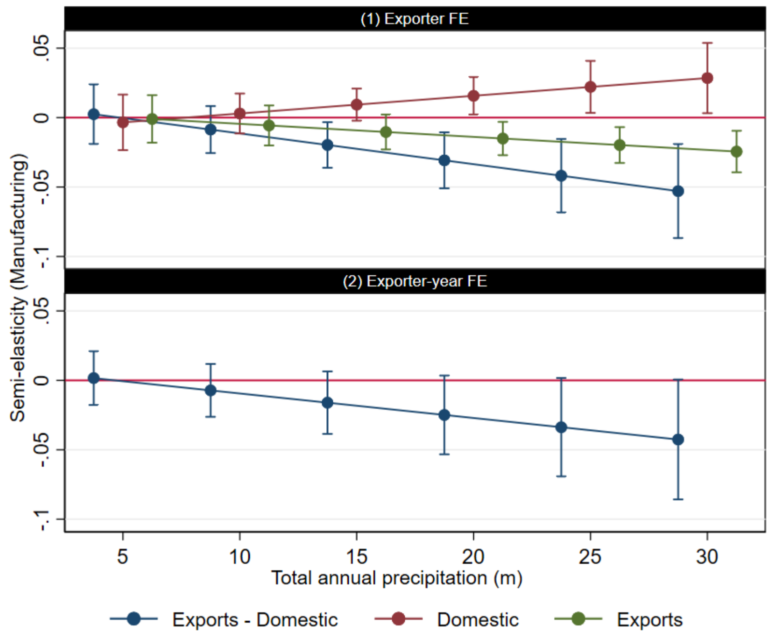

Notes: Marginal effect estimates shown here correspond to the estimates in the column of the same number in Table 1.3 - i.e. the plot in first row, titled (1), corresponds to estimates from column (1) in Table 1.3. ‘Exports - Domestic’ (the blue line) denotes the difference in the marginal effect for exports versus domestic trade - it is the main effect of interest in this study, telling us if weather shocks affect exports differently from domestic sales.

Figure 1.13: Returning to 1980s weather: Estimated impacts on welfare via changes in trade flows, $\sigma = 4$

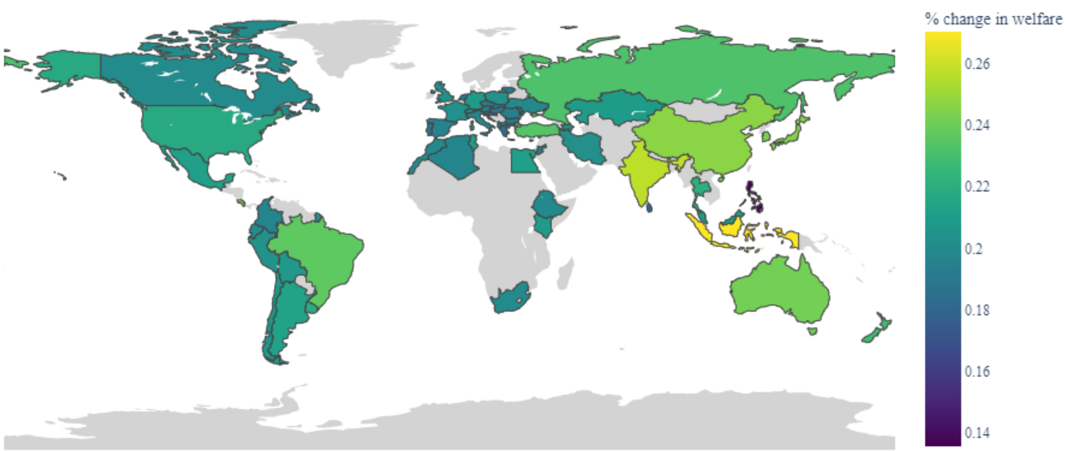

Chapter 2

The role of trade openness in the temperature-growth relationship

2.1 Introduction

In recent years, empirical studies from the climate change economics literature have substantially increased our understanding of how weather and climate change can impact economic outcomes. However, many of the estimates in these studies rely on ceteris parabis assumptions that hold international trade relationships constant. In the context of our globalized, interconnected world, this assumption is obviously one of convenience rather than realism. Indeed, a common hypothesis is that openness to international trade can help to moderate the impact of climate change on economic growth. The idea behind this hypothesis is that climate change will cause shifts in the suitability of different locations for different types of production, and free trade allows for adjustment in the spatial distribution of production in line with these changes in comparative advantage. So far the evidence to support this hypothesis relies heavily on structural modelling techniques, and we lack a strong empirical basis for this theory.

This paper contributes empirical evidence on the role of trade openness in the effect of temperature on economic growth. I use a reduced-form empirical model that brings together developments from the international trade and climate change economics literatures to deliver an estimate of the effect of trade openness on the historical relationship between temperature shocks and GDP per capita growth. I find a negative effect of temperature shocks on aggregate income growth that is mainly present in hot *and* relatively remote countries, which supports the hypothesis that openness to international trade may help to moderate the negative effect of temperature shocks on income growth. However, when I estimate the effects separately for agricultural and manufacturing income growth, the picture is less clear. In particular, I do not find evidence that trade openness lessens the negative impact of temperature shocks on agricultural income growth in hot countries, which challenges the idea that trade openness can be an important mode of adaptation to climate change given that these places are particularly vulnerable to climate change. Overall, this paper makes an important contribution towards an empirical understanding of the role that trade openness might play in the impact of weather and climate on economic outcomes.

This paper is part of a growing literature that uses reduced-form empirical models to estimate the impact of increased temperature on economic outcomes. This body of work provides an evidence base for the structural relationships between climate and the economy in integrated assessment models of climate change, and is also an important source of evidence for climate policy in its own right. In a seminal contribution, Dell et al. (2012) find that increased temperatures have historically led to decreased growth, spurring a large body of work using similar reduced-form panel data methods to assess the economic impacts of weather shocks and climate change. Important contributions such as Burke et al. (2015) and Kalkuhl and Wenz (2020) have confirmed the impact of weather shocks on growth and have illuminated the importance of functional form assumptions such as non-linearities in the temperature-growth relationship. Overall, research over the last decade has built a strong basis of empirical evidence that increased temperatures can lead to decreased economic growth.

Another potentially important channel for economic growth is international trade. From longstanding arguments based on specialization and comparative advantage, to more recent arguments based on increasing returns to scale, economic theory has a long tradition of expounding the benefits of international trade (Krugman 1987). Close links to international markets allows firms to exploit their comparative advantage and reduce costs through increasing returns scale, and so not only does firm productivity increase, but consumers enjoy a cheaper, more accessible range of products. The assumption that trade is good for economic welfare underpins many modern models of international trade; indeed, Arkolakis et al. (2012) demonstrate that in many of the quantitative trade models used for modern applied analyses, economic welfare is decreasing in the share of expenditure spent on domestically produced goods.

Of course, many researchers have tried to verify empirically this link between trade openness and economic welfare. In a seminal contribution on this topic, Frankel and Romer (1999) use geographic variables to construct an instrument for openness to international trade, finding that trade openness has a large impact on income. Alcalá and Ciccone (2004) use the same instrument as Frankel and Romer (1999) but with increased sample coverage and more robust controls for institutional quality. They also find a strong empirical link between trade and productivity. Redding and Venables (2004) further highlight the potential importance of geographical access to international markets for per capita income; they develop empirical measures of market access (i.e. trade openness) in a manner consistent with the relevant parameters from trade theory. Anderson et al. (2020) build on these previous approaches using panel data techniques and recent developments from the structural gravity literature to construct instruments for trade openness that are both time-varying and theoretically consistent. Once again, they find that trade has notable positive impacts on countries' incomes.

We have substantial evidence that temperature and trade openness separately have strong potential growth impacts. Trade openness and temperature might also have an interactive effect on growth: economists often propose that international trade could be an important channel for adaptation to climate change. Interest in this hypothesis that international trade can play a role in global adaptation to climate change goes back to early contributions such as Reilly and Hohmann (1993) and Randhir and Hertel (2000). As discussed by Copeland, Shapiro, and Taylor (2022), the key mechanism underlying this hypothesis is the fact that climate change will have heterogeneous impacts on countries around the world, and therefore the optimal global spatial distribution of production will change. In other words, countries are likely to lose their comparative advantage in some products and sectors but gain comparative advantage in others. The more the global economy can shift towards the new optimal spatial distribution of production, the more we can soften the impact of climate change on productivity. However, large barriers to trade means that the global market signals that incentivize these shifts are weaker. For example, agricultural producers will be more likely to switch to new crops that are better-suited to the changed local climate conditions if they have access to a large demand pool via international markets. Meanwhile, if locally-produced supply is restricted due to climate change, consumers are more able to switch to alternative suppliers if they face low costs of importing. In general, trade barriers increase prices and therefore distort market signals indicating the most efficient, low-cost producer. Low trade barriers ensure that price signals provide a good indication of the suitability of local climates to producing a given product, thereby helping to ensure that the global spatial distribution of production shifts in alignment with climate change.

In the agriculture sector in particular, several papers have found evidence that international trade can play a role in adaptation to the impact of climate change on agricultural yields. For example, Costinot et al. (2016) use a quantitative trade model combined with projections of field-level impacts of climate change on crop yields around the world. They find that the negative impact of climate change on economic welfare is slightly larger when adjustments in exports are restricted compared to when they are not, suggesting that international trade can play a role in improving adaptation to climate change via changes in field-level crop choices. Building on this approach by improving the calibration of demand and supply elasticities in the trade model, Gouel and Laborde (2018) find an even larger role for international trade in adapting to climate change. They find that not allowing bilateral import shares to adjust increases the impact of climate change by 76%. Moore et al. (2017) combine projections from agriculture models with the GTAP model of international trade to assess the global impact of climate change on the agricultural sector. They find that reduced yields due to higher temperatures may be offset by increased prices on global markets for some crops, which could help shield producers from these negative productivity impacts. Moving beyond a singular focus on the agriculture sector alone, Nath (2020) explores the potential for hot countries' comparative advantage to shift away from the agriculture sector and towards the manufacturing and services sectors due to the relative sensitivity of agriculture to climate change in these countries. Using a quantitative trade model, he finds that international trade reduces the impact of climate change on productivity, and even more so when he simulates an increase in trade openness in poor countries. This body of literature makes an important contribution to our understanding of the link between international trade and climate change, providing evidence that trade and temperature can have an interactive impact on productivity and growth. However, these papers rely on structural models and simulations to address this topic, so the link between temperature, trade, and productivity in their analyses is governed heavily by theoretical assumptions.

A key contribution of this paper is to limit reliance on structural assumptions and, in the spirit of the applied climate economics papers discussed above, use a reduced-form approach to investigate the link between trade openness, temperature, and growth. However, taking an empirical approach to measuring the impact of trade openness on growth comes with a notable challenge: trade openness is potentially highly endogenous to income. To tackle this potential endogeneity, I start by constructing an instrument for trade openness. I begin by estimating bilateral trade costs. These estimates reflect geographical and cultural proximity, any other unobserved time invariant bilateral factors, as well as variation over time in globalization and regional trade agreements.[1] Related to the discussion above on how trade

[1]Importantly, I use a robust set of controls to obtain estimates that represent variation in bilateral trade costs is purged of the effects of country-level variables such as economic size and growth that will confound the estimate of the effect of trade openness on growth in the second stage.

openness can aid adaptation to climate change, these trade cost estimates capture a variety of factors that serve to drive a wedge between production costs and prices, thereby dampening market signals that incentivize shifts in production to better exploit comparative advantage as it shifts due to climate change. Next, I aggregate these bilateral costs across all potential trading partners of a given country. The structure of this aggregation is based on theory; in essence, I construct empirical counterparts to the 'multilateral resistance' parameters from standard quantitative trade models. I follow the approach in Anderson et al. (2020) to purge these measures of trade openness from endogeneity. In the second stage I use this instrument for trade openness in a reduced-form empirical model of income growth that is inspired by approaches in the climate econometrics literature. The empirical model in this paper therefore bridges the gap between developments in international trade and climate change economics.

The remainder of the paper is organized as follows. The next section outlines a theoretical framework that provides a basis for an empirical model of income that includes both temperature and trade openness as key explanatory variables. I present this empirical model and then I describe the data used in the analysis. Finally, I present and discuss the results and then offer some concluding remarks.

2.2 Theoretical background

The following section provides a theoretical basis for the inclusion of both temperature and trade openness in a model of income growth. First, start with a standard framework for production in a given country. Nominal income is given by $Y_{j,t} = p_{j,t}Q_{j,t}$, where $p_{j,t}$ is the producer price in country j in year t and $Q_{j,t}$, quantity produced, is given by:

$$Q_{j,t} = h(T_{j,t})F(A_{j,t}, L_{j,t}, K_{j,t}) \tag{2.1}$$

$F(A_{j,t}, K_{j,t}, L_{j,t})$ is a function that describes how technology, $A_{j,t}$, labour supply $L_{j,t}$, and capital stock are combined to produce goods in country j in year t. $T_{j,t}$ is a measure of temperature in country j and year t. $h(T_{j,t})$ reflects the macroeconomic damages of temperature changes, describing how temperature impacts productivity. This simple framework is standard in the climate change economics literature; in particular, empirical investigations of the effect of temperature on growth, such as Burke et al. (2015) and Kalkuhl and Wenz (2020) start with similar specifications for the link between temperature and macroeconomic output.

Nested within this framework for production is a standard framework for bilateral trade flows from the international trade literature. A variety of micro-foundations lead to similar expressions for the value of bilateral trade between exporter i to importer j (Yotov et al. 2016), commonly known as the structural gravity model:

$$X_{ij,t} = \left(\frac{t_{ij,t}}{\Pi_{i,t} P_{j,t}}\right)^{1-\sigma} \frac{Y_{i,t} E_{j,t}}{Y_t} \tag{2.2}$$

The first term is a measure of the trade barriers associated with sales from i to j, and the second term is a measure of relative economic size. $Y_{i,t}$ is gross nominal income of exporter i in year t, $E_{j,t}$ is gross nominal expenditure of importer j in year t, and Y_t is total global income in year t. $t_{ij,t}$ are bilateral trade costs to export goods from exporter i to importer j in year t, and σ is the trade elasticity of substitution (the standard assumption is $\sigma > 1$). Key parameters in this equation for bilateral trade are the 'multilateral resistance' parameters:

$$P_{jt} = \left[\sum_{i=1}^{N} \left(\frac{t_{ij,t}}{\Pi_{i,t}}\right)^{1-\sigma} \frac{Y_{i,t}}{Y_t}\right]^{\frac{1}{1-\sigma}} \quad \text{and} \quad \Pi_{it} = \left[\sum_{j=1}^{N} \left(\frac{t_{ij,t}}{P_{j,t}}\right)^{1-\sigma} \frac{Y_{j,t}}{Y_t}\right]^{\frac{1}{1-\sigma}} \tag{2.3}$$

These terms aggregate bilateral trade costs across all trading partners of a given country, and an increase in bilateral trade frictions, $t_{ij,t}$, with any trading partner implies an increase in these multilateral resistance parameters. In other words, they measure relative remoteness from the global marketplace and are therefore inverse measures of trade openness (Anderson et al. 2020). Another common way to describe these parameters is as measures of global market access of producers and consumers in a country relative to producers and consumers elsewhere (Head and Mayer 2014). $\Pi_{i,t}$ is outward multilateral resistance, summarizing global market access for producers in country i, and $P_{j,t}$ is inward multilateral resistance, summarizing global market access for consumers in country j. Relating back to the motivation for this paper, producers' access to international markets may impact their ability to adapt to climate change. For example, if climate change impacts the suitability of a geographical area to producing a certain crop, producers may be able to adapt to this change by switching to produce a different crop that is better-suited to the new climate conditions. However, access to a large demand pool for this new crop from international markets can help to ensure that this switch is economically viable. Similarly, the impacts of climate change on consumers may be softened if they can readily access international markets and find alternative suppliers if their usual supply becomes restricted and costly due to climate change.

Equation 2.2 makes clear that a country's aggregate level of trade openness affects its volume of trade with any given partner country. To see that trade openness also affects aggregate income in a country, impose the standard market clearing condition that total nominal income in country j equals total expenditure on goods purchased from all producers (domestic as well as foreign): $Y_{j,t} = \sum_i X_{ij,t}$. Substitute equation 2.2 into this condition and rearrange to obtain:[2]

$$p_{j,t} = \frac{(Y_{j,t}/Y_t)^{\frac{1}{1-\sigma}}}{\rho_j \Pi_{j,t}} \tag{2.4}$$

This equation elucidates the inverse relationship between prices and relative remoteness: global competition forces relatively remote producers to have relatively low prices to compensate for the high costs of getting their products to markets. Anderson et al. (2020) point out that price levels are linked to the level of investment in a country, which is an important determinant of growth. More precisely, we can substitute equation (2.4) for prices into the equation for nominal income, $Y_{j,t} = p_{j,t}Q_{j,t}$ and rearrange to obtain an equation for aggregate income that is a function of trade remoteness/openness, as given by $\Pi_{i,t}^{\sigma-1}$. Furthermore, although the main channel through which temperature impacts aggregate income is via its direct effect on productivity (as illustrated by equation 2.1), we would expect productivity impacts could translate into price effects, which may in turn impact investment. Accordingly, temperature and trade openness may interact with each other to impact prices, investment, and growth. These insights from economic theory provide the basis for the empirical model used in this paper.

Finally, the relative remoteness of buyers in country j to sellers around the world (i.e. $P_{j,t}$) might also impact income levels in country j. Anderson et al. (2020) discuss how prices faced by consumers impact the direct and opportunity costs of investment, which impacts income and growth; they show that in a dynamic model with a log-linear capital accumulation function, the capital stock in country j is directly affected by the relative remoteness of consumers in the previous year, $P_{j,t-1}$. This framework provides a theoretical justification to include the (one period lag) of inward multilateral resistance in the income equation as well. I mainly focus on the effect of outward multilateral resistance on income growth, but also briefly explore briefly the effect of inward multilateral resistance in the results below.

2.3 Empirical model

To estimate the role of trade openness in the relationship between temperature and income growth, I use a two-stage empirical model. In the first stage I construct an instrument for trade openness, and in the second stage I use this instrument in a reduced-form model of income growth.

[2] ρ_j is the preference parameter in the CES utility function. For a full derivation see Anderson et al. (2020) or the appendix of chapter 3 of this thesis.

First stage

The goal of the first stage is to obtain estimates of bilateral trade costs, $\hat{t}_{ij,t}$, and then to aggregate these bilateral variables up to the country-level to obtain a measure of the average trade costs faced by a given country, which serves as the instrument for trade openness. To obtain $\hat{t}_{ij,t}$, I follow the standard approach in the structural gravity literature, estimating the the empirical counterpart of equation (2.2):

$$X_{ij,t} = exp[\pi_{i,t} + \chi_{j,t} + \mu_{ij} + \alpha RTA_{ij,t} + \eta INTL_{ij} \times YEAR_t] \times \varepsilon_{ij,t}$$

$X_{ij,t}$ is the gross value of bilateral trade between exporter i and importer j in year t, including domestic trade flows ($i = j$). $\pi_{i,t}$ and $\chi_{j,t}$ are exporter-year and importer-year fixed effects, μ_{ij} is an exporter-importer fixed effect, and $RTA_{ij,t}$ is a dummy variable indicating whether or not i and j are part of the same trade agreement in year t. $INTL_{ij}$ is a time-invariant dummy indicating international (rather than domestic) trade flows (i.e. $i \neq j$), and so $INTL_{ij} \times YEAR_t$ is a time-varying international border dummy. Following standard practice in the literature on gravity equations for trade, I estimate the above model using the PPML estimator (Correia et al. 2020), which accounts for information in zero trade flows and is robust to bias caused by heteroskedasticity (Silva and Tenreyro 2006). From this estimation I recover estimates of bilateral trade costs as follows (Yotov et al. 2016):

$$\hat{t}_{ij,t}^{1-\sigma} = exp(\hat{\mu}_{ij} + \hat{\alpha} RTA_{ij,t} + \hat{\eta} INTL_{ij} \times Year_t)$$

Once again, the aim of this first stage is to obtain a measure of bilateral trade costs that reflects variation in trade due to bilateral-specific variables such as geographic and cultural proximity as well as bilateral trade policy, but which is purged of the effect of country-specific variables such as economic size and growth that will confound the estimate of the effect of trade openness on growth in the second stage. $\pi_{i,t}$ and $\chi_{j,t}$ play an essential role in absorbing any such country-specific confounding variables, including the multilateral resistance parameters as well as the income and expenditure terms in equation (2.2).

A large body of work explores how to obtain unbiased estimates of the effect of regional trade agreements on bilateral trade, and this paper follows several developments in this literature to reduce endogeneity bias on the $\hat{\alpha}$ estimate. First of all, countries' economic size and growth rates may be correlated with their propensity to form regional trade agreements with other countries. The exporter-year and importer-year fixed effects deal with this concern, ensuring that $\hat{\alpha}$ is not upward biased by this correlation between $RTA_{ij,t}$ and economic size and growth. Baier and Bergstrand (2007) discuss additional concerns of endogeneity bias associated with the $RTA_{ij,t}$ variable due to omitted bilateral variables or simultaneity.[3] They recommend dealing with these issues by using panel data and including exporter-importer fixed effects, which absorb any unobserved time-invariant bilateral factors that may be correlated with both $\varepsilon_{ij,t}$ and the propensity to form a regional trade agreement. Following this recommendation, μ_{ij} in the model above controls for any time-invariant bilateral trade costs including observable frictions such as distance as well as unobserved frictions such as cultural and institutional differences

Anderson and Yotov (2016) point out that omitting domestic sales from the model is a potential source of downward bias on the estimated effect of the $RTA_{ij,t}$ variable: regional trade agreements can pull trade away not only from other international partnerships but from the domestic market as well, and the magnitude of this effect may depend on the size of the domestic market. Accordingly, I include domestic as well as international trade flows in the sample to deal with this source of potential bias. Finally, Bergstrand et al. (2015) explore potential endogeneity bias on the estimate of $\hat{\alpha}$ stemming from unobserved time-varying heterogeneity in bilateral trade costs. They find that the effects of international borders on trade has declined notably over time, and therefore recommend including the $INTL_{ij} \times YEAR_t$ dummies in gravity equations for bilateral trade. These time-varying border dummies control for worldwide average trade integration in a given year, and in particular the trend towards increased globalization over the sample period (Bergstrand et al. 2015).

In light of the robust set of controls in this first stage model, remaining potential sources of bias on the estimate of $\hat{\alpha}$ are any time-varying bilateral-specific factors that affect the propensity of two countries to join a trade agreement and are correlated with bilateral trade.[4] More importantly, this bias undermines the identification strategy of this paper if the source of bias is also correlated with income growth, the second stage outcome variable. Such confounding variables seem likely to account for just a small portion of the variation in bilateral trade, which in turn accounts for just part of income growth, so potential endogeneity bias of this nature may not be

[3]Baier and Bergstrand (2007) argue that omitted variable bias is the biggest issue of concern. In their example of bias due to unobserved heterogeneity in countries' levels of domestic regulations, they argue that the bias on the effect of RTAs on trade is likely to be negative. Regarding the issue of simultaneity (reverse causality between bilateral trade and the propensity to form an RTA), they test for this concern by including leads of the RTA variable in their estimation; they find that future regional trade agreements are not a statistically significant predictor of current bilateral trade levels.

[4]Feyrer (2019) provides an example of a potential confounding variable of this nature. He illustrates how changes in the relative cost of air transport compared to sea shipping over time have benefited some bilateral partnerships more than others because for some partners air distance is much shorter than sea distance (e.g. the UK and Japan) while for others air and sea distance are very similar (e.g. the UK and US). If these relative changes in bilateral trade costs are correlated with the propensity to form an RTA as well as income growth, for example through common technology shocks, then this omitted variable would cause bias on the $\hat{\alpha}$ estimate.

very large. Nevertheless, future work could follow the approach in Bergstrand et al. (2015) to control for exporter-importer fixed effects interacted with a time trend or bilateral distance interacted with year fixed effects, which would further alleviate concerns that the estimate of $\hat{\alpha}$ is biased.

International trade data usually has many zero or missing observations for trade flows, and if this occurs over the entire sample period for a given exporter-importer pair, then the exporter-importer fixed effect is not estimable and therefore $\hat{t}_{ij,t}^{1-\sigma}$ is missing. To fill in these missing values and obtain a full set of bilateral trade cost estimates I follow Anderson and Yotov (2016)'s procedure. First, I regress the (non-missing) estimates of the exporter-importer fixed effects on importer and exporter fixed effects as well as a vector of time-invariant determinants of bilateral trade costs, $\boldsymbol{G_{ij}}$. Once again I use the PPML estimator:

$$exp(\hat{\mu}_{ij}) = exp(\pi_i + \chi_j + \boldsymbol{\gamma G_{ij}}) \times \epsilon_{ij}$$

Using the parameter estimates from this estimation I fill in any missing trade cost estimates with predicted bilateral trade costs: $\hat{t}_{ij,t}^{1-\sigma} = exp(\hat{\pi}_i + \hat{\chi}_j + \boldsymbol{\hat{\gamma} G_{ij}} + \hat{\alpha} RTA_{ij,t} + \hat{\eta} INTL_{ij} \times Year_t)$.

Structural instrument for trade openness

Building on the approach in Frankel and Romer (1999), Anderson et al. (2020) develop 'structural instruments' for trade openness by constructing empirical counterparts to the multilateral resistance parameters from the structural gravity framework (given by (2.3)). These parameters aggregate a country's bilateral trade costs across all trading partners; they are akin to a weighted sum of bilateral trade costs where the weights are relative economic size of the trading partner. As discussed above, a natural interpretation of these parameters is that they represent a country's relative openness to international trade.

However, the theoretical structure for these parameters is endogenous to income for a couple of reasons: (i) they are directly affected by a country's own income, and (ii) they are affected by other countries' incomes (directly and indirectly through the multilateral resistance term in the denominator), which may be related to a country's own income (for example, through common technology shocks). To purge these parameters of this endogeneity, Anderson et al. (2020) suggest removing the intra-national component from the sum (to deal with issue (i)) and using a measure of effective labour force in a pre-sample baseline year rather than contemporaneous income (to deal with issue (ii)). These 'structural instruments' are the empirical counterparts to the theoretical parameters in (2.3):

$$\tilde{P}_{i,t}^{1-\sigma} = \sum_{i \neq j} \left(\frac{\hat{t}_{ij,t}}{\tilde{\Pi}_{i,t}} \right)^{1-\sigma} \frac{N_{i,1990}}{N_{1990}} \quad \text{and} \quad \tilde{\Pi}_{i,t}^{1-\sigma} = \sum_{j \neq i} \left(\frac{\hat{t}_{ij,t}}{\tilde{P}_{j,t}} \right)^{1-\sigma} \frac{N_{j,1990}}{N_{1990}} \tag{2.5}$$

I obtain the trade cost estimates, $\hat{t}_{ij,t}^{1-\sigma}$, from the first stage estimation, as explained above. Recall that the first stage estimation employs several measures to ensure that these trade cost estimates represent variation in bilateral trade costs that is purged of the effects of a country's own income. Most importantly, the first stage estimation controls for exporter-year and importer-year fixed effects, but also exporter-importer fixed effects and time-varying international border dummies help to ensure that the estimated effect of regional trade agreements (which is an important component of $\hat{t}_{ij,t}^{1-\sigma}$) does not suffer from endogeneity bias. $N_{i,1990}$ is country i's effective labour force in the year 1990 (one year before the sample period starts), calculated as the number of workers multiplied by a human capital index. N_{1990} is total effective labour force across all countries in the model. I solve this system of equations for the structural instruments using the fsolve package in Matlab. I avoid calibrating the trade elasticity, σ, by solving for the exponentiated versions of these variables. This system is homogeneous of degree 0, which means that it requires a normalization to ensure a unique equilibrium. In other words, remoteness of a given country is measured *relative* to the remoteness of all the other countries in the world rather than in absolute terms. To normalize the system I set inward multilateral resistance equal to 1 in the reference country.[5]

Second stage

The second stage uses the instrument for trade openness in an empirical model of income growth that closely follows the approach in Burke et al. (2015):

$$\begin{aligned} \Delta ln(GDP_{it}) = h(T_{it}) + \lambda_1 R_{it} + \lambda_2 R_{it}^2 + \kappa_1 ln\left(\frac{1}{\tilde{\Pi}_{it}^{1-\sigma}}\right) \\ + \kappa_2 ln\left(\frac{1}{\tilde{\Pi}_{it}^{1-\sigma}}\right) \times h(T_{it}) + \xi_i + \nu_t + \theta_i t + \epsilon_{it} \end{aligned} \tag{2.6}$$

The dependent variable is the first difference of the natural log of GDP per capita, GDP_{it}. Modelling effects on income growth rather than levels follows the approach in many previous papers in the applied climate change economics literature and importantly avoids issues caused by non-stationarity of GDP levels. In the main specifications, the temperature-growth relationship is given by $h(T_{it}) = \beta_1 T_{it} + \beta_2 T_{it}^2$, where T_{it} is annual mean temperature in country i and year t. R_{it} is total annual precipitation in country i and year t. Country fixed effects, ξ_i, control for time-

[5] In the results shown below, the reference country is South Africa.

invariant factors affecting growth in a given country such as geography, history, etc., and year fixed effects ν_t control for worldwide shocks such as a global recession. The time trend, $\theta_i t$, controls for a linear global growth trend.

The structural instrument for trade openness, $\tilde{\Pi}_{it}^{1-\sigma}$, enters as a fraction to remain consistent with the theoretical framework described above. As bilateral trade costs increase, $1/\tilde{\Pi}_{it}^{1-\sigma}$ (i.e. $\tilde{\Pi}_{it}^{\sigma-1}$), also increases, so this instrument indicates the relative remoteness of country i from international markets; in other words, the instrument is an inverse measure of trade openness. $(\kappa_2 \times \beta_1)$ and $(\kappa_2 \times \beta_2)$ are the key parameters of interest in this model; together, they tells us to what extent remoteness from international markets interacts with the effect of temperature on growth. In line with the hypothesis that trade openness moderates the economic damages caused by increased temperatures, we would expect a positive sign on for κ_2: more remoteness makes the impact of temperature on growth more negative.

As discussed above, this instrument is purged of the effects of a country's own income on trade openness, as well as contemporaneous correlations between a country's own income and incomes of its trading partners. Remaining endogeneity concerns stemming from the construction of the instrument should be minimal, and would require systematic correlation between effective labour forces of partner countries in the pre-sample base year (1990) and remaining unobservable factors affecting income growth after controlling for the fixed effects in the empirical growth equation above. Another potential source of endogeneity is via the estimated effect of regional trade agreements in the first stage; however, as discussed, the first stage model employs an array of measures to alleviate this concerns.

2.4 Data

2.4.1 Data sources

The analysis presented below uses an unbalanced panel of 63 countries spanning 1991 to 2017.[6] In this section I explain the data sources and methods used to construct the sample.

Bilateral trade. The first stage dependent variable, the gross value of bilateral trade flows, includes both international and intra-national trade flows. This variable represents total trade in the manufacturing and agriculture sectors; other sectors of the economy are not included due to limited data availability for constructing the

[6]The only imbalance in the panel is that Slovakia and the Czech Republic do not enter the data until 1993. See the appendix to this chapter for the list of sample countries.

intra-national trade flows. Data on international bilateral trade flows are from UN Comtrade for manufacturing and the FAO Trade Matrix for agriculture (United Nations 2021; FAO 2021). I mainly use reported imports to construct this variable, because this data is more reliable than reported exports, except when a country report no imports but a partner country reports exports, in which case I use reported exports. I assume that missing values for international bilateral trade represent zero trade.

I construct domestic (intra-national) trade flows as $X_{ii} = Y_i - \sum_{i \neq j} X_{ij}$, where Y_i is gross value of production, and the sum $\sum_{i \neq j} X_{ij}$ is taken over all partner countries reported in the data, not just those included the final sample. For Y_i, I use gross value of production (rather than value-added) to ensure consistency with the international trade data, which is reported in in gross (rather than value-added) terms. Gross production data for manufacturing is from the UNIDO and CEPII TradeProd databases and from FAOSTAT for agriculture (UNIDO 2020; Sousa et al. 2012; FAO 2021). I avoid overlaps between these datasets by dropping 254 items from the FAO data which are also reported in the UNIDO data. I do not include FAO items if they cannot be matched to an HS code or do not have corresponding production data available (e.g. live animals).

Trade costs. I use the CEPII gravity database for data on bilateral trade costs (Head and Mayer 2014). Specifically, I use the RTA dummy variable, which equals 1 if the two countries have a regional trade agreement in a given year, and zero if they do not (intra-national trade observations have a value of zero for this variable). As explained above, to obtain the full matrix of bilateral trade costs I use some time invariant bilateral trade cost variables: the natural log of population-weighted distance between the most populated cities of each partner country; a dummy variable indicating if they share a border; a dummy variable indicating if they share a language spoken by at least 9% of the population in both countries; and a dummy variable indicating if they have ever been in a colonial relationship.

Penn World Table. I use the Penn World Table 10.0 for the second stage dependent variable, GDP per capita (Feenstra et al. 2015). I use output-side real GDP at chained PPPs (in million 2017 US$), divided by total population (in millions) to obtain a per capita measure of income. I also obtain measures of effective labour force in each country in 1990 (i.e. for $N_{i,1990}$ in Equation (2.5)) as the number of persons engaged in employment (in millions) multiplied by the human capital index, which reflects both years of schooling and returns to education.

Sector-specific GDP. To obtain measures of GDP specific to the agriculture and manufacturing sectors, I use the Macro Statistics series from the FAOSTAT database (FAO 2021). For the agricultural sector, I use the share of GDP attributable to agriculture, forestry and fishing and multiply this share with GDP per capita from the Penn World Table.[7] For the manufacturing sector, I use the same approach with the share of GDP attributable to total manufacturing.

Weather. Temperature data is from ECMWF Reanalysis v5 (ERA5) hourly grid-level temperature and precipitation data is from the University of Delaware Terrestrial Precipitation: 1900-2017 Gridded Monthly Time Series (V 5.01) (Hersbach et al. 2020; Willmott and Matsuura 2018). I do all transformations (e.g. constructing the square of temperature) at the grid level before spatial aggregation. I spatially aggregate the grid cell observations to the country level using the population-weighted average across grid cells in the country, using weights from the Gridded Population of the World v4 dataset for the year 2000.

2.4.2 Descriptive statistics

Table 2.1: Descriptive statistics for main variables

	Mean	Median	Std deviation	Min	Max
$X_{ii,t}$	320258.14	45054.95	1118888.93	420.34	14508635.64
$X_{ij,t}$	1327.82	32.70	8516.16	0.00	453716.69
$RTA_{ij,t,i\neq j}$	0.24	0.00	0.43	0.00	1.00
$N_{i,1990}$	68.09	18.43	186.82	0.53	1345.06
$GDP_{i,t}$	18270.30	13147.85	14258.96	580.50	73624.98
$T_{i,t}$ (Kelvin)	288.76	288.59	5.81	276.63	301.29
$R_{i,t}$ (m)	9.54	8.20	6.10	0.12	35.33

Notes: Units for trade flows and GDP are million USD.

Table 2.1 shows summary statistics for the key variables used in the analysis. The economic variables are quite skewed (with the median smaller than the mean), particularly the bilateral trade variable. As expected, the sample contains many observations of zero for international trade. The largest bilateral relationships in the sample are exports from China to the USA and trade between NAFTA members (USA, Mexico, and Canada). Intra-national trade tends to be much larger than international bilateral trade; the largest observation for domestic trade occurs in China (14.5 trillion USD) and the smallest occurs in Moldova (420 million USD). Regional trade agreements between countries in the sample become more prevalent

[7]GDP share for agriculture alone is not available in the data, and measures of the level of value-added for agriculture alone are missing in many instances.

over the sample period: the share of international bilateral relationships with an RTA is 9% in 1991 and 40% in 2017. On average over the entire sample period 24% of international bilateral partnerships have an RTA. The countries with the largest and smallest effective labour forces, $N_{i,1990}$, are China and Cyprus, respectively.

The average mean annual temperature in the sample is about 289°Kelvin or 15°C, and the average annual precipitation is 9.4 metres. The minimum temperature observed in the sample (3.5°C) occurs in Russia in 1993 and the maximum temperature (28.1°C) occurs in Thailand in 1998. The minimum annual precipitation (0.12 m) occurs in Egypt in 1999 and the maximum annual precipitation (35.33 m) occurs in Costa Rica in 1993. See the appendix of this chapter for maps of the weather variables.

2.5 Results

2.5.1 Instrument validity

The coefficient estimate on the $RTA_{ij,t}$ variable in the first stage is 0.174 (robust standard error clustered by exporter-importer pairs is 0.034), which implies that entering a regional trade agreement increases bilateral trade between the partner countries by $[exp(0.174) - 1] \times 100 = 19\%$. This result is in line with estimates in previous papers: Anderson et al. (2020)'s coefficient estimate on this parameter is 0.11 and Baier et al. (2019)'s coefficient estimate is 0.293. This comparison provides a sanity check for the first stage results.

Another way to validate first stage results is to examine trade cost estimates, $\hat{t}_{ij}$. To do so, I assume that the trade cost elasticity $\sigma = 6$ (as suggested by Head and Mayer (2014)'s meta-analysis), and then I average the bilateral trade cost estimates across all years that a given pair is observed in the sample: $\frac{1}{T}\sum_t \hat{t}_{ij,t}$. Many intuitive patterns emerge. Pairs with some of the lowest average trade cost estimates include geographically and culturally close countries such as Cyprus-Lebanon, Slovakia-Czech Republic, and Germany-Netherlands. Countries with some of the highest trade cost estimates include pairs such as Israel-Iran, Yemen-Peru, Ethiopia-Kyrgyzstan: countries that are geographically, politically, and/or culturally distant from each other.

Finally, Figure 2.1 shows scatter plots of the instruments for multilateral resistance, which are the reciprocals of the terms given by the equations in (2.5) (i.e. $\tilde{P}_{i,t}^{\sigma-1}$ and $\tilde{\Pi}_{i,t}^{\sigma-1}$). These instruments are inverse measures of trade openness, so larger values imply more remoteness from international markets. Countries with the smallest values for the instrument for outward multilateral resistance, which measures global

Figure 2.1: Scatter plots of the instruments for multilateral resistances

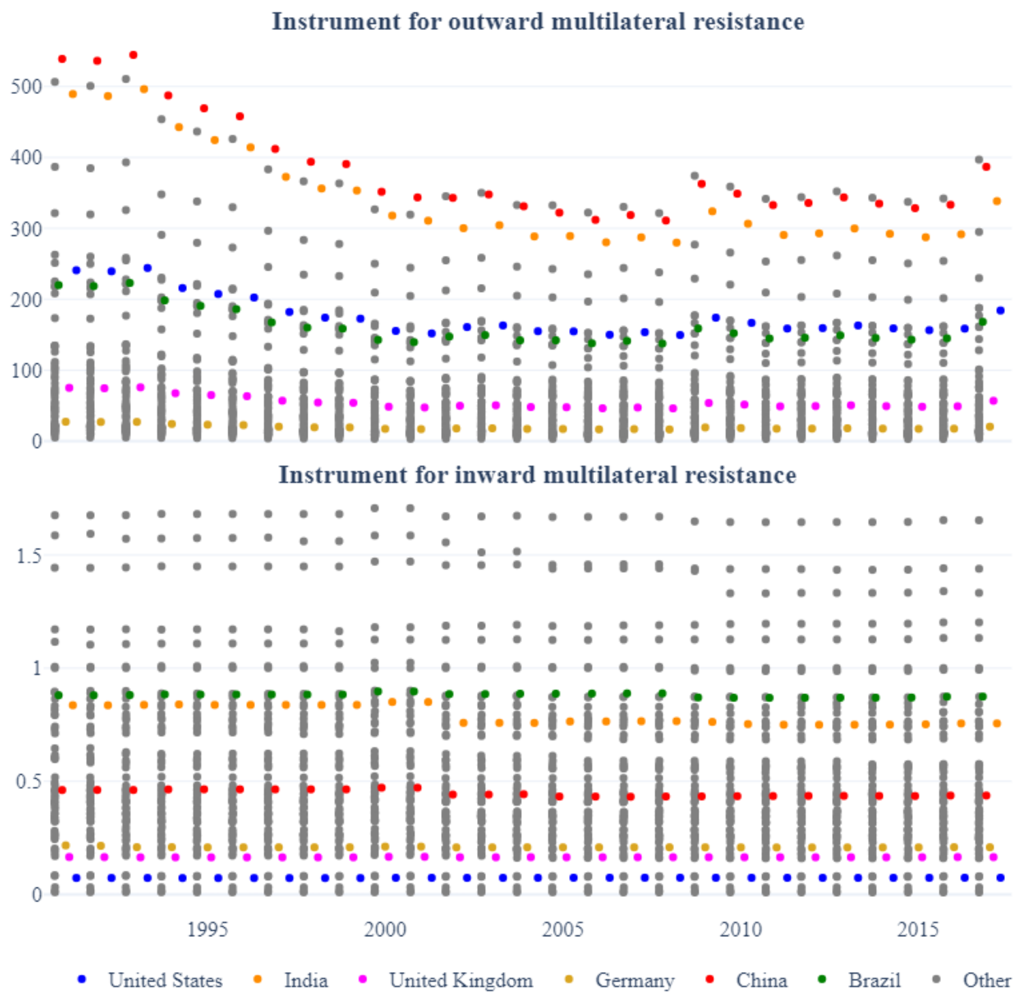

Notes: Each point represents the value of the instrument for a single country-year.

market access for producers in a given country, include small open economies such as Pakistan, Switzerland, and Malaysia. Economies that are large and relatively more self-sufficient and/or are relatively less integrated with other countries have large values for this parameter; for example, Iran, China, and India have some of the highest values in many years of the sample period. The smallest values for the instrument for inward multilateral resistance, which captures global market access for consumers in a given country, occur in Pakistan, Madagascar, and the Netherlands, and the highest values occur in Palestine, Argentina, and Peru. Well-connected countries such as the United States and the United Kingdom tend to be in the low to middle parts of the distributions of these instruments; trade agreements such as NAFTA and the European Union have historically reduced barriers for buyers and sellers in these countries, but the large domestic markets in these economies can pull economic activity away from international markets.

2.5.2 Main results

Table 2.2 shows the main results for the second stage estimation (equation (2.6)). The dependent variable in all columns is income growth (constructed as the first difference of the log of per capita GDP). The bottom panel shows estimates of the marginal effect of temperature on income growth at 25°C and (when relevant) several points in the distributions of the instruments for multilateral resistance. For the purpose of comparing with estimates from previous literature, column (1) omits the trade openness variables and focuses on the effect of temperature shocks on growth. The estimates suggest that at 25°Celsius, a 1°temperature increase is associated with a 2.9 percentage point decrease in income growth. This marginal effect estimate is in line with previous literature. In similar specifications, Burke et al. (2015) find a marginal effect on income growth of -1.3 at 25°C and Kalkuhl and Wenz (2020) find a marginal effect of -2.6 at 25°C.

Column (2) adds the trade openness variable to the model, but again for the purposes of comparing with previous literature, omits the interaction with temperature. Adding the control for trade openness does not notably change the coefficient estimates on the temperature variables. The coefficient estimate on $ln(\tilde{\Pi}_{it}^{\sigma} - 1)$ suggests that a 1% increase in the relative remoteness of exporters from international buyers leads to a 18 percentage point decrease in income growth.[8] For comparison, Anderson et al. (2020) use income levels (rather than growth) as their dependent variable and obtain coefficient estimates on $ln(\tilde{\Pi}_{it}^{\sigma-1})$ ranging from -0.16 to -0.29.

Column (3) of Table 2.2 presents results for the key effect of interest in this paper: the interactive effect of temperature and trade openness on growth. At 25°C and the sample median for relative remoteness of exporters from international buyers, the marginal effect of temperature on income growth is -2.6 percentage points. As illustrated in Figure 2.2, this effect may increase as remoteness increases. This result aligns with hypotheses and previous findings that trade openness can moderate the negative impact of temperature shocks on income growth. Column (4) adds the instrument for inward multilateral resistance, including its interaction with the temperature variables. I follow the approach in Anderson et al. (2020) to use the 1-period lag of this variable. At roughly 25°C and the sample medians for the instruments for trade openness, the marginal effect of temperature on income growth

[8] Note that this coefficient estimate for the remoteness variable is statistically insignificant. The issue of large standard errors is a persistent challenge in this and subsequent tables presented below. In general, income growth is driven by a myriad of factors so we would expect to face a lot of noise when trying to identify precisely the effect of temperature and trade openness on this variable. Future work on this topic could try to address the lack of precision in these estimates by expanding the sample size (in terms of countries as well as years) as well as exploring ideas for alternative instruments for trade openness.

Table 2.2: Main results

	(1)	(2)	(3)	(4)
$\text{Temp}_{i,t}$	0.8602***	0.8821***	-0.5076	0.2179
	(0.2219)	(0.2104)	(0.8444)	(0.8710)
$\text{Temp}^2_{i,t}$	-0.0015***	-0.0015***	0.0009	-0.0004
	(0.0004)	(0.0004)	(0.0014)	(0.0015)
$\text{Precip}_{i,t}$	-0.0017	-0.0020	-0.0016	-0.0015
	(0.0023)	(0.0019)	(0.0024)	(0.0019)
$\text{Precip}^2_{i,t}$	0.0000	0.0000	0.0000	0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0000)
$ln(\tilde{\Pi}^{\sigma-1}_{it})$		-0.1777	-52.7716	-37.5199
		(0.1181)	(33.3602)	(27.9075)
$\text{Temp}_{i,t} \times ln(\tilde{\Pi}^{\sigma-1}_{it})$			0.3597	0.2549
			(0.2288)	(0.1922)
$\text{Temp}^2_{i,t} \times ln(\tilde{\Pi}^{\sigma-1}_{it})$			-0.0006	-0.0004
			(0.0004)	(0.0003)
$ln(\tilde{P}^{\sigma-1}_{it-1})$				-47.0217
				(33.0989)
$\text{Temp}_{i,t} \times ln(\tilde{P}^{\sigma-1}_{it-1})$				0.3257
				(0.2291)
$\text{Temp}^2_{i,t} \times ln(\tilde{P}^{\sigma-1}_{it-1})$				-0.0006
				(0.0004)
Constant	-126.8349***	-124.7772***	78.2913	-30.5774
	(32.2130)	(29.1333)	(122.4593)	(125.3538)
Observations	1634	1634	1634	1634
R^2	0.193	0.198	0.209	0.214
Marginal effects of temperature at 25°C				
	-0.0292***	-0.0294***		
	(0.00877)	(0.00778)		
at 10th percentiles of $\tilde{\Pi}^{\sigma-1}_{it}$ *and* $\tilde{P}^{\sigma-1}_{it}$*:*			-0.0157	-0.0118
			(0.0127)	(0.0151)
at medians of $\tilde{\Pi}^{\sigma-1}_{it}$ *and* $\tilde{P}^{\sigma-1}_{it}$*:*			-0.0258*	-0.0281**
			(0.0110)	(0.0103)
at 90th percentiles of $\tilde{\Pi}^{\sigma-1}_{it}$ *and* $\tilde{P}^{\sigma-1}_{it}$*:*			-0.0361*	-0.0432**
			(0.0150)	(0.0142)

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < .01$, *** $p < .001$

Notes: Bootstrapped standard errors in parentheses. All specifications include country fixed effects, year fixed effects, and a linear time trend.

Figure 2.2: Estimated marginal effects of temperature and remoteness on income growth: Outward multilateral resistance

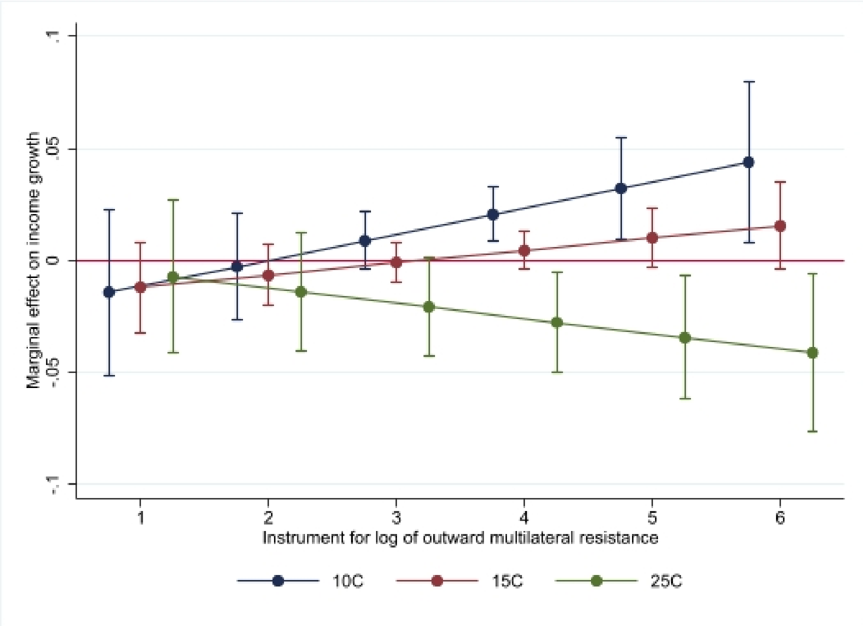

Notes: These estimates correspond to the specification in column (3) of Table 2.2. Larger values on the x-axis imply more relative remoteness of producers to export markets. Recall that remoteness is defined in relation to other countries in the world (discussed in section 2.3). Error bars depict 95% confidence intervals.

is -2.8 percentage points and this effect increases as inward multilateral resistance increases. Once again, this result implies that remoteness from international markets may exacerbate the impact of temperature shocks on income.

In subsequent tables I focus on the instrument for outward multilateral resistance and do not include the instrument for inward multilateral resistance. This choice is motivated by several factors. First, given that within the same country inward and outward multilateral resistance are strongly correlated, controlling for these variables separately may not add a lot of additional information compared to the additional imprecision associated with increasing the number of coefficients to estimate. Next, as discussed in the Introduction, the hypothesized mechanisms through which trade openness impacts the temperature-income relationship often highlight producers' access to international markets. Finally, the structural gravity framework presented above suggests a more direct role for outward multilateral resistance on income compared to the role of inward multilateral resistance.

Not only do these results align with the hypothesis that opennesss to international trade moderates the negative effect of increased temperatures on growth, but they suggest that trade openness may be a key ingredient in the negative temperature-growth relationship established by previous literature. Once the interactive term

Figure 2.3: Estimated marginal effects of temperature and remoteness on income growth: Inward multilateral resistance

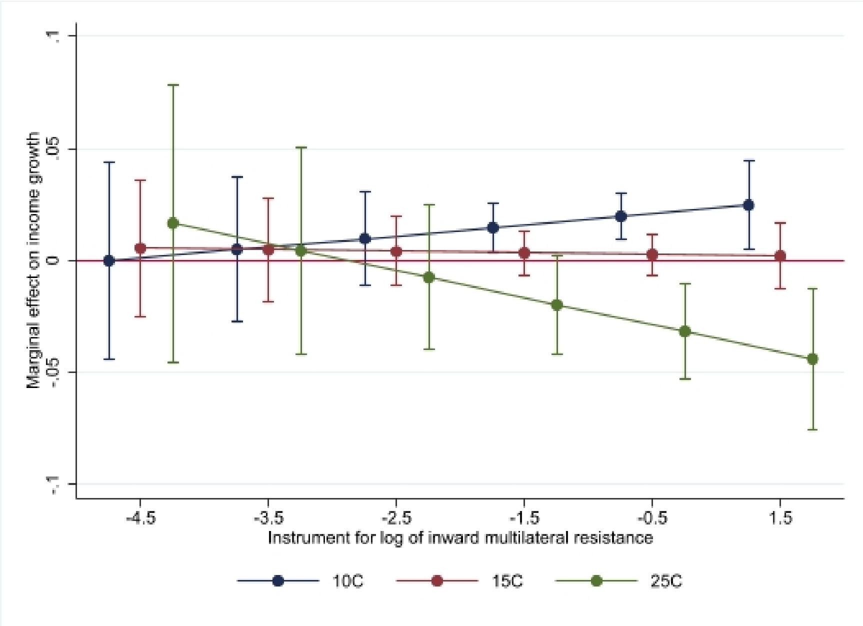

Notes: This figure corresponds to the specification in column (4) of Table 2.2. These marginal effects are computed at the sample average of the instrument for outward multilateral resistance. Larger values on the x-axis imply more relative remoteness of consumers to international markets. Recall that remoteness is defined in relation to other countries in the world (discussed in section 2.3). Error bars depict 95% confidence intervals.

between trade openness and temperature is included in columns (3) and (4), the non-interactive effect of temperature becomes statistically insignificant. Furthermore, as illustrated in Figures 2.2 and 2.3, the estimated marginal effect of temperature on growth is strongest at high levels of remoteness. In other words, these results suggest that the negative effect of temperature on growth may occur mainly in hot *and* relatively remote places. Previous papers have found that the negative effects of temperature shocks on income growth are stronger in hotter places, and this paper contributes evidence that trade openness may be a key moderating factor in this non-linear temperature-growth relationship.

Table 2.3 explores alternative functional forms for this remoteness-temperature-growth relationship. Following the results in Table 2.2 that the interactive effect of remoteness and temperature shocks on income growth seems to occur at high levels for both these variables, I focus on specifications that highlight this relationship. Column (1) interacts the quadratic temperature function with a Remote$_{it}$ dummy variable instead of the continuous variable $ln(\tilde{\Pi}_{it}^{\sigma-1})$. The Remote$_{it}$ dummy takes a value of 1 when $\tilde{\Pi}_{it}^{\sigma-1}$ is in the fourth quartile of the distribution of this variable within year t. Contrary to the results in the previous table, the estimates for this specification suggest that remoteness does not affect the temperature-growth relationship.

Column (2) forgoes the quadratic assumption and instead uses a flexible functional form based on degree days. The effect of temperature on income growth is given by $h(D_{it}) = \sum_{b=0}^{B} \theta_b D_{it,b}$, where $D_{it,b}$ is the number of days during year t in which the mean temperature in country i falls into bin b, and each bin represents a 10°C span in the temperature distribution. The reference bin is days in the 10°to 20°C range. The coefficient estimate $\hat{\theta}_b$ indicates the marginal effect on income growth of an additional day in bin b relative to an additional day in the reference bin. Each bin is interacted with the Remote_{it} dummy variable, which provides an estimate of the additional effect of relative remoteness on this marginal effect. The results are imprecise with a lot of statistically insignificant coefficient estimates. However, the estimates provide some suggestive evidence that extreme temperature days are important drivers of the negative impacts of temperature. More precisely, with the the effect of remoteness separated out, the effect of an additional day in the 20°to 30°C bin is positive. This result is consistent with the hypothesis that hot days may not be harmful to income growth when a country has close links to international markets.

Column (3) of Table 2.3 focuses on extreme temperature days, specifically those below 5°C and above 25°C. Compared to the specification in column (2), this specification hones in on where the main effect seems to be while reducing the number of coefficients to be estimated. These results imply that cold days have a negative impact on income growth, but remoteness does not exacerbate this effect. Meanwhile, hot days have a negative impact on income growth, and this effect seems to be present mainly in particularly remote places. Overall, although imprecise estimates make interpreting the results challenging, Table 2.3 provides some corroborating evidence for the results in Table 2.2 that suggest that trade openness may be an important ingredient in the temperature-growth relationship.

Table 2.3: Alternative functional forms

	(1)	(2)	(3)
$\text{Temp}_{i,t}$	0.8128***		
	(0.1974)		
$\text{Temp}^2_{i,t}$	-0.0014***		
	(0.0003)		
$\text{Remote}_{i,t}$=1	48.0126	0.0088	-0.0057
	(127.2072)	(0.0803)	(0.0317)

Remote$_{i,t}$=1 × Temp$_{i,t}$	-0.3190		
	(0.8851)		
Remote$_{i,t}$=1 × Temp$^2_{i,t}$	0.0005		
	(0.0015)		
$D_{it,b\in(-\infty,0]^{\circ}C}$		-0.0008**	
		(0.0003)	
Remote$_{i,t}$=1 × $D_{it,b\in(-\infty,0]^{\circ}C}$		0.0008	
		(0.0005)	
$D_{it,b\in(0,10]^{\circ}C}$		-0.0000	
		(0.0002)	
Remote$_{i,t}$=1 × $D_{it,b\in(0,10]^{\circ}C}$		0.0002	
		(0.0004)	
$D_{it,b\in(20,30]^{\circ}C}$		0.0004*	
		(0.0002)	
Remote$_{i,t}$=1 × $D_{it,b\in(20,30]^{\circ}C}$		-0.0006	
		(0.0006)	
$D_{it,b\in(30,\infty)^{\circ}C}$		-0.0003	
		(0.0004)	
Remote$_{i,t}$=1 × $D_{it,b\in(30,\infty)^{\circ}C}$		-0.0015	
		(0.0025)	
$D_{it,b\in(-\infty,5]^{\circ}C}$			-0.0009**
			(0.0003)
Remote$_{i,t}$=1 × $D_{it,b\in(-\infty,5]^{\circ}C}$			0.0005
			(0.0005)
$D_{it,b\in(25,\infty)^{\circ}C}$			-0.0000
			(0.0002)
Remote$_{i,t}$=1 × $D_{it,b\in(25,\infty)^{\circ}C}$			-0.0011^{+}
			(0.0006)
Precip$_{i,t}$	-0.0018	-0.0018	-0.0032
	(0.0024)	(0.0025)	(0.0031)
Precip$^2_{i,t}$	0.0000	0.0000	0.0001
	(0.0001)	(0.0001)	(0.0001)
Constant	-120.5750***	-2.9044**	-2.6851*
	(28.7564)	(1.0866)	(1.2484)

Observations	1634	1634	1634
R^2	0.196	0.191	0.192

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < .01$, *** $p < .001$

Notes: Bootstrapped standard errors in parentheses. All specifications include country fixed effects, year fixed effects, and a linear time trend.

2.5.3 Sector-specific results

Table 2.4 estimates the effect of temperature and remoteness separately for the agriculture and manufacturing sectors. The dependent variables are the first difference in the log of sector-specific GDP per capita, and I re-estimate the instruments for multilateral resistance using data on sector-specific bilateral trade. The first two columns show results for the agricultural sector and columns (3) and (4) show results for the manufacturing sector. The bottom panel shows estimated marginal effects of temperature for quadratic specifications.

The marginal effects estimates for the quadratic specification in column (1) suggest that at 25°C, the marginal effect of temperature on agricultural income growth actually becomes less negative as a country becomes more remote from international markets. Meanwhile, in cool countries trade openness may be beneficial for the temperature-growth relationship. At 10°C, temperature shocks tend to increase agricultural income growth, and this effect is strongest for low values of relative remoteness. Intuitively, access to international markets may help producers find buyers for the increase in supply caused by increased temperatures. Figure 2.4 illustrates these relationships. Next, the results in column (2) imply that remoteness does not exacerbate the negative impact of additional extremely hot or cold days on agricultural income growth. As highlighted in the Introduction of this chapter, the agriculture sector is a common subject of the hypothesis that trade openness can improve the temperature-growth relationship. However, these results do not point to a robust relationship between trade openness and the impact of temperature shocks on economic growth for the agriculture sector. In particular, in hot countries this sector is very vulnerable to temperature shocks, but trade openness does not seem to mitigate this risk. Meanwhile, the agriculture sector in cold countries is less vulnerable to temperature shocks overall and trade openness may play some role in reducing this vulnerability further.

Results in column (3) indicate that at 25°C, increased remoteness may increase the manufacturing sector's vulnerability to hot temperatures. However, this trend is imprecisely estimated and the results in column (4) do not provide corroborating

Table 2.4: Sector-specific results

	(1) Ag	(2) Ag	(3) Manu	(4) Manu
$\text{Temp}_{i,t}$	1.6307***		0.0108	
	(0.4529)		(0.8643)	
$\text{Temp}^2_{i,t}$	-0.0028***		-0.0001	
	(0.0008)		(0.0015)	
$ln(\tilde{\Pi}^{\sigma-1}_{it})$	13.3702$^+$		-51.2908	
	(7.7304)		(41.8402)	
$\text{Temp}_{i,t} \times ln(\tilde{\Pi}^{\sigma-1}_{it})$	-0.0915$^+$		0.3451	
	(0.0532)		(0.2881)	
$\text{Temp}^2_{i,t} \times ln(\tilde{\Pi}^{\sigma-1}_{it})$	0.0002$^+$		-0.0006	
	(0.0001)		(0.0005)	
Remote_{it}=1		-0.0361		-0.0018
		(0.0586)		(0.0530)
$D_{it,b\in(-\infty,5]^\circ C}$		-0.0010**		-0.0011**
		(0.0003)		(0.0004)
$\text{Remote}_{it}=1 \times D_{it,b\in(-\infty,5]^\circ C}$		0.0009$^+$		0.0010
		(0.0005)		(0.0009)
$D_{it,b\in(25,\infty)^\circ C}$		-0.0011*		-0.0000
		(0.0004)		(0.0003)
$\text{Remote}_{it}=1 \times D_{it,b\in(25,\infty)^\circ C}$		0.0001		-0.0007
		(0.0006)		(0.0007)
$\text{Precip}_{i,t}$	0.0085	0.0056	0.0022	0.0006
	(0.0055)	(0.0053)	(0.0037)	(0.0037)
$\text{Precip}^2_{i,t}$	-0.0002	-0.0001	-0.0001	-0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Constant	-238.3592***	-2.3417	4.8140	-3.6211*
	(66.0458)	(2.2516)	(124.9815)	(1.5705)
Observations	1634	1634	1620	1620
R^2	0.101	0.097	0.182	0.157
Marginal effects of temperature at 25°C				
at 10th percentile of $\tilde{\Pi}^{\sigma-1}_{it}$:	-0.0453***		-0.0340*	
	(0.0124)		(0.0151)	
at median of $\tilde{\Pi}^{\sigma-1}_{it}$:	-0.0388***		-0.0371**	
	(0.0096)		(0.0142)	
at 90th percentile of $\tilde{\Pi}^{\sigma-1}_{it}$:	-0.0344***		-0.0407$^+$	
	(0.0087)		(0.0224)	

$^+$ $p < 0.10$, * $p < 0.05$, ** $p < .01$, *** $p < .001$

Notes: Bootstrapped standard errors in parentheses. All specifications include country fixed effects, year fixed effects, and a linear time trend.

Figure 2.4: Estimated marginal effects of temperature and remoteness on agriculture income growth

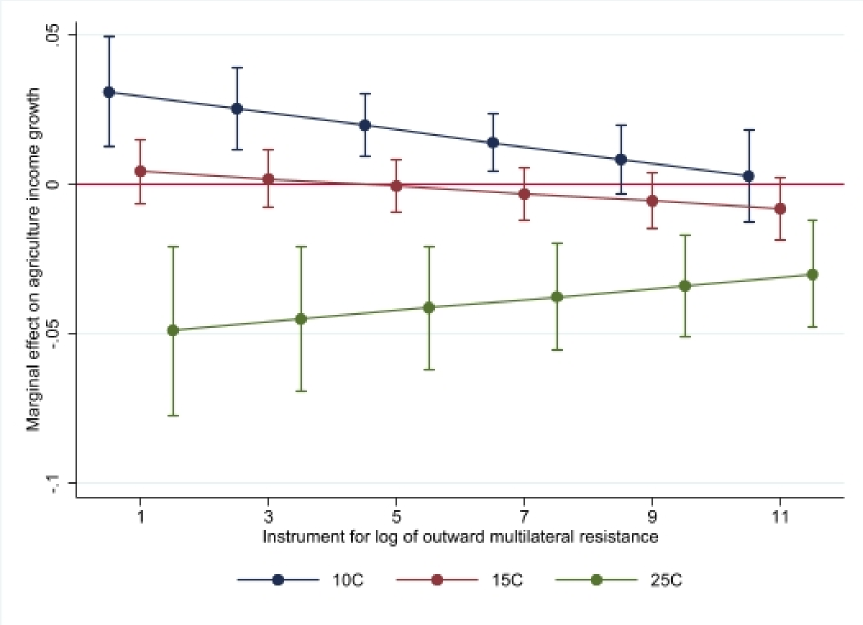

Notes: This figure corresponds the specification in column (1) of Table 2.4. Larger values on the x-axis imply more relative remoteness of agricultural producers to export markets. Recall that remoteness is defined in relation to other countries in the world (discussed in section 2.3). Error bars depict 95% confidence intervals.

evidence. Once again large standard errors are a challenge for making inferences from these results, and it's difficult to say whether the lack of a statistically significant impact for relative remoteness is due to imprecision or lack of an underlying effect.

Overall, these sector-specific results challenge some of the findings for aggregate income discussed above, suggesting that historically trade openness has not always emerged as a mechanism through which macroeconomic growth adapts to temperature shocks. These findings point to a continued need to continue to explore ways to test the empirical validity of the hypothesis that trade openness can help moderate the impacts of climate change, and to better understand the conditions under which trade openness can indeed be an adaptive mechanism to temperature shocks.

2.6 Conclusion

This study tests the hypothesis that openness to international trade can help to moderate the negative impact of increased temperatures on growth. Previous papers that have examined this topic have relied on structural models and numerical simulations; this paper contributes to this literature by examining the question using a reduced-form empirical model and a theoretically-consistent instrument for

Figure 2.5: Estimated marginal effects of temperature and remoteness on manufacturing income growth

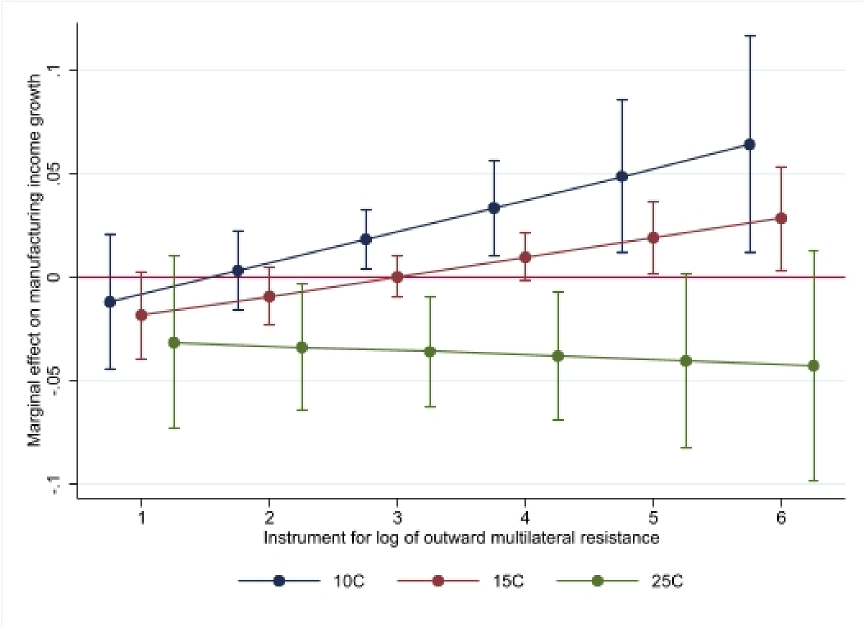

Notes: This figure corresponds to the specification in column (3) of Table 2.4. Larger values on the x-axis imply more relative remoteness of manufacturing producers to export markets. Recall that remoteness is defined in relation to other countries in the world (discussed in section 2.3). Error bars depict 95% confidence intervals.

trade openness inspired by Anderson et al. (2020). The results suggest that historically the negative impact of temperature increases on aggregate income growth has occurred mainly in hot countries that are relatively remote from international markets. In other words, remoteness from international markets could be a key mediating factor in the historical temperature-growth relationship. Nevertheless, the sector-specific results imply that this mediating role of trade openness has not historically been present for the agriculture sector in hot places, which challenges the idea that trade openness can be a key channel for adapting to climate change, given that these places may be particularly vulnerable to climate change. Overall, this study is an important step towards an empirical understanding of the links between trade openness, temperature, and economic growth, and future research can build on this work by uncovering more precisely where in the global economy and under what conditions trade openness emerges as a mechanism to improve the temperature-growth relationship.

A key caveat of this analysis relates to the ongoing debate over whether we can project responses to future climate change from responses to historical climate and weather. To make such projections from the results of this study we need to assume that past responses to temperature shocks are indicative of future responses to gradual, expected temperature change. This assumption is potentially quite strong, and as such this paper does not claim to forecast the role of trade openness in the effect of future climate change on income growth. The quadratic functional form for the temperature-growth relationship allows marginal effects to vary across the temperature distribution and may help to alleviate concerns that the estimates are entirely weather responses and have no basis in local climate (Kolstad and Moore 2020). Nevertheless, the limitations of the reduced-form model to estimate future responses to climate change remain. Future work could explore other functional forms and models (e.g. long difference models) that may get closer to pinning down the response to long term climate change.

Despite these limitations of this reduced-form approach to investigating the role of trade openness on the economic impacts of climate change, this study serves as an important empirical counterpart to structural work on the topic. This paper makes an academic contribution to the field of climate change economics by empirically testing the hypothesis that trade openness can help economies to adapt to the impacts of climate change. In addition, this paper contributes empirical evidence on the potential synergies between trade policy and climate change adaptation policy.

2.A Appendix to Chapter 2

2.A.1 List of sample countries

2.A.2 Additional summary statistics

Figure 2.6: Average annual mean temperature, 1991-2017

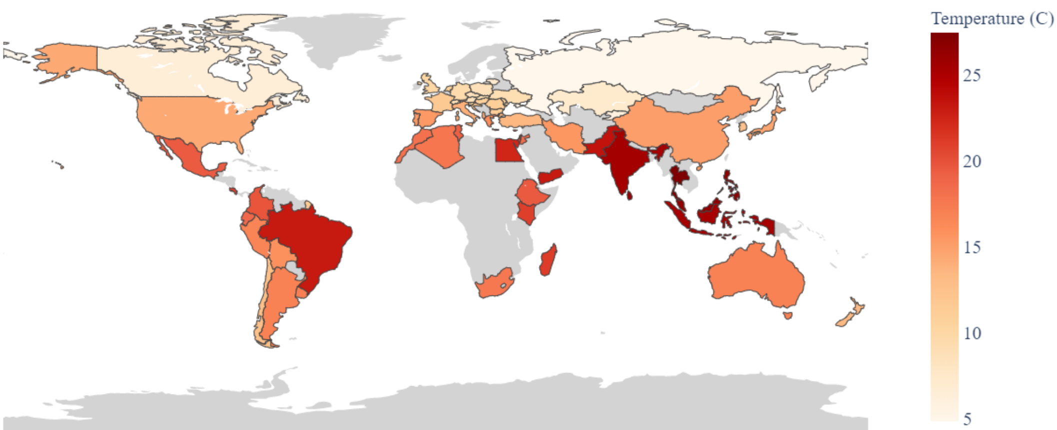

Table 2.5: Sample countries

Albania	France	Mexico
Argentina	United Kingdom	Malaysia
Australia	Greece	Netherlands
Austria	Croatia	New Zealand
Bulgaria	Hungary	Pakistan
Bolivia	Indonesia	Peru
Brazil	India	Philippines
Canada	Iran	Poland
Switzerland	Israel	Portugal
Chile	Italy	Romania
China	Jordan	Russian Federation
Colombia	Japan	Slovakia
Costa Rica	Kazakhstan	Slovenia
Cyprus	Kenya	Thailand
Czech Republic	Kyrgyzstan	Tunisia
Germany	Republic of Korea	Turkey
Algeria	Sri Lanka	Ukraine
Ecuador	Lithuania	Uruguay
Egypt	Morocco	United States
Spain	Republic of Moldova	Yemen
Ethiopia	Madagascar	South Africa

Figure 2.7: Average total annual precipitation, 1991-2017

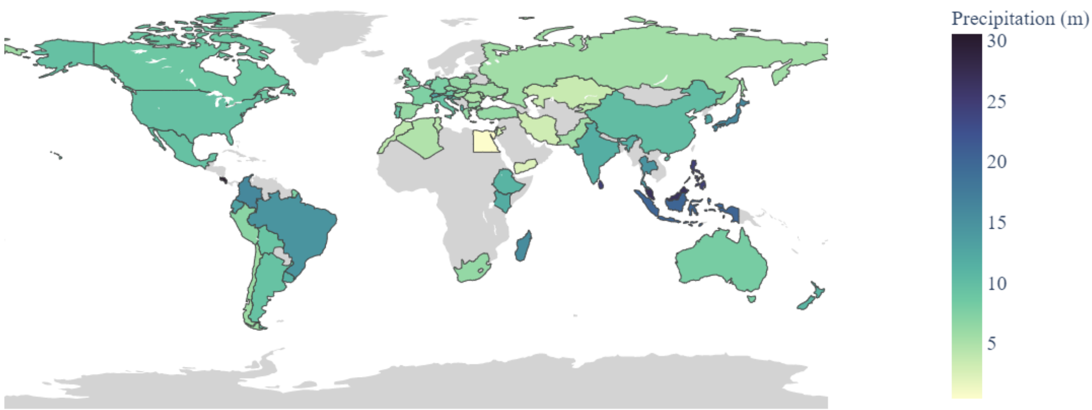

Chapter 3

Exploring the global general equilibrium consequences of a tax on meat consumption in the EU: A structural gravity approach

3.1 Introduction

The UN Food and Agriculture Organisation (FAO) estimates that the livestock sector alone contributes 14.5% of human-induced greenhouse gas (GHG) emissions (Gerber et al. 2013). Similarly, Poore and Nemecek (2018) find that the food supply chain is responsible for 26% of GHG emissions, and 52% of these food-related emissions come from the livestock sector. Meanwhile, Sandström et al. (2018) find that 83% of the carbon footprint associated with European Union (EU) diets are attributable to meat, dairy, and egg consumption. GHG emissions related to livestock production are clearly a large negative externality, and standard Pigouvian theory suggests that pricing the carbon embodied in meat could be an efficient means to correct this market failure. The agriculture sector is currently exempt from the European Union (EU)'s carbon pricing policy, but meat consumption taxes have garnered increasing interest in recent years, particularly consumption-side taxes that overcome constraints on the technical and political feasibility of production-side measures (Wirsenius et al. 2011; Springmann et al. 2016; Funke et al. 2021). A growing literature investigates the potential environmental effectiveness of carbon taxes on meat consumption; however, these studies have mostly focused within the borders of a given country or region, holding international trade relationships constant. Given the interconnectedness of the global economy and global nature of climate change, the effects of a unilateral tax on meat consumption in one region would likely reverberate throughout the global economy, and these ripple effects may impact the effectiveness of the policy in tackling global emissions. However, quantitative assessments of these global effects are a gap in our current evidence base on meat consumption taxes.

This paper addresses this gap in the literature on meat consumption taxes by using a structural model of international meat trade to explore quantitatively the potential global implications of a tax on meat consumption in the EU. A few important features of the policy scenario under consideration are important to keep in mind. First, the aim of the tax is to address a global environmental externality; more precisely, climate change depends on the total stock of carbon in the atmosphere, regardless of where the emissions originate. Second, the policy is unilateral, meaning that it is only imposed in the EU, and meat consumption elsewhere remains unregulated. Next, the tax rate is non-discriminatory; the same tax rate is applied to all meat, regardless of where it was produced or the carbon intensity of its production. This type of carbon policy is more likely to be WTO-compliant than a tax rate that varies based on where or how the meat was produced (Branger and Quirion 2014). Moreover, since the tax does not discriminate based on country of origin (both imports and domestically-produced meat are taxed), the policy does not expose EU meat producers to "unfair" competition from un-taxed meat imports. This feature distinguishes the policy from the EU Emissions Trading Scheme, which does not apply to imported goods and has raised concerns of negative competitiveness impacts on EU producers. Finally, the jurisdiction imposing the tax - the EU - is economically large and so shifts in this market might feasibly shift the global supply and demand equilibrium for meat.

To explore the global consequences of this policy, I extend a standard structural gravity framework to include consumption taxes and then apply this model to international trade data for the meat sector. I simulate an ad valorem tax on meat consumption in the EU, focusing mainly on a 1% tax rate for simplicity of exposition, but I also discuss the sensitivity of my main results to higher tax rates. The model holds meat production quantities fixed across the pre- and post-tax scenarios, so the analysis abstracts from supply-side adjustments and focuses on adjustment mechanisms that occur through international trade. This simplifying assumption means that the results of the tax simulation do not quantify the emissions mitigation achieved by the policy and are thus an incomplete assessment of the effectiveness of the tax. However, my results illustrate how the transmission of the tax through the international trade network generates market signals that create the potential for increases in meat consumption outside the EU. I explore the heterogeneity in these international trade effects across countries, and the changes in countries' relative positions in the international trade network. Examining this heterogeneity provides some insight into how the spatial distribution of meat production and consumption may redistribute in response to the tax.

Under this assumption of fixed production quantities, a simple graphical analysis of global meat supply and demand suggests a few hypotheses for the impact of the tax. We would expect global meat demand to shift downward in response to the tax due to decreased demand from EU consumers, and as a result global meat prices would decrease. Consumers enjoy increased accessibility of meat through decreased prices, while the fixed production assumption suggests that the tax squeezes meat producers' surplus. However, this basic supply-demand analysis does not provide insight into the heterogeneity of these impacts across countries. The tax is essentially an additional barrier to trade, and the introduction of this policy shakes up the distribution of trade frictions across the globe. In this respect, we would expect that some countries are more exposed than others to the impacts of the tax depending on their position in the global trading network and in particular their relative connectedness to the EU market.

A structural gravity-type model has not yet been used to study the topic of meat taxes, but these models are ideal for going beyond a simple graphical supply-demand analysis and assessing these relative changes in trade barriers and opportunities due to the tax. Structural gravity models have become the 'workhorses' of applied quantitative trade analysis for good reason (Head and Mayer 2014): they offer a parsimonious theoretical framework in which most parameters are calibrated or estimated directly from the data, and they combine theoretical consistency with the longstanding empirical power of gravity regressions. In other words, this type of model allows for empirically-grounded assessments of the heterogeneity in the effects of the tax across countries while also providing insight into the international trade mechanisms underlying these effects. As discussed in Head and Mayer (2014), structural gravity models offer a natural decomposition of policy responses in partial trade impacts, modular trade impacts, and general equilibrium trade impacts.[1] This modularity of the structural gravity system opens the 'black box' of general equilibrium trade models to some degree, and allows researchers to examine different aspects of policy impacts separately.

The partial trade impact is the direct effect of the tax on sales of meat in the EU; this response has been studied in several previous papers in the context of single countries as well as the EU as a whole but are not a focus of this paper. This paper assesses the modular trade impact (MTI) and general equilibrium trade impact

[1] Yotov et al. (2016) refer to these effects as partial equilibrium effects, conditional general equilibrium effects, and full general equilibrium effects, respectively.

(GETI) of the tax. The MTI holds prices and income constant across the baseline (pre-tax) and counterfactual (post-tax) scenarios but allows the tax impact to 'ripple out' and affect trade frictions beyond EU borders. Yotov et al. (2016) discuss the advantages of studying the MTI (versus looking only at the GETI). They point out that the MTI isolates policy effects via the trade frictions channel from policy effects via adjustments in economic size. Moreover, the market access parameters (which are the basis of the MTI) allow the researcher to assess the incidence of the tax across meat producers and consumers across the globe. The GETI allows global prices and income to adjust alongside these adjustments in trade frictions. The stylized framework used in this paper holds underlying production quantities constant, so prices bear the full brunt of adjustment to the tax in the full general equilibrium.[2] We can think of the GETI in this model as the very short term impact of the tax when substitution of factors of production away from the meat sector remains limited.

My results for the MTI of the tax illustrate that the policy leads to a diversion of meat sales away from EU markets and towards other markets. As we might infer from a simple global supply-demand analysis in which production quantities are fixed, the incidence of the tax falls on meat producers, particularly EU meat producers, while consumers enjoy increased accessibility of meat and emerge as the relative 'winners' under the policy. Going beyond the simple supply-demand analysis, the structural gravity system highlights how even EU consumers, despite facing a tax on meat consumption, can benefit from the increased competition amongst meat producers due to their relatively close ties to EU meat producers. Moreover, the MTI highlights how not only does the tax 'destroy' trade opportunities between meat producers and EU buyers, but it also creates new opportunities well beyond the orbit of the EU market.

The results for the GETI of the tax confirms the hypothesis from the simple supply-demand analysis that the tax leads to a decrease in meat prices around the world. Under a 1% ad valorem tax, meat producers' prices decrease between 1.5% and 3.5% relative to pre-tax levels, but as the tax rate increases these price adjustments increase; under a 10% tax rate meat producer prices decrease by as much as 25% in some countries. Furthermore, changes in countries' shares of global meat income reinforce the changes in trade frictions caused by the tax, leading to a shift in meat

[2]The GETI is 'general equilibrium' in the sense that international trade patterns fully adjust to the tax, as opposed to an economy-wide general equilibrium in which factor markets also adjust. Furthermore, the model used in this paper is consistent with the assumption of a fixed allocation of a single factor of production, such as labour, used to produce meat using constant returns to scale technology, and so under this assumption the GETI price adjustments can be interpreted as adjustments in wages (i.e. marginal costs).

income away from the EU orbit. In effect, the market signals generated by the international trade effects of the tax suggest that some of the economic activity that the tax aims to regulate may shift to other jurisdictions. The countries that benefit most from this shift are the relatively large and well-connected producers in North America and north Asia. Overall, although the fixed production assumption in the model does not allow a quantitative assessment of the mitigation potential of the tax, the results of this study highlight the potential for global trade impacts to generate market signals that undermine to some extent the tax's aim to reduce emissions related meat consumption, and the magnitudes of the effects estimated in this paper suggest that this risk is non-negligible.

A growing body of literature aims to quantify the efficiency, environmental effectiveness, and equity of carbon taxes on food. As mentioned above, for the most part these studies focus within the borders of a given country or region, implicitly holding international trade relationships constant. Studies have been conducted for: the EU (Wirsenius et al. 2011), the UK (Kehlbacher et al. 2016), Denmark (Edjabou and Smed 2013), Sweden (Säll and Gren 2015; Säll 2018; Andersson 2019), Norway (Abadie et al. 2016), France (Caillavet et al. 2016; Bonnet et al. 2018), and Spain (Garcıa-Muros et al. 2017). Overall, the consensus among these studies is that consumption taxes on food may be an effective, low cost method to reduce emissions from the agricultural sector. Moreover, although trade-offs between health, environmental, and equity outcomes may exist in some cases, a carefully designed policy can ensure that a reduction in meat consumption is compatible with improved health outcomes, especially in Western countries.

A few studies assess carbon taxes on food from a global perspective. Henderson et al. (2018) simulate a global tax on ruminant meat producers to identify supply-side adjustments that would occur in response to the policy, and Springmann et al. (2016) use a large computable general equilibrium model to study the effects of implementing carbon-based food taxes on consumers worldwide. However, these studies do not consider unilaterally-imposed taxes. Zech and Schneider (2019)'s study is the closest in the literature to this paper; they use the EUFASOM model to quantify carbon leakage associated with a carbon tax on food in the EU, finding that carbon leakage may significantly undermine the effectiveness of such policies. The current paper contributes to this evidence base for the global consequences of unilateral taxes on meat consumption using an alternative, more parsimonious and transparent general equilibrium framework.

The study of global general equilibrium effects of unilateral carbon pricing policy is longstanding in environmental economics. Felder and Rutherford (1993) was the first paper to try to quantify what they call fossil fuel price channel carbon leakage, that is, "price-induced substitution" towards carbon intensive products outside of the region implementing a unilateral carbon policy. Branger and Quirion (2014) note that these general equilibrium price effects seem to be large and dominant in studies that assess the global general equilibrium impacts of unilateral carbon policy. This literature often employs large computable general equilibrium models to make these assessments. More recently, however, structural gravity models have become a popular method in the environmental economics literature to assess the global consequences of climate change and carbon pricing policies. Important contributions include Aichele and Felbermayr (2015), Shapiro (2016), and Sager (2019). Larch and Wanner (2017) use a framework that is particularly close to the model used in this paper, but they study economy-wide carbon tariffs rather than sector-specific consumption taxes. Costinot et al. (2016) focus on the agricultural sector, using a gravity-like framework to quantify the negative economic impacts of climate change due to changes in crop productivity. Domínguez-Iino (2021) focuses specifically on agricultural trade between Brazil and Argentina and the rest of the world; he finds that an EU environmental tariff on beef imports from these countries is significantly undermined by increases in exports to non-EU countries. Overall, these studies from the trade and environment literature highlight the potential significance of global general equilibrium impacts of climate policy and demonstrate the usefulness of parsimonious gravity-like frameworks to assess such policies on a global level.

This paper builds on previous literature by employing a structural gravity framework to assess the global consequences of a tax on EU meat consumption. This type of model has not been applied to this particular issue and therefore is an important complement to previous studies that have used large, relatively opaque computable general equilibrium models to address this topic. The next section describes in detail the theoretical framework employed in this paper, and in the following section I explain how I estimate and calibrate the parameters of this model as well as my data sources. Then I present and discuss the results for the global general equilibrium effects of the tax and finally I offer some concluding remarks.

3.2 Theoretical framework

To explore the potential global consequences of an EU tax on meat consumption, I use start with a standard Anderson and Wincoop (2003) structural gravity framework and extend it to accommodate hypothetical consumption taxes. The Anderson and Wincoop (2003) framework is isomorphic to a wide class of structural gravity models with various different micro-foundations and so the welfare implications of my model are not reliant on this choice of framework (Arkolakis et al. 2012). My model has just one sector - meat - and makes an 'endowment economy' assumption that holds production levels constant. This type of model is highly stylized but nevertheless brings insight into some key global trade impacts of the tax in a tractable environment. In particular, the model represents the very short term scenario in which substitution opportunities for production and consumption are limited.

Model set up

In a model with n countries, aggregate utility of the representative consumer in country j is a Cobb-Douglas combination of consumption quantities of meat, Q_j, and a composite outside good, O_j:

$$U_j = Q_j^{\eta_j} O_j^{1-\eta_j} \tag{3.1}$$

η_j is a fixed parameter representing the share of income spent on meat in country j. In this simple framework, trade and production for the outside good are exogenous, while the representative consumer in country j chooses over varieties of meat according to the following CES utility function:

$$Q_j = \left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

Under the standard Armington assumption, each country produces a differentiated meat product - for example, meat from Germany is differentiated from meat from Brazil - and consumers prefer variety across this product space. q_{ij} is consumption in country j of country i's variety of meat, σ is the elasticity of substitution between different varieties (assume $\sigma > 1$), and β_i is an exogenous quality parameter for country i's variety of meat.

Consumers in country j pay the following price for meat imported from country i:

$$p_{ij} = p_i t_{ij} \tau_j$$

p_i is the producer price for meat in exporting country i. t_{ij} is the cost to ship meat from country i to country j, which takes the standard 'iceberg' structure: a proportion $1/t_{ij}$ of the shipment 'melts' on the way from country i to country j such that to deliver one unit of meat to buyers in country j, producers in country i must ship $1 + t_{ij}$ worth of meat (Anderson 2011). Trade costs are normalized such that $t_{ij} = 1$ when $i = j$, and thus can be interpreted as the cost to ship meat from country i to country j *relative* to the cost of shipping the good within country i's domestic market. τ_j is 1 + the ad-valorem tax on meat consumed in country j. This tax is non-discriminatory: it is applied to all goods regardless of their origin, domestic or foreign. In the policy under consideration in this paper, $\tau_j = 1$ if $j \notin$ EU. This structure for preferences implies the standard CES price index for meat consumed in country j:

$$P_j = \left(\sum_i (\beta_i p_{ij})^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{3.2}$$

Total meat expenditure in country j, E_j, is given by:

$$E_j = \sum_{i=1}^{N} p_i t_{ij} \tau_j q_{ij} \tag{3.3}$$

Following the assumption of a Cobb-Douglas aggregate utility function, meat expenditure is a fixed share (η_j) of total income. Consumers maximize utility subject to the following budget constraint for meat expenditure, E_j:

$$E_j = \eta_j \big(R_j + Y_j + T_j \big) \tag{3.4}$$

$$E_j = \eta_j \big(R_j + Y_j + (\tau_j - 1) \sum_i X_{ij} \big) \tag{3.5}$$

R_j is (exogenous) nominal income from production of the outside good in country j, Y_j is nominal income from meat production, T_j is tax revenues in country j, and X_{ij} is the value of bilateral trade in meat from exporter i to importer j. Analogously, expenditure on the outside good is a $(1 - \eta_j)$ share of total income; while outside good income is exogenous, outside good expenditure implicitly adjusts as meat sector income and tax revenues adjust. However, substitution possibilities remain limited in this model given that the outside good prices and output remain fixed and the Cobb-Douglas structure assumes a constant elasticity of substitution of 1. Finally, this framework assumes balanced trade (income = expenditure) for the aggregate economy of each country.

Bilateral trade

Utility maximization leads to a standard structural gravity expression for the nominal value of meat trade between exporter i and importer j:[3]

$$X_{ij} = \frac{Y_i E_j}{Y} \left(\frac{t_{ij}}{\Pi_i P_j} \right)^{1-\sigma} \tau_j^{-\sigma} \tag{3.6}$$

Y_i is country i's nominal income from meat production, Y is global nominal meat income, $\sum_n Y_i$, E_j is country j's expenditure on meat, and Π_i and P_j are the structural parameters known as 'multilateral resistance' terms (defined below). Importantly, X_{ij} includes intra-national ($i = j$) as well as international ($i \neq j$) sales, allowing the

[3]For a detailed derivation, see section 3.A.1 in the appendix to this chapter

model to capture the full range of choices that consumers face between domestic as well as foreign varieties.

Equation 3.6 makes clear the first-order partial trade impact of the tax (τ_j) on sales (X_{ij}). For bilateral relationships (including intra-national trade) in which the buyer is from the EU, the tax decreases the value of trade by $(\tau_j - 1)^{-\sigma}$. Most studies of carbon taxes on meat have focused on this direct effect of the tax on meat sales within the taxing jurisdiction (i.e. the direct effect of τ_j on X_{ij} for $j \in EU$). This partial trade impact holds international trade linkages constant; however, the tax will have a ripple effect on outcomes beyond EU borders, which is the focus of this paper. The structural gravity model allows us to decompose these 'ripple effects' into the modular trade impact (MTI) and general equilibrium trade impact (GETI) (Head and Mayer 2014; Yotov et al. 2016).[4]

Modular trade impact

The modular trade impact of the model is the solution to the system of equations for multilateral resistances, which nests within the structural gravity system. Inward and outward multilateral resistance terms (Π_i and P_j in equation (3.6)) are key parameters in the structural gravity framework and summarize market access and the incidence of trade frictions. Anderson (2011) explains that these multilateral resistance terms capture the idea that bilateral trade depends on bilateral trade costs *relative* to trade costs with the rest of the world. Country i's outward multilateral resistance, Π_i, increases as meat producers in this country face increased barriers to trade:

$$\Pi_i^{1-\sigma} = \sum_j \tau_j^{-\sigma} \left(\frac{t_{ij}}{P_j} \right)^{1-\sigma} \frac{E_j}{Y} \tag{3.7}$$

Country j's inward multilateral resistance, P_j, increases as consumers in this country face increased barriers to buying meat:

$$P_j^{1-\sigma} = \tau_j^{1-\sigma} \sum_i \left(\frac{t_{ij}}{\Pi_i} \right)^{1-\sigma} \frac{Y_i}{Y} \tag{3.8}$$

Equations 3.7 and 3.8 represent a $2 \times n$ system of equations that can be solved numerically (subject to a normalization) in the baseline (pre-tax) scenario and the tax scenario.[5] Comparing outcomes across these two scenarios yields the MTI of the tax. This impact is 'modular' because it holds prices and incomes constant across the baseline and tax scenarios while allowing for adjustments in relative trade fric-

[4] Yotov et al. (2016) refer to these effects as conditional general equilibrium effects and full general equilibrium effects, respectively.

[5] This system is homogeneous of degree zero and therefore requires a normalization to obtain a unique solution. I normalize by setting $P_j = 1$ in the reference country.

tions to ripple through the global trade system.

This system of equations illustrates how the model allows the tax to impact more than simply sales of meat to EU buyers. Equation 3.8 shows that the tax has a first order impact on EU buyers of decreasing their access to the global marketplace (P_j increases); for buyers outside the EU, the tax has no first order impact on P_j (because $\tau_j = 1$ in this case). Equation 3.7 shows that all producers face a first order impact in which the tax decreases their market access (Π_i increases), but those that sell a lot to the EU (and therefore have smaller t_{ij} for $j \in$ EU), are the most strongly impacted. So the first order impact of the tax is to increase trade frictions faced by EU consumers and all producers.

The second order effect of the tax counteracts these first order effects. The increase in P_j for EU countries means that EU countries have less weight in outward multilateral resistances, Π_i, and this effect is strongest for countries with strong ties to the EU (low t_{ij} when $j \in$ EU). So the second order effect decreases multilateral resistance for producers, particular those strong ties to the EU. Meanwhile, the first order increase in Π_i means that sellers (especially those with strong ties to EU buyers) have less weight in inward multilateral resistance, P_j, and this effect is stronger when the two countries have close ties (t_{ij} is low). So the second order effect of the tax improves market access for buyers, particularly those with strong ties to sellers with strong ties to the EU. Overall, the second order effects of the tax reduce trade frictions.

These mechanisms behind the system of equations given by 3.7 and 3.8 reflect an intuitive hypothesis of the effect of the tax: by increasing the barriers to meat sales in the EU, meat sales outside the EU become *relatively* easy. In other words, the 'trade destruction' effect on meat sales in the EU may be offset to some degree by a 'trade creation' effect outside the EU. The tax has a ripple effect beyond EU borders simply by shaking up the distribution of relative trade frictions across countries. Moreover, by separating the effects of these changes in trade frictions on producers versus consumers in each country, the changes in multilateral resistances due to the tax bring insight into the incidence of the tax on producers and consumers around the globe (Yotov et al. 2016).

Equation 3.6 illustrates that the amount of bilateral trade flowing from exporter i to importer j is driven by two channels: the first channel relates to economic size (given by the term $(Y_i E_j)/Y$), and the second channel relates to trade frictions (given by $(t_{ij}/\Pi_i P_j)^{1-\sigma}\tau_j^{-\sigma}$). The MTI closes the first channel by definition, and so the MTI of the tax on bilateral trade can be summarized by changes in the second 'trade frictions' channel.[6] Therefore in the results below I use the percent change in this trade frictions channel as a key tool to assess the global MTI of the tax.

General equilibrium trade impact

Under the general equilibrium trade impact, p_i, and therefore nominal output and expenditure, Y_i and E_i, adjust to the tax. Nominal output is given by $Y_i = p_i Q_i$ where Q_i is the quantity of meat production in country i. In this simple framework I make an endowment economy (a.k.a. exogenous supply) assumption so underlying meat production, Q_i, does not adjust to the tax, and therefore the production-side adjustment to the tax occurs fully through producer prices.[7] Obviously this assumption is highly restrictive, but it offers a tractable and transparent framework in which to consider the GETI of the tax. We can think of it as a very short term effect of the tax before factors of production are re-allocated away from the meat sector and meat production quantities adjust. To close the model and obtain expressions for the GETI of the tax, impose market clearance, $Y_i = \sum_{j=1}^{N} X_{ij} \forall i$, and substitute the expression for bilateral trade, equation 3.6, into this condition:

$$Y_i = \sum_j \tau_j^{-\sigma} \left(\frac{\beta_i p_i t_{ij}}{P_j} \right)^{(1-\sigma)} E_j \tag{3.9}$$

Divide the above equation by global meat income, $Y = \sum_i Y_i$, and rearrange to obtain the following expression for meat producer prices in country i:[8]

$$p_i = \frac{(Y_i/Y)^{\frac{1}{1-\sigma}}}{\beta_i \Pi_i} \tag{3.10}$$

Equation 3.10 elucidates the link between the MTI and the price effects under the GETI scenario. As discussed above, the tax has a first order effect of increasing outward multilateral resistance (corresponding to an increase in Π_i), and more so for producers with close ties to EU buyers. Referring to equation 3.10, this increase in outward multilateral resistance leads to a decrease in producer prices. Intuitively, producers decrease their prices to compensate for the increase in trade frictions they face when selling meat in the EU. This effect is strongest for countries with closest ties to the EU, and so the buyers that stand to gain the most from this first order decrease in prices are non-EU consumers that have close ties to producers with close ties to the EU. Similarly, the producers that stand to gain the most from these price

[6]This channel is my model's analogue to the constructed trade bias (CTB) index established by Agnosteva et al. (2014).

[7]This assumption is equivalent to assuming perfect competition and a single factor of production with constant returns to scale in meat production, i.e the production function is $Q_i = w_i L_i$ where w_i is the wage rate and L_i is labour used in meat production in country i. Perfect competition implies that producer prices equal the marginal cost of production, which in this case is the wage rate, so $p_i = w_i$.

[8]I show this derivation in more detail in section 3.A.1 of the appendix to this chapter.

effects are those with strong alternatives to the EU market. Through this price mechanism meat sales are redirected away from the EU market and towards other buyers. These GETI price effects are a key concern regarding the global impact of a tax on meat consumption in the EU: though the tax incentivizes decreased meat consumption in the EU, the global GETI effects of decreased prices incentivize increased meat consumption outside the EU, and to some extent meat consumption and production may simply shift from the EU to other regions. Given that the goal of the tax is to reduce global emissions and thereby tackle the global climate change externality, these GETI price effects could potentially undermine the goal of the tax to some extent.

In this simple model, producer prices bear the full brunt of adjusting to the tax, so we would expect these GETI price effects to be very strong. In a more complex model that better represents medium to long term responses, both prices and quantities will adjust towards a new market equilibrium. The degree of adjustment in quantities depends on the extent to which consumers can substitute away from meat and producers can substitute factors of production to other sectors. In the highly restricted model presented above these substitution opportunities do not exist.[9]

3.2.1 Solving the model

In the policy simulation presented below, I use this framework for global meat trade to assess the global trade effects of a 1% tax on meat consumption in the EU. Solving the model for the impact of this policy involves two main steps. First, I solve the baseline (pre-tax) model, in which $\tau_i = 1 \forall i$, and then I solve the counterfactual (post-tax) model, in which $\tau_i = 1.01$ if $i \in EU$ and $\tau_i = 1$ otherwise.

To solve the baseline model I start by solving the $2 \times n$ system of equations given by 3.7 and 3.8 to obtain baseline multilateral resistances. This system requires information on the bilateral trade cost parameter, which I estimate (explained further below), and baseline meat income and expenditure, which I calibrate directly from data for the base year (2015). This system of equations is homogeneous of degree zero and therefore requires a normalization to obtain a unique solution. In both the baseline and counterfactual models I normalize the system by setting inward multilateral resistance equal to 1 in the reference country (Australia). Accordingly, all results should be interpreted with this normalization in the mind, that is, as changes relative to changes in inward multilateral resistance in Australia. After obtaining the baseline multilateral resistances, I solve for β_i using equation 3.10 and assuming

[9] The 'general equilibrium' price effects in this model are potentially upper estimates of the general equilibrium price effects that could result from the tax, reflecting a global adjustment to the tax that occurs fully through prices while quantities cannot adjust.

that pre-tax baseline prices $p_i = 1 \forall i$.[10]. β_i remains constant across the baseline and counterfactual scenarios.

Next I solve for the modular trade impact of the tax. I assume meat income and expenditure remain constant at baseline levels and solve the system of equations for multilateral resistances with the EU tax imposed. This solution represents the MTI of the tax. Finally, I solve for the general equilibrium trade impact of the tax on multilateral resistances as well as prices using the $3 \times n$ system of equations given by 3.7, 3.8, and 3.10. Although quantity produced remains exogenous in this system, nominal income is endogenous via adjustments in producer prices, and the budget constraints for EU countries also adjusts via tax revenues. In the results presented below, I show the effects of the tax as changes relative to the baseline pre-tax scenario. I solve both baseline and counterfactual systems using a trust-region-type algorithm from Matlab's fsolve package.

3.3 Estimation and Data

3.3.1 Estimating trade costs

Using the standard approach in the structural gravity literature, I estimate the (exponentiated) trade cost parameter, $t_{ij}^{1-\sigma}$, from the empirical counterpart to equation 3.6:

$$X_{ij,t} = exp(\pi_{i,t} + \chi_{j,t} + \mu_{ij} + \gamma_k RTA_{ij,t}) \times \varepsilon_{ij,t}$$

The value of bilateral meat trade in period t, $X_{ij,t}$ is a function of an exporter-time fixed effect, $\pi_{i,t}$, an importer-time fixed effect, $\chi_{j,t}$, a country pair fixed effect, μ_{ij}, a dummy variable indicating whether exporter i and importer j have a regional trade agreement in period t, $RTA_{ij,t}$, and random error term, $\varepsilon_{ij,t}$. The importer-time and exporter-time fixed effects absorb the effects of multilateral resistance and economic size, as well as any other country-specific features affecting bilateral trade in a given period. The exporter-importer fixed effect absorbs any time invariant features of bilateral trade costs such as distance between the two countries and whether they share a common border. The $RTA_{ij,t}$ dummy accounts for variation over time in bilateral trade costs related to changes in trade agreements over the sample period.

Following the recommendations of Yotov et al. (2016), I estimate this equation with intra-national as well as international sales and use panel data in intervals, specifically 2-year intervals from 2001 to 2015, which helps to account for slow adjustment of bilateral trade to trade policy changes. I use the the Poisson pseudo maximum

[10]Since I am interested in price changes relative to the baseline rather than price levels, this assumption does not affect interpretation of the results presented below

likelihood (PPML) estimator (Correia et al. 2020); unlike the log-linear OLS estimator, the PPML estimator allows for the inclusion of zero values for bilateral trade and is robust to heteroskedasticity (Silva and Tenreyro 2006; Yotov et al. 2016). Consistent with the structural gravity model outlined in the previous section, I construct estimates of power-transformed trade costs from the above equation as follows:

$$\hat{t}_{ij,t}^{1-\sigma} = exp(\hat{\mu}_{ij} + \hat{\gamma} RTA_{ij,t})$$

For the tax policy simulation, I use trade cost estimates for the final year of the sample (i.e. $t = 2015$). For domestic trades flows ($i = j$) I assume $\hat{t}_{ij,t}^{1-\sigma} = 1$, so the trade cost parameter represents the cost of trade relative to domestic sales. As is common in trade data, the sample includes many observations of bilateral trade flows that are missing or zero across the panel, which means that some exporter-import fixed effects cannot be identified (i.e. $\hat{\mu}_{ij}$ is missing). To obtain the complete set of country-pair fixed effects I follow Anderson and Yotov (2016)'s method: I restrict the sample to a cross-section in the chosen reference year (i.e. 2015) and regress the estimates of the pair fixed effects (from the above regression) on importer and exporter fixed effects as well as a vector static bilateral trade variables ($\boldsymbol{G_{ij}}$). Once again I estimate this equation using the PPML estimator:

$$exp(\hat{\mu}_{ij}) = exp(\pi_i + \chi_j + \boldsymbol{\rho G_{ij}}) \times \varepsilon_{ij}$$

Using my parameter estimates I construct predicted bilateral trade costs and use these predictions to fill in bilateral trade cost estimates not identified in the first stage:

$$\hat{t}_{ij}^{1-\sigma} = exp(\hat{\pi}_i + \hat{\chi}_j + \boldsymbol{\hat{\rho} G_{ij}} + \hat{\gamma} RTA_{ij})$$

3.3.2 Calibrating other parameters

In addition to the trade cost parameter, solving the model as outlined above also requires values for the trade cost elasticity parameter, σ, and the country-specific budget share for meat expenditure, η_i. The trade cost elasticity, σ, is the only parameter that I calibrate from an external source. Head and Mayer (2014)'s review of estimates of this parameter for the entire economy suggests a value of 6 for σ, but Larch and Wanner (2017)'s study suggests that this elasticity is lower for the food and agriculture sectors compared to the rest of the economy; they use 4.76 and 5.01 for the agriculture and food sectors respectively. Accordingly, I use a value of 5 for σ in my model.

This η_i parameter represents the share of aggregate income in country i allocated to meat expenditure. I calibrate this parameter directly from the data for 2015. Meat sector income and expenditure in country i are simply total sales and purchases (respectively) across all trading partners, which I obtain from the data on bilateral trade flows. Non-meat sector income, R_i, is also taken directly from the data. Accordingly I construct η_i as follows:

$$\eta_i = \frac{\sum_j X_{ji}}{R_i + \sum_j X_{ij}}$$

3.3.3 Data

The countries included in the model are the top 20 producers of meat from 2006-2017 in terms of gross production value over this period; together they account for almost 80% of the global value of meat production over this period. Table 3.1 lists these countries and their 3-digit ISO codes.

Table 3.1: Model countries and their ISO3 codes

Argentina	ARG	Australia	AUS
Brazil	BRA	Canada	CAN
China	CHN	France	FRA
Germany	DEU	India	IND
Indonesia	IDN	Iran	IRN
Italy	ITA	Japan	JPN
Mexico	MEX	Netherlands	NLD
Republic of Korea	KOR	Russian Federation	RUS
Spain	ESP	Turkey	TUR
United Kingdom	GBR	United States of America	USA

FAOSTAT. The main data source for this study is the FAOSTAT database (FAO 2021), which provides data to construct the bilateral meat trade variable, $X_{ij,t}$. I use the FAO's 'Item Group' definitions to restrict the sample to products in the Meat Item Group. Data on international trade flows are from FAO's trade matrix. I mainly use reported imports, which are more accurate than reported exports (Yotov et al. 2016), but when reported imports are missing but reported export data is available, I use reported exports. I fill in remaining missing values for international trade with a value of zero. I construct intra-national sales as the difference between gross production value (also from the FAO data) and gross total export value to all partners (not just those included in the model).

CEPII Gravity. As explained above, the procedure to estimate the full set of bilateral trade costs requires data on bilateral trade barriers. I use the following variables on bilateral relationships from the CEPII Gravity database (Head and Mayer 2014): the regional trade agreement dummy variable, population-weighted distance, the shared border dummy variable, the common official language dummy variable, the common religion index, and the dummy variable indicating if the countries have ever been in a a colonial relationship.

ITPD-E. To construct a measure of income from the composite outside good, $R_{i,t}$, I use the International Trade and Production Database for Estimation (ITPD-E) (Borchert et al. 2021). This database spans the agriculture, mining, energy, manufacturing, and service sectors and provides information on both domestic and international trade measured in gross nominal values, which is consistent with the measure of bilateral trade used in this paper. From the ITPD-E data I compute total income (sum of all sales, exports plus domestic) in a given country across all sectors. Then I subtract total meat income (as given by my bilateral trade variable) to obtain non-meat sector income.[11]

Table 3.2: Summary statistics of main variables

	Mean	Std. dev.	Min	Max
Meat income, Y_i	28.20	63.19	1.84	294.10
Outside good income, R_i	2450.54	4517.67	56.34	20468.49
Meat budget share, η_i	0.0464	0.0676	0.0003	0.2202
International trade costs, $\hat{t}_{ij,i\neq j}$	15.37	17.03	0.79	125.67
Domestic meat trade, X_{ii}	33.15	62.30	1.42	277.80
International meat trade, $X_{ij,i\neq j}$	0.13	0.40	0.00	3.13
Regional trade agreement, $RTA_{ij,i\neq j}$	0.33	0.47	0.00	1.00

Notes: Units for income and trade variables are billions USD. All summary statistics are for 2015 (which is the year used in the tax simulation).

Descriptive statistics

Table 3.2 summarizes the main variables used to estimate the trade costs and simulate the tax. The largest meat producers in the model (in terms of 2015 value of production) are China and Brazil, and the smallest are the UK and the Netherlands. The largest economies in terms of outside good income are the USA and Japan and the smallest are Argentina and Iran. Argentina and Iran have the largest share of total income allocated to meat expenditure (η_i), while the USA and the Netherlands have the smallest share. The most prohibitive trade costs in the model occur for exports from Russia to Indonesia and to India, and the lowest trade costs are for exports from Germany to the Netherlands and from Australia to Japan. Unsurprisingly, domestic meat sales are much larger than international meat trade, and international trade has quite a few missing observations which I assume are implicit

[11]The ITPD-E database's source for agriculture sector data is also the FAOSTAT database.

zeros. In 2015, about a third of international exporter-importer partnerships in the model have regional trade agreements. The exporters with trade agreements with the most other countries in the model are Korea and Mexico, and those with the fewest are Russia and Iran.

3.4 Results

3.4.1 Modular trade impact

The following section presents and discusses the modular trade impact (MTI) of a 1% ad valorem tax on meat consumption in the EU. Recall that the MTI holds prices, income, and expenditure constant at baseline (pre-tax) levels, but allows for adjustments in the distribution of global trade frictions in response to the tax. The MTI therefore isolates effects of the tax due to changes in trade frictions from effects of the tax due to changes in economic size, and also allow us to assess the incidence of the tax on producers and consumers around the globe (Yotov et al. 2016). The modular trade impact is the solution to system of equations given by 3.7 and 3.8, and the MTI of the tax are the changes in these multilateral resistance terms under the policy compared to their baseline (pre-tax) levels. Finally, recall that outward multilateral resistance (OMR), Π_i, aggregates trade costs for producers and measures their access to the global market, and inward multilateral resistance (IMR), P_i, measures consumers' market access. These parameters are often referred to as multilateral resistance because they increase in magnitude as market access decreases.

Three key results emerge from the MTI of the tax simulation:[12] (i) the tax decreases market access for all meat producers (OMR increases); (ii) Meat becomes more accessible to consumers (IMR decreases); and (iii) although the tax deteriorates some trade opportunities, it also creates new opportunities for meat trade elsewhere. Results (i) and (ii) are unsurprising in light of a simple global supply-demand analysis of the tax; however, the structural gravity model helps to shed light on heterogeneity in these impacts between countries. Result (iii) is not immediately obvious from a simple 1-region supply-demand analysis, and is an important insight into the trade mechanisms that create market signals that can undermine the effectiveness of the unilateral tax.

Result 1: The tax decreases market access for all meat producers (OMR increases).

The top panel of Figure 3.1 shows the effect of the tax on outward multilateral resistance for each country in the model. Outward multilateral resistance increases

[12] For detailed quantitative results, see Table 3.3 in section 3.A.2 of the appendix to this chapter.

Figure 3.1: MTI effects on MRs: 1% tax on EU meat consumption

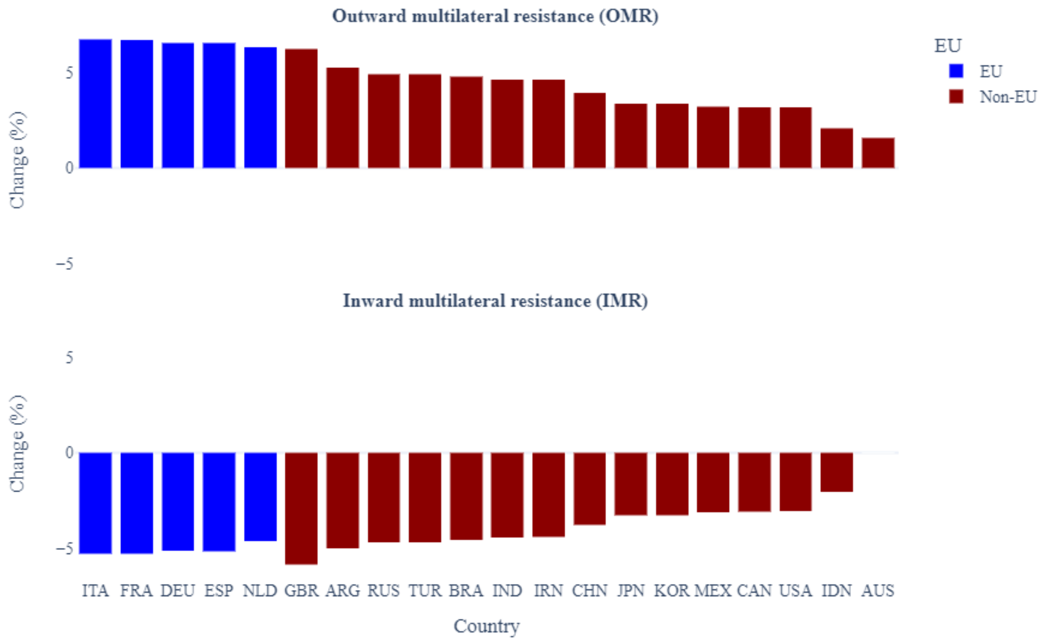

Notes: In both the baseline (pre-tax) and counterfactual scenarios the system of multilateral resistances is solved under the normalization that inward multilateral resistance in Australia, $P_{AUS} = 1$.

for every country, which implies that meat producers everywhere are negatively impacted by this policy. The EU tax is essentially an additional trading friction in the global market, and accordingly all producers now compete in a tougher environment. However, this increased competition does not impact all producers equally. The left panel of Figure 3.2 illustrates that stronger increases in OMRs tend to correlate with lower relative costs of exporting to the EU compared to non-EU destinations. Intuitively, the producers most hurt by the tax tend to be those with the closest relative ties to the EU market.

A potential motivation for taxing all meat consumption, regardless of where the meat is produced, is to protect EU meat producers from negative effects on their relative competitiveness. This aspect of the policy design contrasts with the EU Emissions Trading Scheme, which only prices the carbon from production within EU borders. The decrease in market access across all producers (not just EU producers) confirms this feature of the tax - no producers benefit under the policy through increased market access. However, while the tax does not expose EU producers to "unfair" competition from untaxed imports, EU producers are amongst the most impacted meat producers simply due to the their close relationship with EU consumers. This result highlights that even by taxing all EU meat consumption regardless of the location of production, the incidence of the tax weighs relatively heavily on EU producers. Meanwhile, producers in Argentina, the UK, Turkey, Brazil, and Russia are also strongly impacted by the tax. Compared to other non-EU countries, these countries have much closer ties to EU compared to non-EU buyers, and therefore diverting sales away from the EU market is relatively costly for them.

Result 2: Meat becomes more accessible to consumers (IMR decreases).

The bottom panel of Figure 3.1 shows the effects of the tax on the inward multilateral resistance of each country in the model. The main result is that meat becomes more accessible to consumers in all countries. The increased competition amongst global producers (as discussed above) puts downward pressure on meat prices, allowing consumers around the world to enjoy easier access to meat.

For EU consumers, this increased accessibility of meat may seem somewhat surprising, because the motivation of the tax is to reduce EU meat consumption in absolute terms. Equation (3.8) confirms that the first-order direct effect of the tax is to increase inward multilateral resistance for EU consumers. However, the tax also has a second order indirect effect via the increases in outward multilateral resistances: intuitively, increased competition amongst meat producers leads to downward pressure on meat prices and therefore increased relative access to meat for consumers. As the results in Figure 3.1 illustrate, these second order effects of the tax dominate the first order effects of the tax. To understand this result, note that producers' OMRs have more weight in consumers' IMRs when the the two countries have close ties (trade costs are low), so when EU producers face relatively strong decreases in their market access due to the tax, EU consumers are relatively well-placed to take advantage of this shift. EU consumers face a tax on meat no matter where they buy it from, so the first order effects of the tax do not change their relative preferences between different meat varieties; EU producers only face the tax when they sell to EU consumers, but since it is so much easier for them to ship their products to EU compared to non-EU markets, the tax is not enough to overcome EU producers' preference for EU markets. Despite facing the additional barrier of the consumption tax, EU producers would still rather sell to EU markets compared to non-EU markets, and EU consumers are able to benefit from this situation.

Non-EU consumers also enjoy increased access to meat under the MTI of the tax. Since the tax is only imposed in EU countries, non-EU consumers experience no first order impact of the tax on their inward multilateral resistance, enjoying only the second order effect of the tax through increased competition amongst meat producers. The right panel of Figure 3.2 shows the positive correlation between the relative costs of importing from EU compared to non-EU sources and the magnitude of the decrease in inward multilateral resistance. Consumers in the UK, for example, have low relative barriers to import from the EU compared to non-EU countries, and are therefore well-placed to take advantage of the diversion of sales away from the EU market.[13] More generally, consumers benefit most from the tax if they have close ties with exporters that have close ties with EU buyers. For example, meat producers in Argentina enjoy a relatively low cost of exporting to the EU, and so consumers in Argentina see a large increase in the accessibility of meat because they are well-placed to absorb the meat supply that is diverted away from the EU market due to the tax.

Setting aside potential health and environmental externalities associated with meat consumption, this increased accessibility of meat suggests that consumers emerge as the relative winners under this tax policy, while the incidence of the tax falls heavily on meat producers. Moreover, the MTI for consumers outside the EU illustrate how the impact of tax reverberates beyond EU borders and could potentially incentivize increased meat consumption outside the EU.

Result 3: Although the tax deteriorates some trade opportunities, it also creates opportunities.

Figure 3.3 is a heatmap of the MTI of the tax on trade flows for every bilateral relationship in the model. As discussed above, MTI effects on bilateral trade flows are driven purely by changes in trade frictions, with economic size held constant. In the figure, green and blue cells indicate an increase in bilateral frictions and therefore a decrease in bilateral trade, and pink and red cells indicate a decrease in bilateral frictions and therefore an increase in bilateral trade.

Sales of EU-produced meat in EU markets mainly decrease, but these decreases are relatively small. The strongest decreases in bilateral trade are exports from non-EU to EU countries, which decrease by over 15% in some cases. Related to the discussion above, the tax reinforces the tight knit relationship between EU producers and consumers. Amidst the increased competition in the EU market, EU producers can maintain a strong position in this market because they enjoy relatively low trade barriers to ship meat to EU consumers. Meanwhile, the strongest increases in bilateral trade are exports from EU to non-EU countries, as well as exports from countries with close ties to EU markets (e.g. UK, Argentina, Turkey) to non-EU

[13] An important caveat to this result is that the trade cost parameter is estimated using data from before the UK left the European Union, so these results do not reflect changes trade barriers related to Brexit. Unfortunately this limitation is difficult to overcome until post-Brexit global trade data is more available.

Figure 3.2: The relative cost of trade with EU countries and MTI effects on MRs: 1% tax on EU meat consumption

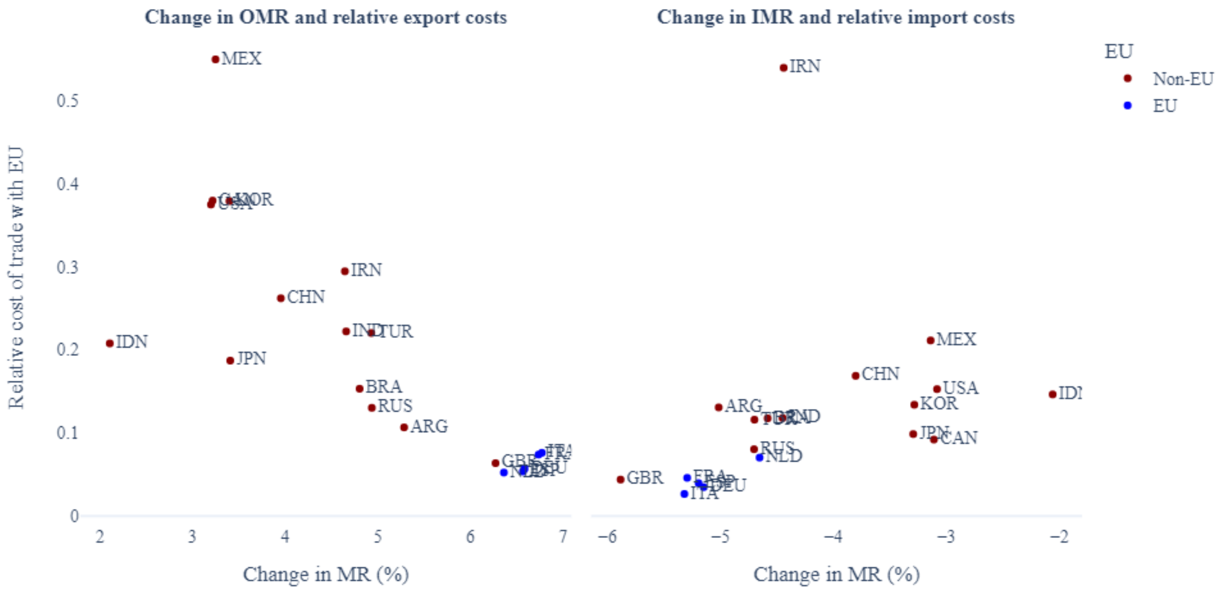

Notes: The relative cost of exporting to EU countries is computed as $\sum_{j \in EU} \hat{t}_{ij} / \sum_{j \notin EU} \hat{t}_{ij}$ and the relative cost of importing from EU countries is defined analogously as $\sum_{i \in EU} \hat{t}_{ij} / \sum_{i \notin EU} \hat{t}_{ij}$. In both the baseline (pre-tax) and counterfactual scenarios the system of multilateral resistances is solved under the normalization that inward multilateral resistance in Australia, $P_{AUS} = 1$.

markets; these trade flows increase by over 15% in some cases. These shifts reflect the diversion of trade away from the EU, as producers try to find alternative markets for meat that was previously sold in the EU.

These changes in bilateral trade flows discussed so far reflect an intuitive global impact of the tax: a portion of trade in the EU market is ‘destroyed’ and diverted to other non-EU markets. However, another somewhat less intuitive ‘trade creation’ effect of the tax emerges: the tax affects even countries that do not have particularly strong ties to the EU as importers or exporters. In some cases partnerships between non-EU countries experience a decrease in trade frictions and therefore an increase in bilateral trade; for example, consumers in Indonesia, India, and Korea import more from Brazil, China, and Mexico. Although they lie outside of the immediate orbit of the EU market, importers in these countries are able to buy more meat because the tax increases global competition amongst meat producers. These results underline the importance of considering the global trade impacts of unilateral climate policy. The first order effects of the tax only impact sales of meat in the EU, but simply through changes in trade frictions the tax reverberates through the international trading network and may have consequences that reach well beyond EU borders and undermine the aim of the tax to decrease carbon emissions related to meat

Figure 3.3: MTI effects on bilateral trade: 1% tax on EU meat consumption

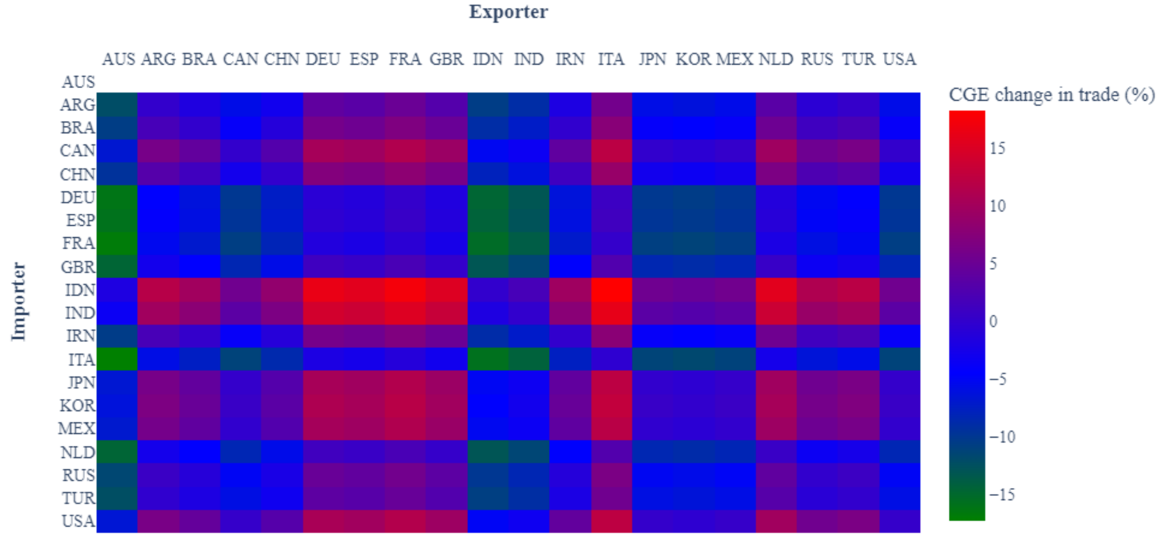

Notes: Australia (as an importer) not shown because it is the reference country. In both the baseline (pre-tax) and counterfactual scenarios the system of multilateral resistances is solved under the normalization that inward multilateral resistance in Australia, $P_{AUS} = 1$.

consumption.

3.4.2 General equilibrium trade impact

Building on the results for the modular trade impact effects of the policy, this section discusses the general equilibrium trade impact (GETI) of a 1% ad valorem tax on meat consumption in the EU. The GETI allows for adjustments in prices and relative economic size as well as changes in trade frictions. In this model, the GETI holds production levels constant, so the results are an assessment of the global trade impacts of the tax in the very short term before any substitution of factors of production away from the meat sector. Furthermore, the model is in nominal terms and is not calibrated to real-world prices or production levels, so the tax simulation does not assess the absolute changes in producer prices in response to the tax, but rather relative changes compared to the baseline (pre-tax) scenario and compared to other countries in the model. Detailed quantitative results of both the MTI and GETI of the tax are available in Table 3.3 in section 3.A.2 of the appendix to this chapter.

Result 4: Meat producer prices decrease.

Once again, a simple graphical analysis of global meat supply and demand in which

Figure 3.4: GETI effects on producer prices and relative economic size: 1% tax on EU meat consumption

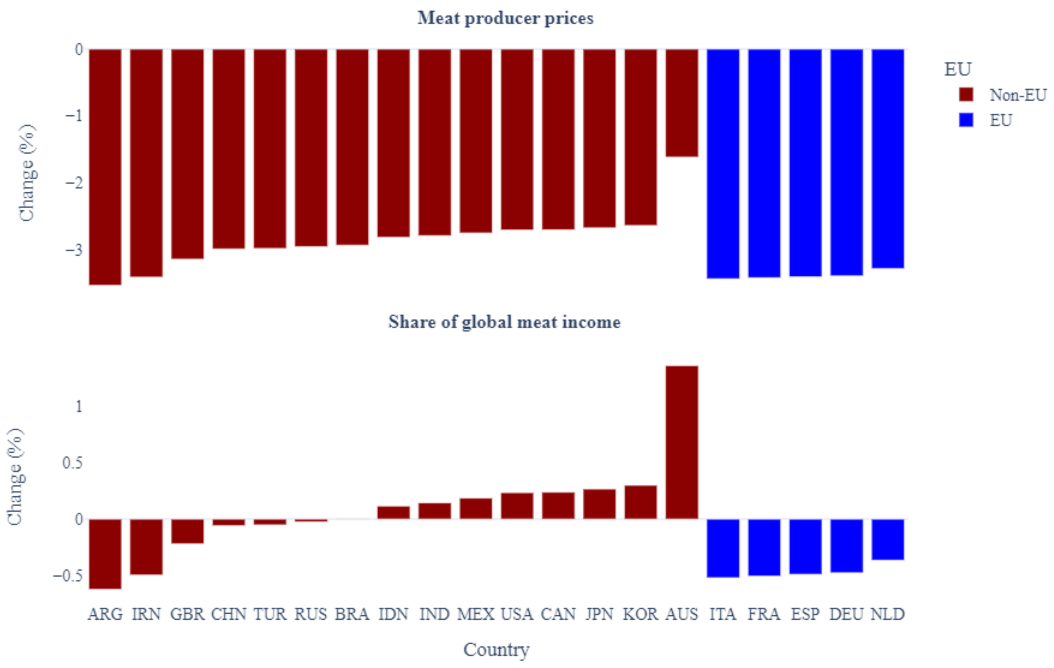

Notes: In both the baseline (pre-tax) and counterfactual scenarios the system of multilateral resistances is solved under the normalization that inward multilateral resistance in Australia, $P_{AUS} = 1$.

production quantities are fixed would suggest that global meat prices decrease in response to the tax. Indeed, the top panel of Figure 3.4 illustrates that every country experiences a decrease in meat producer prices due to the tax, ranging from a 1.5% to 3.5% fall in prices relative to the baseline (pre-tax) scenario. The structural gravity framework brings insight into the heterogeneity underlying this result: producers that are particularly close to the EU market (EU producers themselves, for example) tend to experience large price decreases. This result for the GETI of the tax follows intuitively from the result for the MTI of the tax. Meat producers around the world face increased competition due to the tax (OMR increases), which puts downward pressure on meat prices. Equation (3.9) illustrates this close link between outward multilateral resistance and producer prices in the model. Producers that see the largest increases in outward multilateral resistance due to the tax (Figure 3.1) also see the largest decreases in meat prices. These decreases in meat prices around the globe due to the EU meat tax demonstrate the risk that such a policy may lead to increased meat consumption outside the EU, and the reduction in carbon emissions associated with EU meat consumption may be offset to some extent by an increase in carbon emissions related to meat consumption outside the EU.

To further illustrate these risks, Figure 3.5 shows results for the changes in producer

Figure 3.5: GETI effects on producer prices: 1%, 5%, and 10% tax rates on EU meat consumption

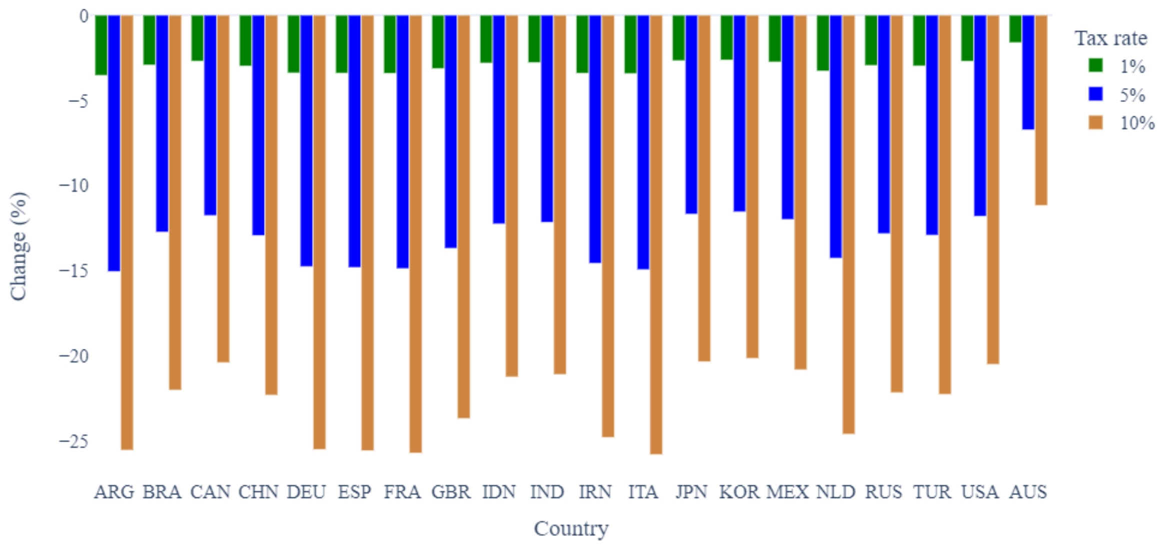

Notes: In both the baseline (pre-tax) and counterfactual scenarios the system of multilateral resistances is solved under the normalization that inward multilateral resistance in Australia, $P_{AUS} = 1$.

prices under 1%, 5%, and 10% ad valorem tax rates.[14] A higher tax rate leads to larger decreases in producer prices. Under a 5% tax, producer prices decrease by up to 15%, and under a 10% tax they decrease by up to 25%. Once again, the assumption of fixed production quantities in this model does not allow for a quantitative assessment of carbon leakage associated with a tax on EU meat consumption. Nevertheless, the magnitudes of these price decreases suggest that the market signals for the spatial redistribution of meat production and consumption away from the EU may be large, and therefore the risk of price-induced carbon leakage is non-negligible. This result aligns with findings in previous papers that global general equilibrium effects of unilateral carbon pricing policies may be a significant risk to the the environmental effectiveness of unilateral climate action.

Result 5: Global meat sector income shifts away from the EU orbit.

The bottom panel of Figure 3.4 shows the percentage change in each country's share of global meat sector income. Since the model holds underlying meat production quantities constant, these results reflect the impact of changes in producer prices

[14]For context, previous studies of meat taxation have found that quite high tax rates are necessary to internalize the environmental damage costs associated with meat. For example, Edjabou and Smed (2013) find price increases of 2% to 11% on meat products are necessary to internalize climate damages associated with meat consumption in Denmark; in the Swedish context, Säll and Gren (2015) account for damages associated with several pollutants (not just greenhouse gases) and find that taxes of 8.9% to 33% on meat products are necessary.

on relative economic size. Countries that experienced the largest producer price decreases see a decrease in their relative size, while countries whose meat producer prices are relatively resilient see an increase in relative size. A clear pattern emerges: meat sector income shifts away from the EU and countries that sell a lot to the EU and towards other large regional markets, namely North America and north Asia. The magnitudes of these economic size effects appear small, with less than a 1% change in shares of global meat income for most countries, but we should keep in mind that these countries are very large meat producers: in the baseline (pre-tax) scenario, a 1% share of global meat income corresponds to 5.65 billion USD. An important feature of this result is that the countries that grow most in terms of relative economic size (e.g. USA and Japan) are those that are relatively large and well-connected producers before the tax. Countries that are relatively small and remote (e.g. Iran, Russia) are not able to benefit strongly from this shift away from the EU.

These GETI effects of the tax on prices and relative economic size feed back into bilateral trade and the system of multilateral resistances. Economic size improves market access, allowing countries that see an increase in relative economic size improved trading opportunities, and vice versa for EU countries and others that see a decrease in relative economic size. In effect, the impacts of the tax on multilateral resistances via the economic size channel further reinforce the shift in meat trade away from the EU orbit.

The aim of the tax is to reduce meat production and consumption, thereby mitigating carbon emissions related to meat. However, this result highlights the challenge of unilateral climate policy: some of the economic activity that the EU is trying to regulate slips away from their regulatory grasp and moves away to other jurisdictions. This effect occurs even in the case of a non-discriminatory tax imposed on all meat consumption in the EU - it is a general equilibrium market adjustment that is very difficult to regulate unilaterally. Overall, the magnitudes of these GETI effects of a 1% tax in this stylized model are not large enough to suggest that they would completely undermine the aims of the policy, but they are enough to suggest we should not dismiss these risks.

3.5 Discussion and concluding remarks

This paper contributes to our evidence base on the global effects of a hypothetical tax on EU meat consumption using a structural gravity framework, which has not yet been used to assess this policy. The results for the modular trade impacts of the 1% tax on EU meat consumption illustrate how sales of meat in the EU are ‘destroyed’ and diverted towards non-EU markets. Consumers benefit from the increased competition amongst meat producers, enjoying increased accessibility of meat. Even EU consumers, despite facing a tax on meat consumption, maintain their relative position in the international trade network and benefit from their close ties to EU producers. The general equilibrium trade impacts of the tax simulation confirm that this increased global competition amongst meat producers leads to decreases in meat producer prices in all countries in the model. Producers in EU countries, the UK, and Argentina see some of the largest price decreases. Furthermore, the effects of the tax on relative economic size demonstrate that global meat income shifts away from the EU orbit, and so some of the economic activity that the policy aims to regulate moves to other jurisdictions. Large and well-connected countries in North America and north Asia benefit most from this shift. These results suggest that the decrease in emissions related to EU meat consumption in response to the tax may be offset to some extent by an increase in meat consumption elsewhere. Given the global nature of climate change, the market signals reflected in these GETI effects undermine the environmental effectiveness of the policy to some extent. The magnitudes of the estimates of these GETI effects suggest that we should not dismiss this risk of carbon leakage when designing policy approaches to tackle meat-related emissions.

An important caveat to keep in mind regarding these results is that the model is in nominal terms and holds underlying production quantities constant, so the analysis does not assess of the effectiveness of the tax in terms of GHG mitigation. The results for prices and relative economic size provide an indication of the global market adjustments that will occur in the meat sector in response to the tax in the very short term, but in the medium to long term these nominal adjustments may be tempered by shifts of factors of production away from the meat sector. If meat producers can readily shift to producing alternative products, then the increase in competition amongst meat producers and corresponding downward pressure on prices will not be so intense. This point highlights an important area for continued research: understanding the extent to which land, labour, and capital can readily shift away from the meat sector. An understanding of the global supply-side adjustments that might occur in response to this policy is important not only for understanding the efficiency and environmental effectiveness of the tax, but also to shed light on its global distributional impacts. Another important takeaway from this paper is that the global incidence of the tax falls on all meat producers, but especially EU producers. Previous papers on carbon taxes on food consumption in EU countries have studied distributional impacts on consumers, addressing concerns that such policies may be regressive. This paper complements this discussion by pointing out the relatively heavy burden of adjustment likely to fall on producers under a meat taxation policy, further highlighting the need for research to understand these supply-side impacts.

Overall, the analysis in this paper emphasizes the limitations of unilateral market-based policy to deal with a global environmental externality such as climate change. The EU cannot control these GETI effects that happen largely beyond its borders. Border-carbon adjustments and other market-based policies that target buying and selling of meat in the EU are irrelevant. This risk of counteractive GETI effects underscores the importance of continuing to pursue other avenues such as international climate agreements and trade agreements with conditions for environmental protections as potential complementary policies alongside meat taxation. Unilateral meat taxation in a large region such as the EU may go some way to reduce global meat consumption but it is not a silver bullet. The economic logic behind this insight remains the same for any region or country large enough to influence the global supply and demand equilibrium for meat, and so we would expect this result to maintain external validity if the taxing jurisdiction was not EU countries but some other 'climate club'.[15] Moreover, the larger the club (in terms of meat consumption), the more meat supply needs to be diverted elsewhere, and therefore the GETI effects of increased accessibility of meat outside the club are potentially larger.

Given the result that the tax makes meat more accessible outside the EU, this finding may raise the question of whether the policy could have potential synergies with food security in low income countries. In many high income countries, meat consumption is so high that it is associated with negative health outcomes. This health externality provides additional motivation (on top of the environmental motivation) for meat taxation (Springmann et al. 2016). On the other hand, a concern regarding meat taxation is that it could negatively impact food security and health outcomes if implemented in low income countries.[16] Taken together, these concerns suggest that a unilateral tax on meat consumption in high-income countries could possibly be preferable to the globally-cooperative approach of imposing a meat tax in all countries: the tax decreases meat consumption in countries that over-consume, and GETI effects potentially increase accessibility of meat in countries that suffer from poor food security and under-consume relative to optimal levels. The tax simulation results in this paper confirm that consumers in all countries may benefit from GETI effects of the policy to some extent. However, another feature of these results is that some countries benefit much more than others; those that tend to see the largest increase in their share of global meat income are economically large and

[15] The high level of integration of EU countries may be a unique factor behind the result that EU consumers do not face a decrease in relatively accessibility of meat due to the tax. Even after the tax is imposed EU consumers remain a highly preferred market in terms of trade costs for EU producers, and this ease of access may help to shield EU consumers from bearing any incidence of the tax.

[16] Springmann et al. (2016)'s simulation finds that this result is not very prevalent and only occurs for the very poorest countries in their model.

well-connected countries. Low income food-insecure countries likely do not have the trading connections and institutions necessary to be strong alternative markets when meat supply is diverted away from the taxing jurisdiction. Therefore the results of this paper do not imply that a tax on meat consumption in the EU may have a beneficial second-order effect of increasing food security in low income countries.

Another important point to note regarding the analysis in this paper is that it holds constant all other food and agricultural policies - in the EU as well as non-EU countries. In particular, reducing agricultural subsidies may be another approach (as an alternative to meat taxation) to bring the price of meat closer to its true cost (Funke et al. 2021). Using this structural gravity framework to explore the potential consequences of removing or adjusting these subsidies compared to layering a meat consumption tax on top of the production subsidies could be a fruitful area for future research. In general, understanding how best to align agricultural and climate policies aimed at reducing meat consumption remains an important area for policy-relevant research.

The results in this paper confirm findings in previous papers that unilateral climate policy may suffer from counteractive GETI effects. In the face of a global externality such as climate change, the sub-optimality of unilateral policy comes as no surprise. However, given the immense obstacles to coordinated global climate policy, research is needed to understand the magnitude of the risks posed by the GETI effects of unilateral policy. This paper is one of the first to address this gap in the context of meat taxation, and does so using a framework that has not been applied to this topic. Future research can build on these results, potentially estimating these GETI risks in extended structural gravity frameworks that allow for long term adjustments in consumption and production.

3.A Appendix to Chapter 3

3.A.1 Deriving the bilateral trade equation

Consumers in country j maximize utility (equation 3.1) subject to the budget constraint (equation 3.4). The Langrangian is:

$$\mathcal{L} = \left[\left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{\eta_j} O_j^{1-\eta_j} - \lambda \left(\sum_i p_{ij} q_{ij} - E_j \right)$$

$n + 1$ first order conditions:

$$\frac{\partial \mathcal{L}}{\partial q_{ij}} = \eta_j \left[\left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{\eta_j - 1} O_j^{1-\eta_j} \frac{\sigma - 1}{\sigma} \left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{\sigma-1}} \frac{\sigma - 1}{\sigma} \beta_j^{\frac{1-\sigma}{\sigma}} q_{ij}^{-1/\sigma} - \lambda p_{ij}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \sum_i p_{ij} q_{ij} - E_j$$

Take the ratio of the first order conditions with respect to varieties q_{ij} and q_{mj}:

$$\begin{aligned} \frac{\frac{\partial \mathcal{L}}{\partial q_{mj}}}{\frac{\partial \mathcal{L}}{\partial q_{ij}}} &= \frac{\eta_j \left[\left(\sum_m \beta_m^{\frac{1-\sigma}{\sigma}} q_{mj}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{\eta_j - 1} O_j^{1-\eta_j} \left(\sum_m \beta_m^{\frac{1-\sigma}{\sigma}} q_{mj}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{\sigma-1}} \beta_j^{\frac{1-\sigma}{\sigma}} q_{mj}^{-1/\sigma} - \lambda p_{mj}}{\eta_j \left[\left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{\eta_j - 1} O_j^{1-\eta_j} \left(\sum_i \beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{\sigma-1}} \beta_j^{\frac{1-\sigma}{\sigma}} q_{ij}^{-1/\sigma} - \lambda p_{ij}} \\ &= 0 \end{aligned}$$

Simplify this expression:

$$\begin{aligned} \frac{\beta_m^{\frac{1-\sigma}{\sigma}} q_{mj}^{-1/\sigma}}{\beta_i^{\frac{1-\sigma}{\sigma}} q_{ij}^{-1/\sigma}} &= \frac{p_{mj}}{p_{ij}} \\ \left(\frac{q_{mj}}{q_{ij}} \right)^{-1/\sigma} &= \frac{p_{mj}}{p_{ij}} \left(\frac{\beta_m}{\beta_i} \right)^{\frac{\sigma-1}{\sigma}} \\ q_{mj} &= q_{ij} \left(\frac{p_{mj}}{p_{ij}} \right)^{-\sigma} \left(\frac{\beta_m}{\beta_i} \right)^{1-\sigma} \end{aligned}$$

Multiply by p_{mj} and sum over all varieties:

$$\sum_m p_{mj} q_{mj} = \sum_m p_{mj}^{1-\sigma} q_{ij} p_{ij}^{\sigma} \left(\frac{\beta_m}{\beta_i} \right)^{1-\sigma}$$

Substitute in the budget constraint, $E_j = \sum_m p_{mj} q_{mj}$, and rearrange to obtain country j's Marshallian demand for country i's variety of meat:

$$\begin{aligned} E_j &= q_{ij} p_{ij}^{\sigma} \left(\frac{1}{\beta_i} \right)^{1-\sigma} \sum_m (\beta_m p_{mj})^{1-\sigma} \\ q_{ij} &= \frac{p_{ij}^{-\sigma} \beta_i^{1-\sigma} E_j}{\sum_m (\beta_m p_{mj})^{1-\sigma}} \end{aligned}$$

The value of meat exports sold by exporter i to importer j is given by: $X_{ij} = p_i t_{ij} q_{ij}$. Substitute the above expression for Marshallian demand for q_{ij} into this equation, as well as the CES price index, P_j (equation 3.2):

$$\begin{aligned} X_{ij} &= p_i t_{ij} (p_{ij} t_{ij} \tau_j)^{-\sigma} \left(\frac{\beta_i}{P_j}\right)^{1-\sigma} E_j \\ X_{ij} &= \tau_j^{-\sigma} \left(\frac{\beta_i p_i t_{ij}}{P_j}\right)^{1-\sigma} E_j \end{aligned}$$

To obtain the equations for multilateral resistances and obtain the expression for bilateral trade given by equation 3.6, start with the market clearance assumption and substitute the above expression for bilateral trade:

$$\begin{aligned} Y_i &= \sum_j X_{ij} \\ Y_i &= \sum_j \tau_j^{-\sigma} \left(\frac{\beta_i p_i t_{ij}}{P_j}\right)^{1-\sigma} E_j \end{aligned}$$

Divide this expression by global meat income, $Y = \sum_i Y_i$, and rearrange:

$$\begin{aligned} \frac{Y_i}{Y} &= \sum_j \tau_j^{-\sigma} \left(\frac{\beta_i p_i t_{ij}}{P_j}\right)^{1-\sigma} \frac{E_j}{Y} \\ \frac{Y_i}{Y} &= (\beta_i p_i)^{1-\sigma} \sum_j \tau_j^{-\sigma} \left(\frac{t_{ij}}{P_j}\right)^{1-\sigma} \frac{E_j}{Y} \\ (\beta_i p_i)^{1-\sigma} &= \frac{Y_i/Y}{\sum_j \tau_j^{-\sigma} \left(\frac{t_{ij}}{P_j}\right)^{1-\sigma} \frac{E_j}{Y}} \end{aligned}$$

Define $\Pi_i^{1-\sigma} = \sum_j \tau_j^{-\sigma} \left(\frac{t_{ij}}{P_j}\right)^{1-\sigma} \frac{E_j}{Y}$:

$$(\beta_i p_i)^{1-\sigma} = \frac{Y_i/Y}{\Pi_i^{1-\sigma}} \tag{3.11}$$

Substitute this expression into the above equation for bilateral trade:

$$\begin{aligned} X_{ij} &= \tau_j^{-\sigma} \left(\frac{t_{ij}}{P_j}\right)^{1-\sigma} \left(\frac{Y_i/Y}{\Pi_i^{1-\sigma}}\right) E_j \\ X_{ij} &= \frac{Y_i E_j}{Y} \left(\frac{t_{ij}}{\Pi_i P_j}\right)^{1-\sigma} \tau_j^{-\sigma} \end{aligned}$$

The above expression is the expression for bilateral trade, equation 3.6.

To obtain the expression for inward multilateral resistance as given by equation 3.8, substitute 3.11 into the CES price index:

$$P_j = \left[\sum_i \frac{Y_i/Y}{\Pi_i^{1-\sigma}} (\tau_j t_{ij})^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$$

$$P_j = \left[\tau_j^{1-\sigma} \sum_i \left(\frac{t_{ij}}{\Pi_i} \right)^{1-\sigma} \frac{Y_i}{Y} \right]^{\frac{1}{1-\sigma}}$$

Finally, to obtain the expression for producer prices (equation 3.9), rearrange 3.11:

$$p_i^{1-\sigma} = \frac{Y_i/Y}{(\beta_i \Pi_i)^{1-\sigma}}$$

$$p_i = \frac{(Y_i/Y)^{\frac{1}{1-\sigma}}}{\beta_i \Pi_i}$$

3.A.2 Detailed results

Table 3.3: Detailed results: 1% tax on EU meat consumption

		MTI effects		GETI effects			
		IMR	OMR	IMR	OMR	Producer prices	Output share
		P_i	Π_i	P_i	Π_i	p_i	Y_i/Y
EU	DEU	-8.31	10.23	-6.40	7.96	-7.32	-0.23
	ESP	-8.23	10.07	-6.35	7.89	-7.27	-0.17
	FRA	-8.49	10.46	-6.39	7.97	-7.33	-0.23
	ITA	-8.62	10.68	-6.42	8.03	-7.37	-0.27
	NLD	-7.87	10.06	-6.39	7.85	-7.24	-0.14
Non-EU	ARG	-8.38	9.15	-8.01	8.98	-8.02	-0.97
	AUS	0.00	5.66	0.00	5.21	-5.39	1.86
	BRA	-7.97	8.66	-7.37	8.03	-7.37	-0.27
	CAN	-7.02	7.55	-6.61	6.96	-6.63	0.52
	CHN	-7.69	8.33	-7.15	7.72	-7.16	-0.05
	GBR	-9.03	9.96	-7.22	7.71	-7.15	-0.04
	IDN	-5.78	6.13	-6.78	7.22	-6.81	0.32
	IND	-6.20	6.61	-6.65	6.99	-6.65	0.50
	IRN	-7.94	8.62	-8.03	9.02	-8.05	-1.01
	JPN	-7.00	7.53	-6.55	6.87	-6.56	0.59
	KOR	-6.89	7.40	-6.43	6.73	-6.47	0.69
	MEX	-7.04	7.58	-6.64	6.98	-6.64	0.51
	RUS	-8.22	8.96	-7.05	7.55	-7.04	0.08
	TUR	-8.42	9.20	-7.10	7.64	-7.10	0.01
	USA	-6.98	7.56	-6.57	6.93	-6.61	0.54

Notes: All results expressed in terms of the percentage change under the tax compared to the baseline (pre-tax) scenario. The modular trade impact (MTI) effects hold prices and income constant but allow trade frictions to adjust to the tax, and general equilibrium trade impact (GETI) effects allow prices and income to adjust as well. IMR denotes inward multilateral resistance and OMR denotes outward multilateral resistance.

Chapter 4

Exploring trends in innovation and demand growth for plant-based substitute products in the US[1]

4.1 Introduction

Recent studies have found that the food industry accounts for a quarter to a third of total global CO2-eq. emissions (Crippa et al. 2021; Poore and Nemecek 2018), and a very large portion of these emissions are attributable to animal-based food products. Poore and Nemecek (2018) find that livestock and fish farms are responsible for 31% of emissions from food, with an additional 16% of food emissions coming from land use for livestock and 6% from producing crops for livestock feed. Moreover, many studies of potential pathways to decarbonise the food sector have found that dietary shifts towards reduced meat consumption plays an important role in this transition (Hedenus et al. 2014; Springmann et al. 2018). In particular, the IPCC Special Report on Climate Change and Land highlights the strong mitigation potential offered by dietary change towards reduced consumption of animal products, and notes the potential role that plant-based substitutes for meat could play in this dietary shift (Mbow et al. 2019).

Meanwhile, media outlets and reports from the private sector in recent years have painted a picture of a flourishing market for plant-based meat substitutes that is ripe with investment opportunities (Bashi et al. 2019); for example, the FAIRR Initiative reported a 300% year-on-year increase in private investment flowing to

[1]Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

alternative proteins in 2020 compared to 2019, though they note that investment levels are still very low compared to investment in electric vehicles (FAIRR Initiative 2021). The ramping up of private investment in these foods (with public funding lagging behind) comes somewhat as a surprise given our understanding of the green innovation process in other sectors such as energy and transport, which have relied heavily on policy intervention (Grubb et al. 2021). Media and private sector reports suggest that consumer-driven change is spontaneously driving green innovation in the food sector, despite a lack of policy intervention. Accordingly, the food sector is an important subject of study for understanding green innovation process because - if stories in the media are correct - this sector is not following our usual understandings of this process.

The aim of this paper is to use rich microdata on US food sales to empirically assess the validity of these popular stories about plant-based substitute products. Our research question is twofold: (i) is the market for plant-based substitutes rapidly expanding, and (ii) is demand an important driver of this expansion. We tackle the first part of this research question by examining and comparing descriptive statistics across time, food groups, and geography. To do so, we start by finding the plant-based substitute products in the data and then categorizing products into custom food groups. Then we examine US-level trends in spending and innovation across these food groups. Our analysis suggests that plant-based substitute products are not gaining traction at an aggregate US-level. Contrary to the picture often painted in the media, we conclude that the market for these products is not trending in such a way as to play a meaningful role in decarbonizing the food sector. Next, we examine geographical heterogeneity in these statistics, and find some evidence that underneath these lacklustre aggregate statistics may be some pockets of rapid expansion in the plant-based substitutes market.

To address the second part of our research question, we undertake some exploratory analysis of the link between demand and innovation across our food groups. First, we examine the non-causal association between store-level product entries and local spending across our food groups, finding that this correlation may be particularly strong for plant-based substitute products. Next, we do a household-level analysis to discover whether any household characteristics are particularly strongly associated with high demand for any of our food groups. We find that young and highly educated households are particularly dominant in the market for plant-based meat substitutes. Finally, we bring together the store-level and household-level analyses, and using data on US demographic trends we assess the potential that demand growth for plant-based meat substitutes has been exogenous to the innovation process. The aggregate US demographic trend towards increased education levels provides some suggestive evidence of exogenous demand growth and therefore that demand may be driving innovation in plant-based meat substitutes rather than vice versa. Overall, although we find some weakly suggestive evidence that demand and innovation for plant-based substitute products are closely linked, the main takeaway from this exploratory analysis is that these products are still a very niche product group. Future work can build on theses analyses and move towards a causal analysis of the link between demand and innovation, but this exploratory analysis suggests that popular stories about these products likely overstate that extent to which demand is flourishing and driving innovation. Given what we know from the literature on environmental innovation about the importance of policy intervention to drive green transitions in other sectors of the economy, these findings suggest that policy intervention may be necessary to spur a widespread dietary shift towards reduced meat consumption.

This paper relates to the economics literature on the role of demand in driving markets and consumer product innovation. For example, Acemoglu and Linn (2004) study innovation in pharmaceutical drugs. They construct an instrument for demand for specific drug categories based on shifts in population size and incomes of different age groups in the US and find that potential market size has a sizeable impact on introductions of new drugs. DellaVigna and Pollet (2007) use demographics to generate forecasts for demand shifts for different categories of consumer goods; they find that these demand shifts are strong predictors of industry returns. Jaravel (2019) studies inflation inequality in consumer products in the US. He finds that finds that relatively high demand growth for products purchased by high income households leads lower inflation for these products compared to products that low income households tend to purchase. Furthermore, he presents evidence that innovation (defined as an increase in product variety), driven by demand growth, is a key mechanism behind this inflation inequality. This paper contributes to this literature by considering consider consumer product innovation in the context of pathways towards climate change mitigation.

The literature on the environmental economics of the food sector has studied the potential to mitigate GHG emissions associated with food through carbon pricing (Säll and Gren 2015; Abadie et al. 2016; Bonnet et al. 2018), as well as dietary change scenarios such as widespread vegetarianism (Simon 2020; Springmann et al. 2018). However, studies in this area have largely taken supply and demand parameters as given, assuming underlying demand parameters and product offerings remain static. As discussed above, the food sector is a dynamic industry with innovative new substitutes for animal products offering a potential pathway towards decarbonization. However, the potential for such innovation to play a role in mitigating food sector emissions remains relatively under-analyzed in this literature.

Finally, this paper also relates to the literature on environmental innovation. Studies on this topic have focused heavily on the role of government intervention in driving the innovation process, specifically 'technology push' and 'demand-pull' policies as mechanisms to correct the market failures that lead to the under-provision of green innovation (Popp 2019; Grubb et al. 2021). In general, this literature has paid relatively little attention to consumers themselves as the genesis of demand change and a driver of innovation. Exceptions include Jens Horbach, Christian Rammer, and Klaus Rennings (2011) and Veugelers (2012) who study green innovation using firm survey data from Europe and identify consumer demand as a potential driver of this innovation. Meanwhile, Popp et al. (2007) conclude that consumer pressure alone is not enough to drive sustained green innovation in the pulp and paper sector - complementary government policy is needed. Overall, the literature on demand as a driver of green innovation remains sparse, and our understanding of green innovation processes in the food sector, where policy intervention has been minimal, is also under-explored.

This paper makes several important contributions. First, we develop an approach to find plant-based substitute products in the data and construct customized food groups for our analysis. Second, we use a data-driven approach to interrogate popular understandings of the plant-based substitute market. Third, we offer some exploratory insights into the green innovation process in the food sector. Finally, our analysis contributes evidence in support of climate policy intervention in the food sector.

The remainder of the paper is structured as follows. The next section describes or data sources and our approach to classifying products into our custom food groups. Then we describe our main variables of interest and discuss trends in the markets for our food groups at an aggregate US-level. The following sections present our household-level analysis and then store-level analysis. We offer a brief discussion of insights from demographic trends in light of our findings, and then the final section concludes.

4.2 Data description

4.2.1 Data sources

Our main data sources for this analysis are from the The Kilts Center Archive of the Nielsen Company. We use two microdata sets from this source, one for data on household purchases and the other for data on retailer sales.

Nielsen Retail Scanner Dataset. This data source is a panel of weekly product-level sales at individual stores in the US, covering 30000-50000 stores each year and representing around 90 retail chains. Products are observed at the level of unique Universal Product Codes (UPCs). As well as UPC-level information (such as price and quantity sold as well as the product brand), we also observe geographic information about the particular store.

Nielsen Consumer Panel Dataset. This data source is panel of UPC-level purchases by about 60000 US households each year. Some households are only in the sample for one year, while others remain for multiple years. We observe annual demographic and geographic data about each household, and for each shopping trip we observe data on the store as well as each product purchased. Nielsen recruits households to this sample using proportionate stratified sampling techniques to maintain correspondence with US demographics. They provide sampling weights ("projection factors") to enable scalability of purchases to the aggregate US level. An important spatial dimension in this data are Scantrack market areas, which are Nielsen-defined geographical areas that correspond roughly to 54 major metropolitan areas - e.g. New York, Kansas City, Sacramento. These areas often cross state boundaries but not county boundaries.

US Bureau of Labour Statistics. We deflate nominal spending measures constructed from the Nielsen data into 2010 dollar terms using consumer price index (CPI) data from the US Bureau of Labour Statistics. In particular, we use the series "Food in U.S. city average, all urban consumers, not seasonally adjusted" and take the simple average of monthly CPI observations for annual CPI.

American Community Survey (5-year). We use demographic data to explore the potential for demand growth for plant-based meat substitutes via population and income growth. To do so, we construct demographic trends from the 5-year American Community Survey (ACS) Public Use Microdata Sample (PUMS). Focusing on the 'reference person' for each household in the ACS PUMS data, we construct household populations (as well as inflation-adjusted household income) based on 4 characteristics: age, education, race, and the presence of children.

4.2.2 Food group definitions

Nielsen classifies individual UPCs (products) according to a multi-tiered system comprising of departments, product groups, and product modules. Nevertheless, plant-based substitutes for animal-based products are not easily located within this hierarchy using product module and group names. In many cases, these products are in the same modules as animal products or some other products. To deal with this issue, we have developed customized rules to allocate UPCs to food groups. In particular, we look for specific strings - e.g. 'tofu', 'meatless', and 'mtls' - in the abbreviated UPC descriptions of products within specific groups and modules that we expect to contain plant-based substitutes for animal products.

To determine the appropriate Nielsen modules in which to look for plant-based substitutes as well as the strings that can identify these products, we use data from Label Insight. Label Insight offers a detailed cross-section of UPC-level product attributes, as well as a granular classification of products that includes categories specifically for plant-based versions of animal products. From Label Insight we obtain UPC codes of products with a particular attribute we're interested in (e.g. tofu products; plant-based yogurt products); then we search for these UPC codes in the Nielsen data and manually identify patterns in Nielsen's classification and abbreviated descriptions of these products. We define our rules to allocate products to food groups based on these patterns.

We are interested in demand for plant-based substitutes for animal products, so we restrict our analysis to these products as well as their animal product counterparts. We have six food groups in total:

Meat: The meat food group includes all types of meat - e.g. poultry, pork, beef - in various different forms, such as frozen, canned, and refrigerated. We also include meat-based frozen entrees. The Nielsen data includes some meat products that do not have a standard Universal Product Code - they refer to these types of products as "magnet data". We do not include magnet data in our analysis because not all households track purchases of these products and so the sample size is limited. Only the meat and seafood groups are affected by dropping magnet data.

Seafood: Similar to the meat food group, the seafood group includes fresh, frozen, and canned seafood products, as well as seafood-based frozen entrees. We do not include seafood products without a standard UPC (i.e. "magnet data").

Dairy (and eggs): The dairy and eggs food group includes milk, yogurt, and eggs. We combine eggs and dairy because eggs are a very small category, and only a few products exist that are meant to be plant-based substitutes for eggs. We do not include cheese and butter products to maintain comparability with the plant-based dairy substitute food group (which does not include plant-based cheese and butter).

Milk products include refrigerated as well as shelf-stable and powdered milk.

Plant-based meat substitutes: This category is quite broad, spanning traditional vegetarian protein sources such as tofu and tempeh, to bean burgers, to plant-based burgers that aim to closely imitate the taste and texture of meat. The availability of plant-based products meant to imitate seafood is limited, so we do not attempt to separate such products into their own group. In general, although this food group is quite diverse, disaggregating it into subgroups is challenging given the limited information on product attributes in the Nielsen data.

Plant-based dairy substitutes: This category includes plant-based substitutes for milk, yogurt, and eggs. We include both refrigerated and shelf-stable plant-based milks, though the latter is particularly challenging to identify within the Nielsen data. We do not include cheese products in this group due to the difficulty of separately identifying truly plant-based cheese substitute products from "imitation cheese" products (which contain milk derivatives).

Pulses: This category includes dried and canned lentils, beans, split peas, chickpeas, and dal. Due to their high protein content, these products are common components of plant-based diets.

The identification of plant-based substitute products within the Nielsen data is an important contribution of this paper, and a necessary step towards improving our understanding of this market. Nevertheless, our method for identifying these products is not perfect and likely misses some relevant plant-based products. Future work could try to improve on these classifications, perhaps by using machine learning techniques to identify product groups. Finally, an important feature of our method to find these products in the data is that it is not brand-dependent, which ensures that our analysis does not simply reflect the success of a small handful of superstar brands.

4.3 Main variables and summary statistics

Table 4.1 summarizes annual observations of our key variables at the US-level, and Figure 4.1 illustrates trends in several of these variables over 2011 to 2019. We construct these statistics from the Retail Scanner Data. Our first variable of interest in this study is food group-specific real spending. To construct this variable, we aggregate weekly store-level sales for each food group, and then use consumer price index data from the US Bureau of Labour Statistics to deflate this measure into 2010 dollar terms.

Table 4.1: US-level summary statistics, Averages of annual observations 2011-2019

	Dairy	Meat	Plant-based dairy sub.	Plant-based meat sub.	Pulses	Seafood
Real spending (millions)	269.02	5853.71	153.14	382.08	312.25	677.64
	(16.10)	(291.14)	(15.85)	(37.20)	(10.46)	(78.19)
Number of products	469.00	10480.67	463.00	917.78	1965.78	2046.00
	(71.78)	(467.21)	(27.71)	(44.33)	(63.83)	(81.85)
Number of new products	70.67	1526.44	52.67	112.44	150.00	176.00
	(31.90)	(212.44)	(14.88)	(26.33)	(36.98)	(35.66)
Number of brands	140.00	1646.89	181.56	241.22	346.00	401.44
	(23.28)	(75.26)	(9.33)	(17.01)	(27.70)	(25.96)
Avg. number of products per brand	3.36	6.37	2.56	3.82	5.71	5.12
	(0.20)	(0.14)	(0.22)	(0.34)	(0.40)	(0.42)

Notes: Standard deviation in parentheses. Statistics constructed from the Nielsen Retail Scanner Data.

Our second key variable of interest in this study measures food group-specific innovation. We measure of innovation throughout this analysis as the count of new products introduced within a given food group. We determine the first year each product is observed in the Retail Scanner Data, and then count the number of products first observed in a given year. For versions of this measure that vary geographically, a product is defined as new when it first appears in a given location, regardless of whether it has been introduced already in other locations. This approach allows us to account for spatial heterogeneity in the timing of product entries, capturing the diffusion of product innovations over space as well as time. A single product is defined by a UPC (Universal Product Code).[2] UPCs are defined at a very granular level within the product space. For example, each flavour of the same type of yogurt has its own UPC, as do different package sizes for the same item (e.g. a pack of two burgers and a pack of four of the same burgers have different UPCs). Our measure of innovation therefore captures product diversification and differentiation on a granular level, from unique new product inventions to relatively minor changes to existing products such as new flavours or new package sizes.

[2]Occasionally UPCs from old products that no longer exist are recycled and assigned to a new product. Nielsen tracks these reassignments on an annual basis and provides a version number for each UPC, so we define unique products based on the combination of the UPC and this version number.

Table 4.1 highlights the notable differences in the sizes of our food groups. Moreover, the size of a food group in terms of spending levels correlates strongly with the total number of products, the number of new products, the number of brands, and the number of products per brand for the food group. This result is not surprising: a large demand pool offers a strong potential for firms to profit by segmenting the market through granular product differentiation. In a similar vein, Acemoglu and Linn (2004) find that demand levels are strong determinants of introductions of new pharmaceutical drugs. In our context, the meat food group in particular is very large relative to the other food groups according to both spending levels and number of products. On the other hand, the plant-based substitute food groups are relatively small according to these measures. The average number of meat products in a given year is over ten times the average number of plant-based meat substitute products, highlighting them immense gap that has yet to be bridged if these substitute products are going to play a meaningful role in decarbonizing the food sector. The number of products in the plant-based dairy substitute category is on par with its animal-based counterpart food group, but spending levels are still notably higher for conventional compared to plant-based dairy products.

Figure 4.1 plots trends over time for our main variables for each food group. Despite a relatively high rate of product entries for its spending levels (panel (a)), the total number of products in the plant-based meat substitute food group has not increased (panel (b)), suggesting that product introductions in this food group have met with mixed success. Similarly, panel (d) illustrates that this food group has seen relatively volatile spending growth compared to other food groups. By comparison, the plant-based dairy substitute food group has seen more consistent spending growth, but relatively little expansion in the number of products available. Meanwhile, the meat product group has seen notable increases in the number of products available, even following a decline in spending levels after 2015, and product numbers in the seafood group have been relatively robust in light of the decline in its spending levels over the period.

Overall, the trends depicted in Figure 4.1 reveal that plant-based substitute products are not gaining traction in the US market on an aggregate level. The market for these products is not trending in such a way as to play a meaningful role in decarbonizing the food sector. This finding is important to establish given the hype in mainstream media and private sector reports about the potential offered by these products. Moreover, innovation rates in large and well-established food groups such as meat and seafood seem relatively robust compared to innovation in the plant-based substitute food groups.

Figure 4.1: US-level summary statistics

(a) Number of new products

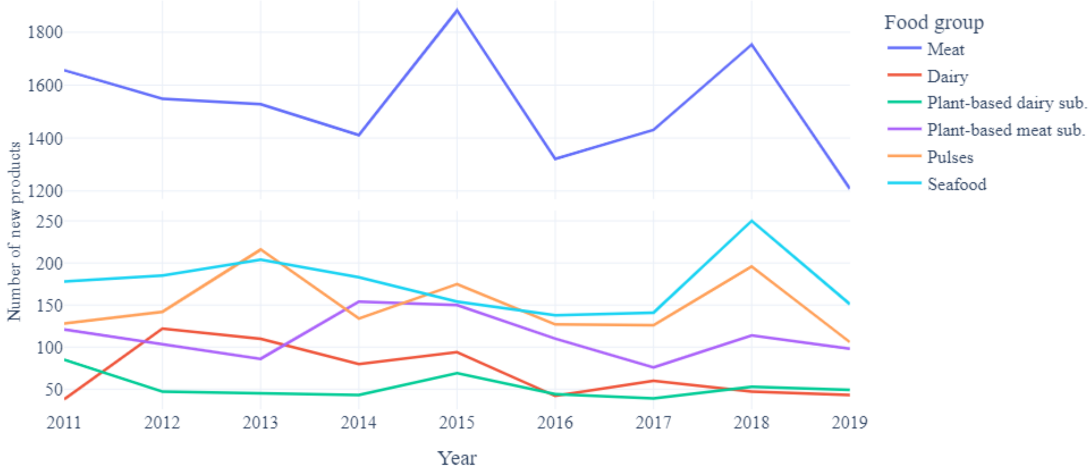

(b) Total number of products

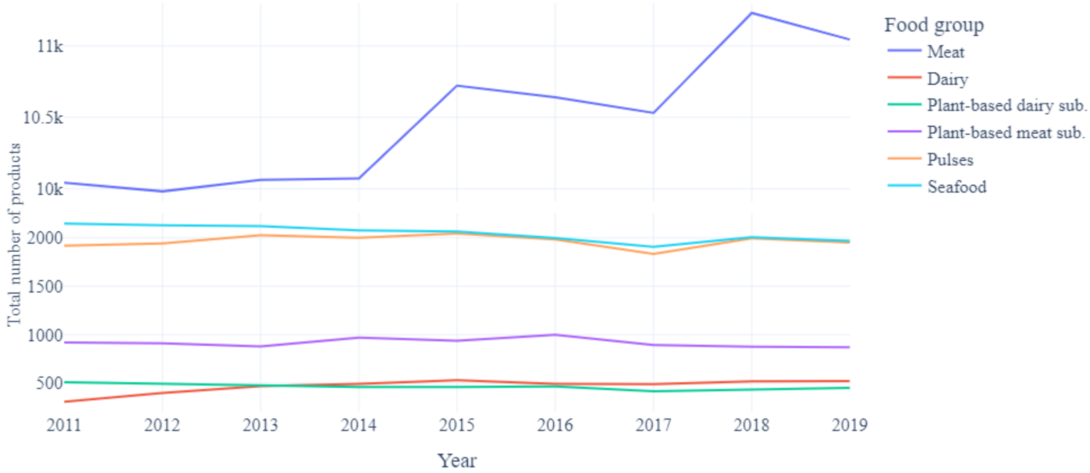

(c) Real spending levels

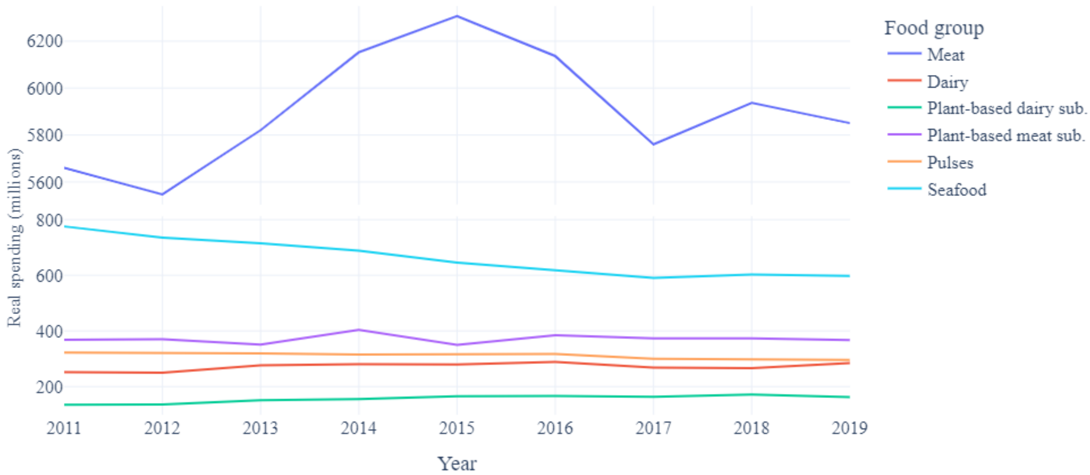

(d) Real spending growth

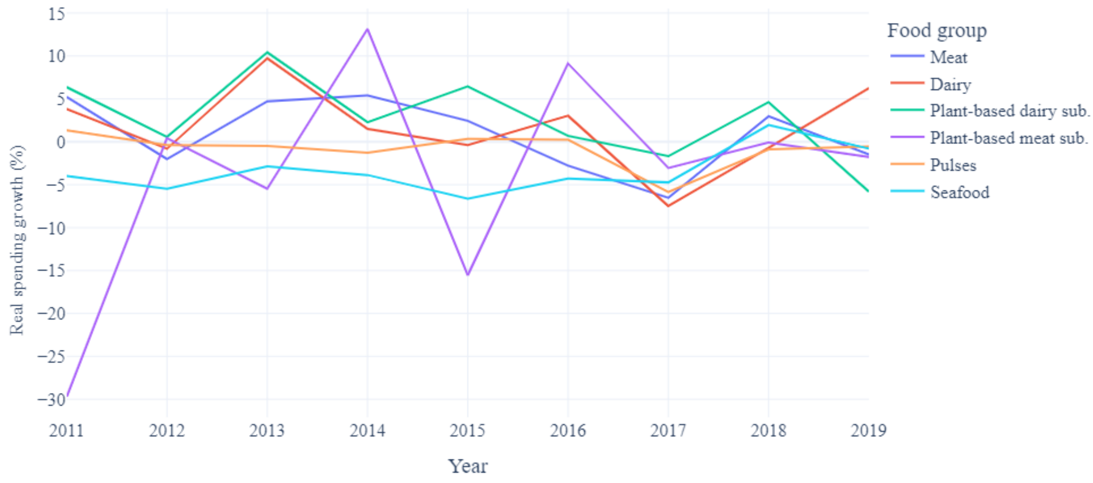

These aggregate US-level statistics may hide spatial heterogeneity within the US market for plant-based substitute products. To understand if this may be the case, we examine the distribution of spending and innovation across granular spatial units. Figure 4.2 plots the distribution of county-level spending growth by food group, revealing that the aggregate US trends discussed above do indeed mask significant spatial heterogeneity. In particular, real spending growth is very heterogeneous for plant-based meat substitutes compared to other food groups, while the plant-based dairy substitute group is relatively strongly skewed towards positive growth rates. Figure 4.9 (in the appendix of this chapter) plots the distribution of store-level product entries by food group. For most food groups, this distribution is quite skewed with a long tail; however, the distribution for plant-based meat substitute product entries is particularly skewed, once again suggesting that heterogeneity is particularly strong for this food group. Overall, these distributions provide some suggestive evidence that the market plant-based substitutes may show signs of acceleration within some pockets of the aggregate US market. In the following section, we take some steps to try to understand the features of these pockets of high growth in plant-based substitutes, and in particular we explore the potential that demand could be an important driver of this expansion.

4.4 Linking demand and innovation

As mentioned, a popular story in the media is that growing demand for plant-based substitutes for animal products is driving a flurry of innovation in this product space. Previous papers from the economics literature that have studied innovation in consumer goods have found that spending levels (Acemoglu and Linn 2004) and spending growth (Jaravel 2019) are important drivers of new product introductions.

Figure 4.2: Box plots of county-level annualized real spending growth by food group, 2010-2019

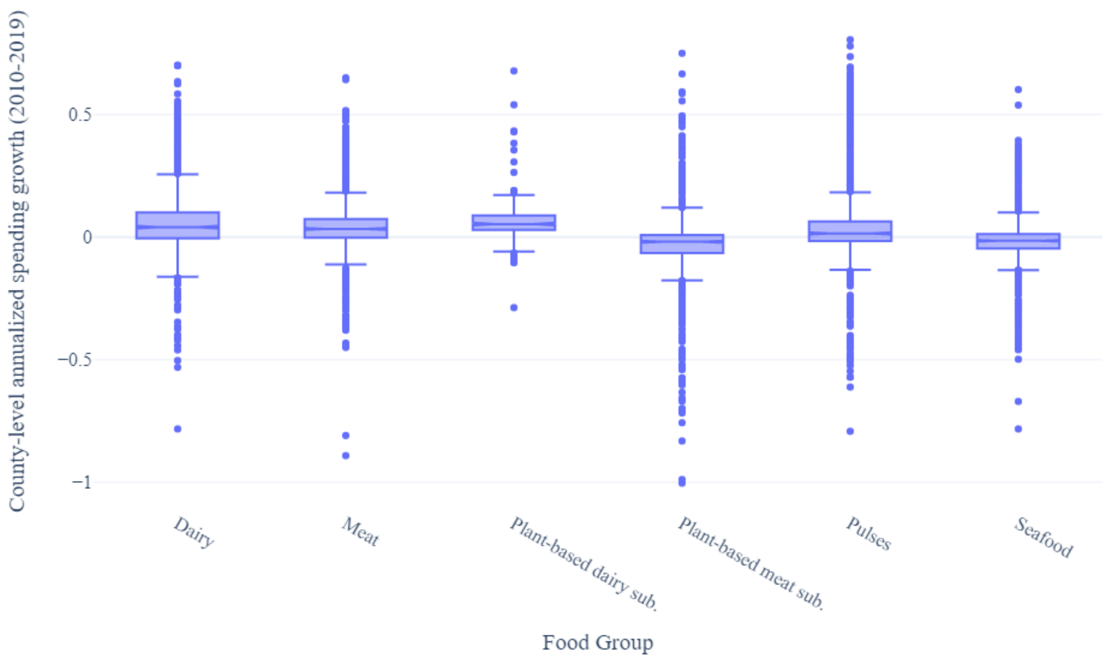

Nevertheless, in our context of green innovation, this story is somewhat surprising given the lack of policy intervention, which has been an important driver of the green innovation process in other sectors. Therefore the aim of this sections is to interrogate popular stories about the role of demand in driving green food innovation and in doing so improve our understanding of the green innovation process in the food sector.

4.4.1 Store-level analysis

We begin this exploratory analysis of the potential link between demand and innovation simply by examining the correlations between local spending and store-level product entries across our food groups. We do so by estimating specifications similar to the following using the PPML estimator:

$$N_{icf} = exp(\alpha + \boldsymbol{\beta}(Q_{if} \times \mu_f) + \mu_f + \gamma S_i + \eta_m + \varepsilon_{icf}) \tag{4.1}$$

N_{icf} is our measure of innovation, which in this context is the average annual number of new products introduced for food group f in store i in county c over the period 2011 to 2019. μ_f is a food group fixed effect, which controls for systematic differences in innovation rates between food groups. S_i is the average annual total nominal sales of store i in millions USD, to control for the likelihood that larger stores have a higher rate of product entries. η_m is a Scantrack market area fixed

Figure 4.3: Relationships between store-level product entries and real spending growth, 2010-2019

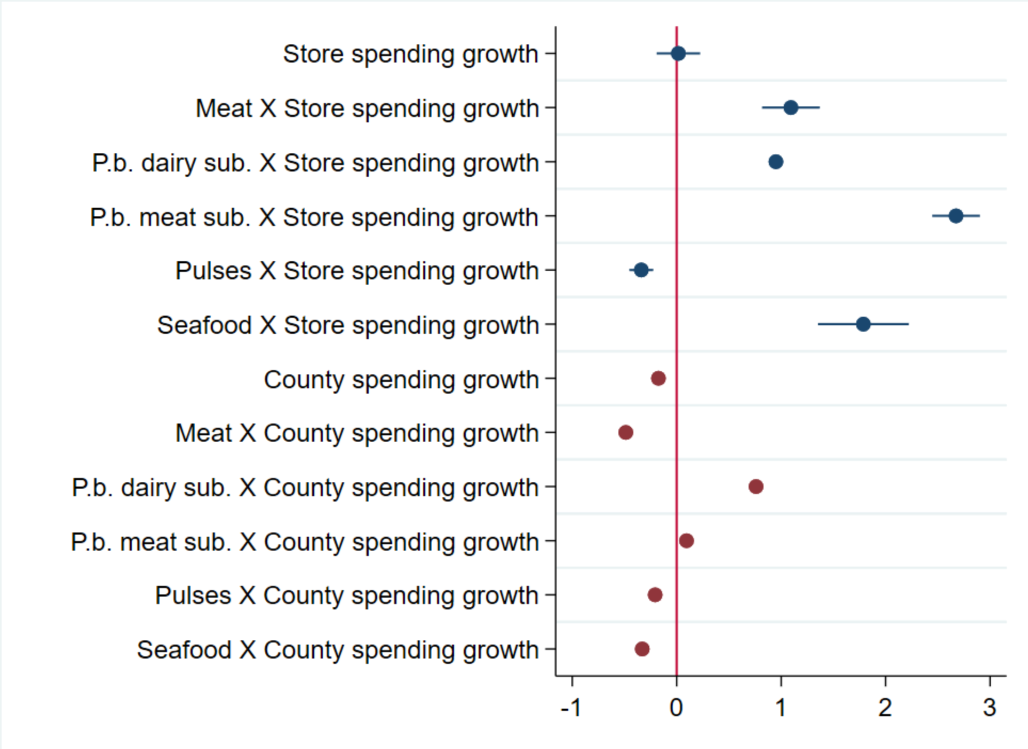

Notes: This figure depicts coefficient estimates for versions of Equation 4.1 that consider spending growth as the variable of interest.

effect, which controls for the average rate of product entries in a market area across all food groups; our identifying variation therefore comes from differences between stores within the same market area.

Our variables of interest are the average annual real spending level, Q_{if}, and annualized real spending growth, $\Delta log(Q_{if})$, for food group f at store i. We also explore the relationship between store-level product entries and the county-level analogues of these variables, Q_{cf} and $\Delta log(Q_{cf})$. $\boldsymbol{\beta}$ is a vector of 6 parameters that represent the food group-specific relationship between spending level (or growth) and product entries for each of the 6 food groups. We expect the relationship between spending and product entries to be highly endogenous. Increased product entries may spur an increase in spending, so the direction of causality can run both ways, and the coefficient estimates for the $\boldsymbol{\beta}$ parameters likely entail significant upward bias. Since we do not deal with potential endogeneity bias in this paper, we should not draw causal conclusions from the magnitudes of our estimates; nevertheless, comparing the non-causal relationships across food groups is informative for understanding how the market for plant-based substitutes may or may not be unique.

Figures 4.3 and 4.4 show the vector of coefficient estimates for the $\boldsymbol{\beta}$ parameters

Figure 4.4: Relationship between store-level product entries and real spending levels, 2010-2019

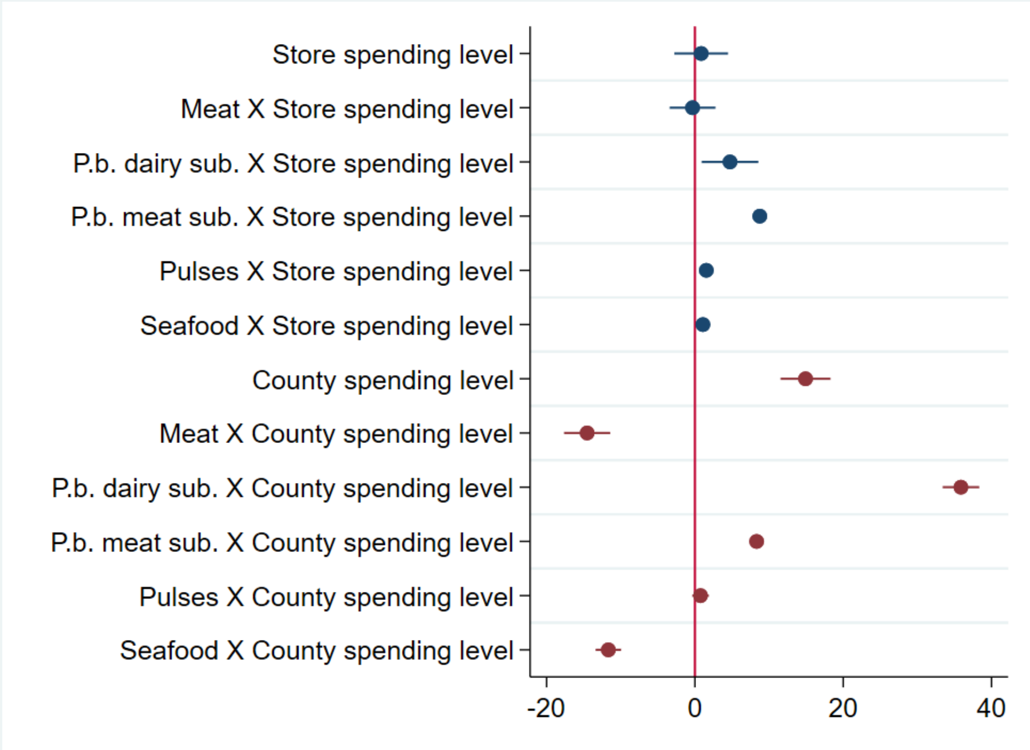

Notes: This figure depicts coefficient estimates for versions of Equation 4.1. To enable visual comparability in this figure, store-level spending is in millions and county-level spending is in billions.

for various specifications. The spending variables are calculated across 2010-2019, and the reference category in all specifications is the animal-based dairy food group. Given the expected upward bias on the coefficient estimates due to endogeneity, the first key takeaway from these estimates is the somewhat surprising result that we do not find a robust positive relationship between new product entries and our measures of spending across all food groups. This result implies that variations in local demand are not always strong indicators of local product entries. For large food groups such as meat, innovation may be related to demand on a more geographically-aggregated scale than individual stores and counties.

Meanwhile, another interesting result of this analysis is that product entries and spending growth and levels seem to be particularly strongly-related for the plant-based meat substitute group (and to some extent the plant-based dairy substitute group) compared to other food groups. This result provides further suggestive evidence that although the trend is not apparent at an aggregate US level, the market for plant-based substitute products is accelerating in localized pockets of the US the market. Furthermore, the result that the relationship between demand and innovation seems to be particularly strong for the plant-based substitutes compared to other food groups lends credence to the stories from the media that demand is driving innovation in this product space.

However, as explained, these estimates likely entail significant endogeneity, so we cannot draw causal inferences from these results. In particular, this relationship is consistent with both directions of causality: localized demand growth could be a driver of store-level introductions of plant-based meat substitutes, or some other factor is driving innovation in plant-based products, and the introduction of these products into stores spurs local demand growth. For example, perhaps socially-motivated entrepreneurs and venture capitalists are driving innovation in plant-based substitutes. However, these supply-side factors behind innovation are unlikely to be as geographically localized as demand, so the fact that we find a relationship between spending growth and product entries at the granular geographic level of stores and counties lends weight to the hypothesis that product entries respond to demand to some extent.[3] An important area for future research is to improve our understanding of the extent to which this link between demand and innovation in plant-based substitute products is causal.

4.4.2 Household-level analysis

The previous section presented evidence that despite the fact that rapid growth in the plant-based substitutes market is not apparent on an aggregate US level, these products may be gaining traction in pockets of the US market where high demand is associated with more product entries for plant-based substitutes. This section uses household-level data to understand the nature of these pockets of high demand and growth, in particular what demographic characteristics are associated with high levels of purchases of these products compared to products from other food groups.

This section uses the rich data from Nielsen's Consumer Panel Dataset to understand demand for our food groups, in particular, what types of households tend to buy these plant-based substitute products, and how these households vary over time and geographically. To do so we estimate a logit model of the share of household h's total food expenditure in year t spent on food group f, s_{hfmt}, where the household lives in Scantrack market area m. Table 4.2 summarizes this outcome variable for each food group. Our six food groups of interest cover 17% of households' annual food expenditures on average. The meat food group has the largest average expenditure share, 7.8%, while the pulses food group has the smallest average share, 0.3%.

[3]When we aggregate to the level of Scantrack market areas, the relationship between spending growth and product entries for this food group is weak and statistically insignificant.

Table 4.2: Average household food group expenditure shares, 2008-2018

Dairy	Meat	Plant-based dairy sub.	Plant-based meat sub.	Pulses	Seafood	Other
0.0651	0.0776	0.0106	0.00878	0.00341	0.0204	0.831
(0.0433)	(0.0474)	(0.0189)	(0.0172)	(0.00452)	(0.0219)	(0.0657)

Notes: This table summarizes households' food group expenditure as a share of total food expenditure reported in the Consumer Panel Data. Standard deviation in parentheses.

Our logit model describes these food group expenditure shares as a function of the age, education level, and race of the head of household as well as a binary indicator for the presence of children in the household. Age, education, and race are categorical variables with 9, 6, and 3 categories respectively. We also control for the number of people in the household, N_h, household income category, M_h, and market area and year fixed effects:

$$logit(s_{hfmt}) = \alpha + \beta_1 Age_h + \beta_2 Edu_h + \beta_3 Race_h + \beta_4 Children_h + \\ \beta_5 N_h + \beta_6 M_h + Year_t + \eta_m + \varepsilon_{hfmt}$$

We estimate this model separately for each food group. Table 4.3 in the appendix to this chapter details the coefficient estimates from this model for each food group. Figure 4.5 illustrates the main patterns of interest that emerge from this analysis. Highly educated households (those with a college degree and especially post-college education) tend to spend relatively less on meat and more on plant-based meat substitutes compared to less educated households. Moreover, young households, particularly under 45 years and especially under 30 years, tend to spend relatively more on plant-based meat substitutes compared to older households. An important feature underlying these results is that the distribution of plant-based meat substitute expenditure shares is highly skewed. Referring to the results illustrated in Figure 4.5, the average predicted expenditure shares for households that are either very young (on the left side of the age axis) or highly educated (the red line) are in the upper quartile of the distribution of household expenditure shares on plant-based meat substitutes. Meanwhile, the average predicated expenditure share of households with low levels of education and in an age category of 35 years and above (the blue line on the right side of the age axis) are below the mean of this distribution.

For the meat food group, age and education are associated with some degree of heterogeneity in average predicted expenditure shares, but overall the distribution of meat expenditure shares is much less skewed than plant-based meat substitute

Figure 4.5: Average predicted expenditure shares by household age and education level

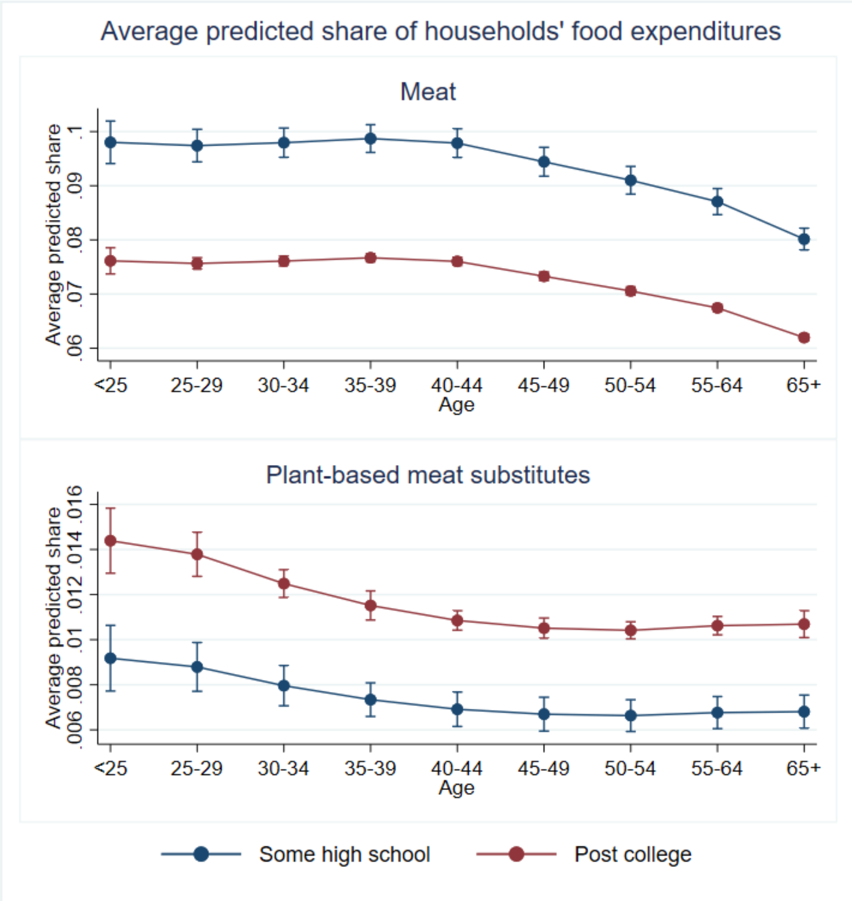

Notes: Error bars represent 95% confidence intervals.

expenditure shares. In particular, the average predicted meat expenditure shares for very young households with some high school education (left side of the blue line in the upper plot of Figure 4.5) and for older households with post college education (right side of the red line) are all within the interquartile range of the distribution of meat expenditure shares. Overall, these results indicate that demand for plant-based meat substitutes is a particularly niche market compared to meat, and is dominated by young and highly educated households.

To explore whether the types of households that tend to spend more on plant-based meat substitute products has changed over time, we estimate the model for this food group separately for three time periods: 2008-2018, 2008-2012, and 2014-2019. Following the insights in Figure 4.5, we simplify the analysis by collapsing the age variable into a binary indicator of whether the household is over 45 years, and by collapsing the education variable into a binary indicator of whether the head of household has post-college education. Figure 4.6 depicts the results of this analysis. The patterns remain broadly the same across all these specifications, highlighting the dominance over time of young and highly educated households in the market for these products. This stability over time in these estimates suggests that these products remain niche and acceleration in the diffusion and adoption of them is not taking place.

Figure 4.6: Average predicted expenditure shares for plant-based meat substitute products by household characteristics

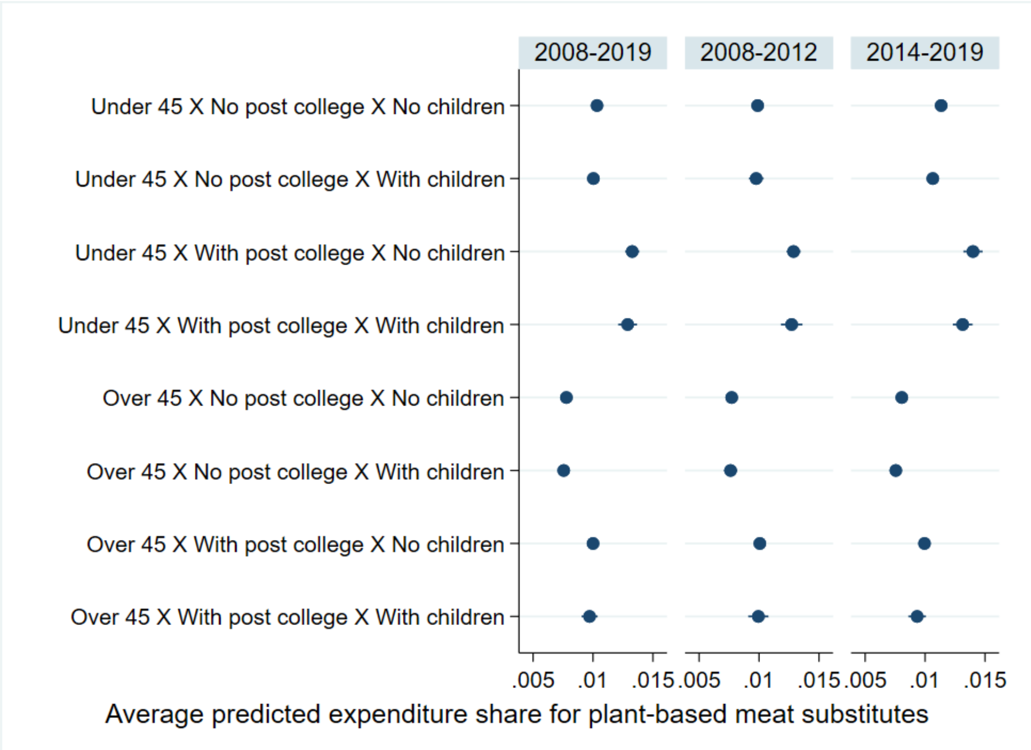

Notes: Under / Over 45 indicates the age of the head of household. "Post college" indicates that a head of household has education beyond the level of a college degree, and "no children" indicates that no children live in the household.

We also run this regression for the plant-based meat substitutes food group separately for each of the 52 Scantrack market areas; Figure 4.10 in the appendix to this chapter illustrates these results, plotting the difference for each Scantrack market in average predicted expenditure shares for households over 45 years versus under 45 years, and with post college education versus without. We find some variation in the magnitudes of these estimated differences in average predicted expenditure shares, but broadly speaking the dominance of young and highly educated households in purchases of plant-based meat substitute products remains consistent across most of these metropolitan areas. We see these results as suggestive evidence that the tendency to buy plant-based meat substitutes is associated less with the local culture of a city, and more with the characteristics of a household regardless of what city they live in.

4.4.3 US demographic trends

The previous sections provide evidence of a positive association between local product introductions and local demand for the plant-based meat substitute food group, alongside insight into the types of households that tend to buy these products. This section brings together these two pieces of analyses by using data on US demographic trends to assess the potential that demand growth for plant-based substitute products has been exogenous to the innovation process. Following Acemoglu and Linn (2004) and Jaravel (2019), we can conceptualize potential demand for a category of products as a combination (among other factors) of (i) the population size of the households that tend to buy these products, and (ii) the income levels of these households. In this spirit, we compare our findings from the household-level analysis of the Nielsen data with demographic trends from the 5-year American Community Survey (ACS). In particular, the household-level analysis suggests that young and highly educated households are dominant buyers of plant-based substitute products. If the population or incomes of these types of households have seen relatively strong growth, then such demographic trends would provide suggestive evidence that increased demand for plant-based meat substitute products may be exogenous to local product introductions to some extent.

Figure 4.7: US household demographic trends, 2009-2018

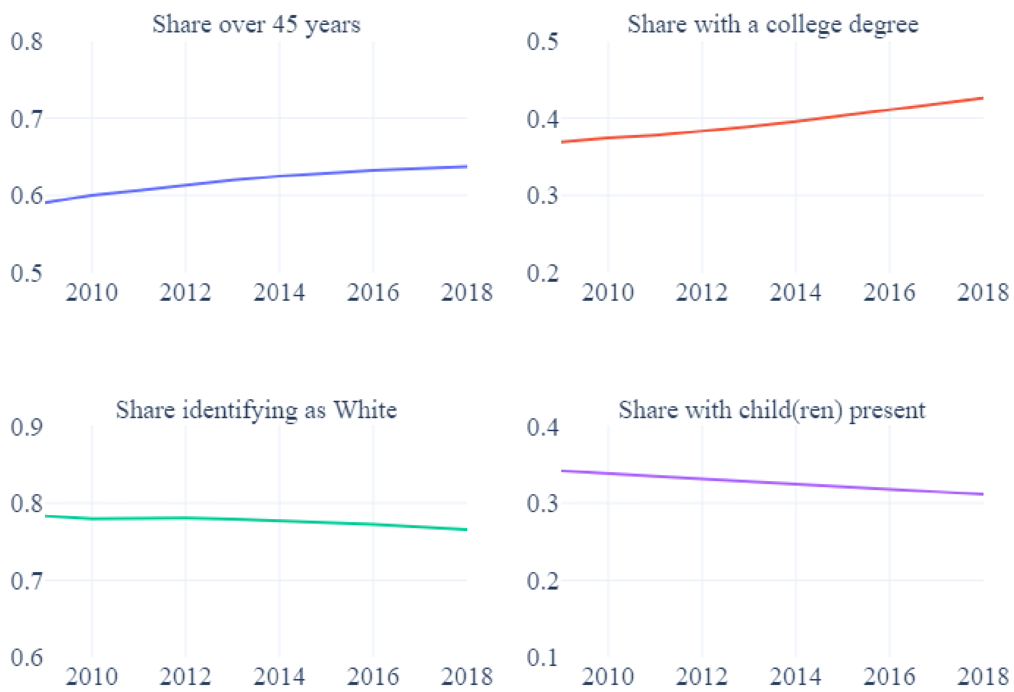

Figure 4.7 illustrates the aggregate demographic trends in the ACS data for each of the characteristics considered in our household-level analysis. Broadly speaking, US households have been getting older, more educated, more racially diverse, and choosing not to have children more often. The trend towards increased education levels aligns with the takeaway from our household-level analysis that highly educated households are associated with relatively high demand for plant-based meat substitutes. Particularly if households sort geographically in such a way that highly educated households tend to live near similar highly-educated households, the demographic trend towards increasing education levels may represent (via channel (i) above) an increase in demand for plant-based meat substitutes that is exogenous to the innovation process. This insight provides suggestive evidence that the association that we found between local spending growth and new product entries discussed may indeed represent (to some extent) a causal link from demand to product introductions. Meanwhile, the potential for exogenous demand growth via increased incomes of households that purchase plant-based meat substitutes seems limited. Figure 4.8 illustrates trends in average household income growth based on the different household characteristics. In recent years young households have seen slightly higher income growth than other households, but otherwise an association between relatively high income growth and a high tendency to buy plant-based meat substitutes is not apparent.

This section provides a bit of evidence that real demand growth via increasing education levels of US households may be causing increased innovation in plant-based substitute products. However, overall US demographic trends do not provide a very hopeful story for increased potential demand for plant-based substitutes, and this analysis suggests that as long a the types of households buying these products does not expand out of its niche, then potential growth for these products is limited, and so too the potential that they can play a meaningful role in decarbonization.

4.5 Conclusion

This paper investigates two popular understandings of the market for plant-based substitute products: (i) this market is rapidly expanding; (ii) growing demand is a key driver of innovation in these products. Using rich microdata on US food sales, we draw comparisons across time, space, and food groups to shed some empirical light on these popular stories. Our analysis suggests that the market for these products seems to be growing in pockets of the US market, which helps to explain that from the perspective of the private investor this trend may offer some exciting investment opportunities. However, from the perspective of decarbonizing the food sector, this growth is not widespread enough to offer hope of meaningful GHG mitigation; our

Figure 4.8: US households income growth trends, 2010-2018

analysis of aggregate US trends in spending and innovation paints a lacklustre picture for the plant-based substitute food groups. Next, we find some evidence that the link between local demand and store-level product entries is particularly strong for the plant-based substitute food groups, but the types of households buying these products hasn't changed over the last decade. Plant-based meat substitute products remain a niche product group popular amongst a subset of young, highly educated households. The demographic trend towards increased education levels in the US may help to broaden the potential demand pool for these products somewhat, but if the types of households buying plant-based meat products does not change, we should not expect a large shift in demand for these products on aggregate.

The literature on environmental innovation points to policy intervention as a key driver of green transitions, and so given the lack of policy intervention so far in the food sector, the result in this paper that plant-based substitute products remain niche may not come as a surprise. Policy intervention may be necessary to help grow the market for plant-based meat substitutes and ultimately reduce meat consumption in the US. The first-best policy to correct the market failure of food-related carbon emissions would be to price the carbon embodied in food, but carbon pricing policies have faced political backlash in other sectors of the economy. Carbon pricing may be even less feasible in the food sector compared to the energy and transport sectors, because of concerns about impacts on small rural farm-based economies alongside the usual concerns about distributional impacts on consumers. With these constraints in mind, policymakers may want to consider alternatives to the first-best approach to decarbonizing the food sector, such as investment in R&D for plant-based substitute products.

Future work can build on the analysis in this paper, for example by using more sophisticated techniques to identify groups within the product space and by using causal inference methods to assess the role of demand as a driver of green innovation in the food sector. A shift-share instrumental variable approach that exploits demographic changes may be a potential avenue to pursue in this regard. As it is, this paper makes several important contributions. First, our allocation of products into customized food groups provides a strong basis for documenting and analysing trends in the plant-based substitute market. Second, we shed some empirical light on popular stories about the market for plant-based substitute products, with our analysis suggesting that from the perspective of decarbonizing the food sector the growth in this market may be overstated. Finally, we make an important policy contribution in this paper by illustrating that without policy intervention plant-based products are likely to stay on the path of being a niche market.

4.A Appendix to Chapter 4

Figure 4.9: Histograms of store-level average annual new product introductions

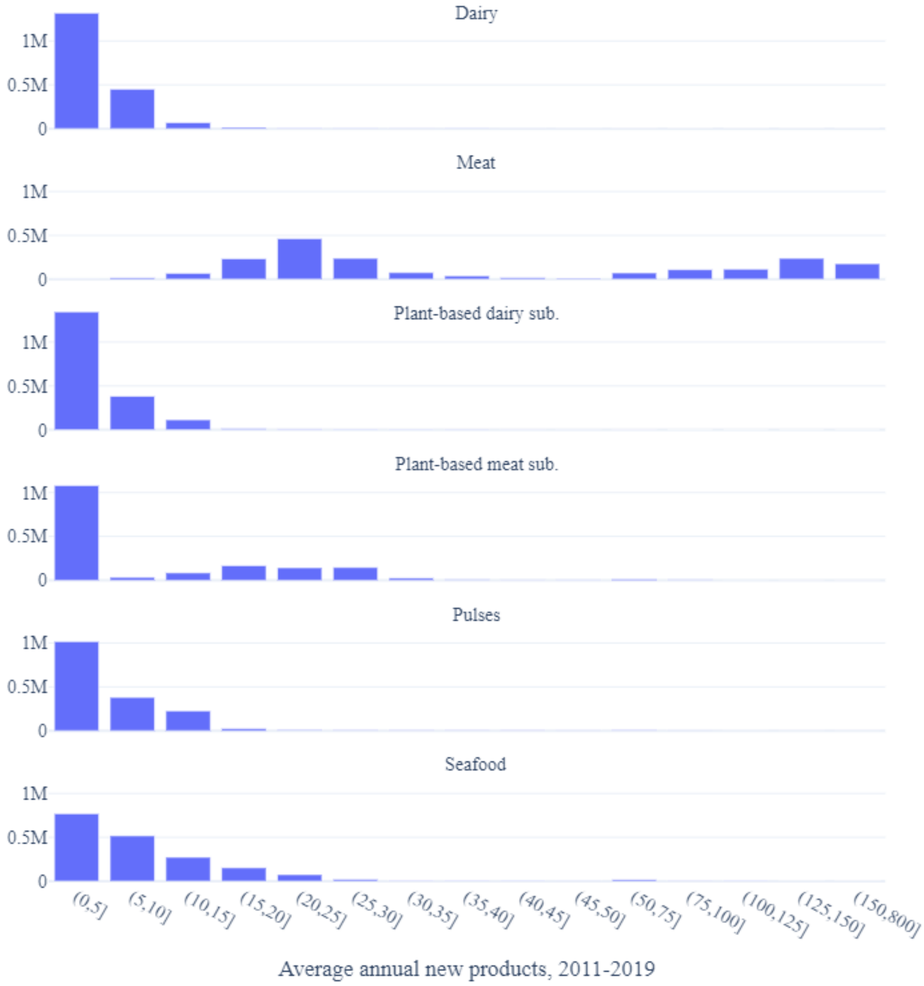

Figure 4.10: Comparing average predicted plant-based meat substitute expenditure shares across Scantrack markets

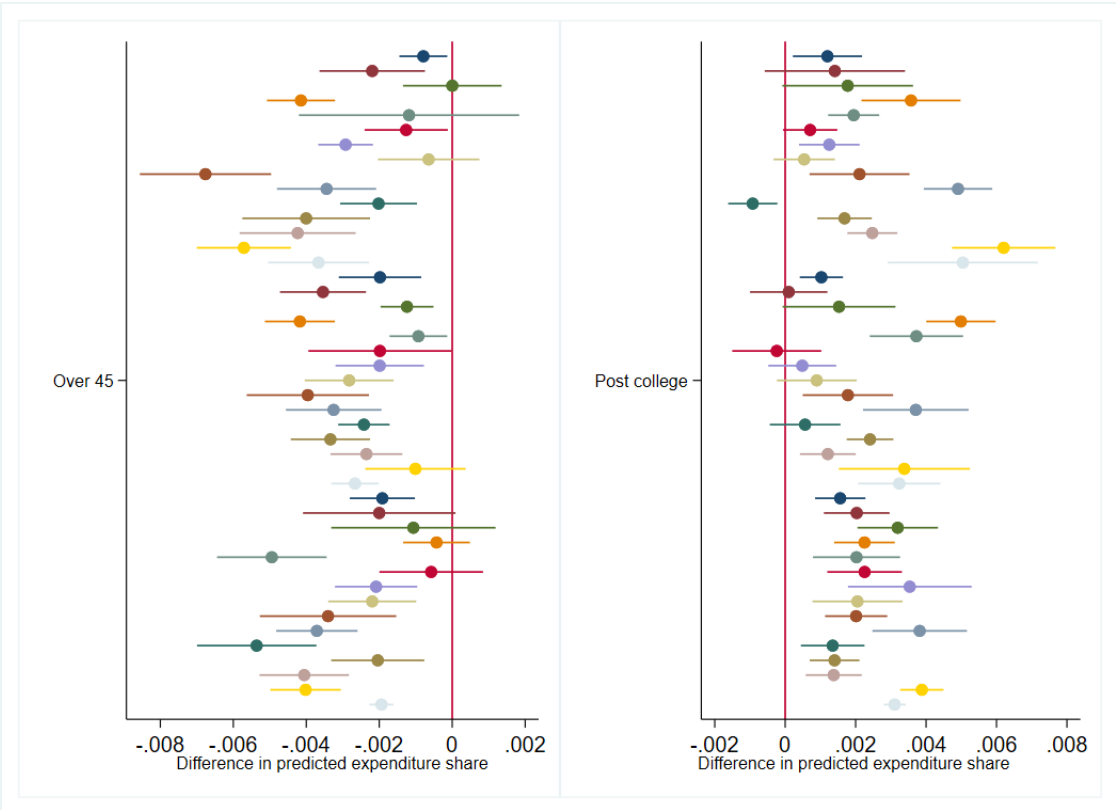

Notes: Each point represents the difference in average predicted expenditure share for a single Scantrack market area. 'Over 45' indicates the difference between households with the head of household over versus under 45 years. 'Post college' indicates the difference between households whose head has post college education versus those who do not. Error bars represent 95% confidence intervals.

Table 4.3: Household-level analysis

Dependent variable: Food group's share of household's total food expenditures

	Dairy	Meat	Plant-based dairy sub.	Plant-based meat sub.	Pulses	Seafood
AGE						
Reference category: <25 years						
25-29 years	0.0296*	-0.000765	0.0874	0.0336	0.0314	-0.0212
	(0.0134)	(0.0160)	(0.0543)	(0.0596)	(0.0312)	(0.0259)
30-34 years	0.0246	-0.0171	0.138**	0.0162	-0.0325	-0.0279
	(0.0146)	(0.0158)	(0.0471)	(0.0586)	(0.0325)	(0.0251)
35-39 years	-0.0202	-0.00956	0.0792	-0.0471	-0.0849**	-0.0126
	(0.0154)	(0.0158)	(0.0454)	(0.0621)	(0.0315)	(0.0242)
40-44 years	-0.0682***	0.000828	-0.0187	-0.126*	-0.119***	0.0249
	(0.0141)	(0.0156)	(0.0428)	(0.0542)	(0.0335)	(0.0230)
45-49 years	-0.100***	-0.0126	-0.0522	-0.211***	-0.151***	0.0460
	(0.0128)	(0.0161)	(0.0464)	(0.0611)	(0.0297)	(0.0236)
50-54 years	-0.0743***	-0.0332*	-0.0880	-0.282***	-0.155***	0.0653**
	(0.0123)	(0.0155)	(0.0458)	(0.0565)	(0.0291)	(0.0238)
55-64 years	-0.0419***	-0.0656***	-0.0669	-0.337***	-0.144***	0.109***
	(0.0121)	(0.0158)	(0.0460)	(0.0573)	(0.0293)	(0.0222)
65+ years	0.0427***	-0.142***	0.0411	-0.407***	-0.139***	0.181***

	Dependent variable: Food group's share of household's total food expenditures					
	Dairy	Meat	Plant-based dairy sub.	Plant-based meat sub.	Pulses	Seafood
	(0.0125)	(0.0177)	(0.0519)	(0.0656)	(0.0331)	(0.0261)
EDUCATION						
Reference category: Grade school						
Some high school	-0.0731*	0.0235	-0.201	-0.227	-0.134	-0.0403
	(0.0361)	(0.0362)	(0.173)	(0.136)	(0.0982)	(0.0686)
High school	-0.0877*	-0.0358	-0.197	-0.200	-0.241**	-0.0549
	(0.0388)	(0.0348)	(0.146)	(0.128)	(0.0916)	(0.0679)
Some college	-0.0718	-0.0923**	-0.0830	-0.115	-0.261**	-0.00195
	(0.0390)	(0.0338)	(0.144)	(0.131)	(0.0896)	(0.0678)
College	-0.00780	-0.168***	0.0712	0.0620	-0.202*	0.0571
	(0.0385)	(0.0353)	(0.143)	(0.128)	(0.0883)	(0.0685)
Post college	0.0584	-0.254***	0.182	0.228	-0.111	0.118
	(0.0380)	(0.0371)	(0.146)	(0.121)	(0.0854)	(0.0697)
RACE						
Reference category: White						
Black	-0.207***	0.328***	0.164***	0.197***	0.244***	0.411***
	(0.00955)	(0.00892)	(0.0262)	(0.0311)	(0.0181)	(0.0208)
Other	0.0768***	0.0357***	0.264***	0.106**	0.226***	0.269***
	(0.0160)	(0.00809)	(0.0252)	(0.0328)	(0.0248)	(0.0157)

Dependent variable: Food group's share of household's total food expenditures

	Dairy	Meat	Plant-based dairy sub.	Plant-based meat sub.	Pulses	Seafood
Children	0.133^{***}	0.00200	0.00969	-0.0614	-0.0704^{***}	-0.0614^{***}
	(0.00615)	(0.00501)	(0.0211)	(0.0324)	(0.0105)	(0.00938)
Household size	0.0102^{***}	0.0605^{***}	-0.152^{***}	-0.157^{***}	-0.0476^{***}	-0.0466^{***}
	(0.00211)	(0.00218)	(0.00923)	(0.0111)	(0.00722)	(0.00437)
Household income	-0.00448^{***}	-0.00633^{***}	-0.0120^{***}	-0.0131^{***}	-0.0113^{***}	0.000942
	(0.000423)	(0.000364)	(0.00131)	(0.00127)	(0.00121)	(0.000596)
Observations	670640	665182	271726	296985	430822	586219

Notes: Standard errors (in parentheses) are clustered by Scantrack market area.All models in include Scantrack market and year fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Bibliography

Abadie, L M et al. (Oct. 2016). “Using food taxes and subsidies to achieve emission reduction targets in Norway”. In: *J. Clean. Prod.* 134, pp. 280–297.

Acemoglu, Daron and Joshua Linn (Aug. 2004). “Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry”. In: *Q. J. Econ.* 119.3, pp. 1049–1090.

Agnosteva, Delina, James Anderson, and Yoto Yotov (Jan. 2014). *Intra-national Trade Costs: Measurement and Aggregation*. Tech. rep. 19872. National Bureau of Economic Research, Inc.

Aichele, Rahel and Gabriel Felbermayr (2015). “Kyoto and Carbon Leakage: An Empirical Analysis of the Carbon Content of Bilateral Trade”. In: *Rev. Econ. Stat.* 97.1, pp. 104–115.

Alcalá, Francisco and Antonio Ciccone (May 2004). “Trade and Productivity”. en. In: *Q. J. Econ.* 119.2, pp. 613–646.

Anderson, James and Yoto Yotov (2016). “Terms of trade and global efficiency effects of free trade agreements, 1990–2002”. In: *J. Int. Econ.* 99.C, pp. 279–298.

Anderson, James E (Sept. 2011). “The Gravity Model”. In: *Annu. Rev. Econom.* 3.1, pp. 133–160.

Anderson, James E, Mario Larch, and Yoto V Yotov (Aug. 2020). “Transitional Growth and Trade with Frictions: A Structural Estimation Framework”. In: *Econ J* 130.630, pp. 1583–1607.

Anderson, James E and Eric van Wincoop (Mar. 2003). “Gravity with Gravitas: A Solution to the Border Puzzle”. In: *Am. Econ. Rev.* 93.1, pp. 170–192.

Andersson, Julius J (Oct. 2019). “The “Meatigation” of Climate Change: Environmental and Distributional Effects of a Greenhouse Gas Tax on Animal Food Products”. en. In: *Working paper*.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodrıguez-Clare (Feb. 2012). “New Trade Models, Same Old Gains?” In: *Am. Econ. Rev.* 102.1, pp. 94–130.

Atkin, David, Amit K Khandelwal, and Adam Osman (May 2017). “Exporting and Firm Performance: Evidence from a Randomized Experiment”. In: *Q. J. Econ.* 132.2, pp. 551–615.

Auffhammer, Maximilian (Nov. 2018). “Quantifying Economic Damages from Climate Change”. In: *J. Econ. Perspect.* 32.4, pp. 33–52.

Baier, Scott L and Jeffrey H Bergstrand (Mar. 2007). “Do free trade agreements actually increase members’ international trade?” In: *J. Int. Econ.* 71.1, pp. 72–95.

Baier, Scott L, Yoto V Yotov, and Thomas Zylkin (Jan. 2019). “On the widely differing effects of free trade agreements: Lessons from twenty years of trade integration”. In: *J. Int. Econ.* 116, pp. 206–226.

Baldwin, Richard and Daria Taglioni (Sept. 2006). “Gravity for Dummies and Dummies for Gravity Equations”.

Bashi, Zafer et al. (Aug. 2019). *Alternative proteins: The race for market share is on.* Tech. rep. McKinsey & Company.

Becker, Austin H et al. (Oct. 2013). “A note on climate change adaptation for seaports: a challenge for global ports, a challenge for global society”. In: *Clim. Change* 120.4, pp. 683–695.

Bergstrand, Jeffrey, Mario Larch, and Yoto Yotov (2015). “Economic integration agreements, border effects, and distance elasticities in the gravity equation”. In: *Eur. Econ. Rev.* 78.C, pp. 307–327.

Beverelli, Cosimo et al. (Feb. 2018). *Institutions, Trade and Development: A Quantitative Analysis.* Tech. rep. 2018-3.

Bonnet, Céline, Zohra Bouamra-Mechemache, and Tifenn Corre (May 2018). “An Environmental Tax Towards More Sustainable Food: Empirical Evidence of the Consumption of Animal Products in France”. In: *Ecol. Econ.* 147, pp. 48–61.

Borchert, Ingo et al. (2021). “The International Trade and Production Database for Estimation (ITPD-E)”. In: *International Economics* 166.C, pp. 140–166.

Branger, Frédéric and Philippe Quirion (Jan. 2014). “Climate policy and the ‘carbon haven’ effect”. en. In: *Wiley Interdiscip. Rev. Clim. Change* 5.1, pp. 53–71.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel (Nov. 2015). “Global non-linear effect of temperature on economic production”. en. In: *Nature* 527.7577, pp. 235–239.

Caillavet, France, Adélaıde Fadhuile, and Véronique Nichèle (Sept. 2016). “Taxing animal-based foods for sustainability: environmental, nutritional and social perspectives in France”. In: *Eur. Rev. Agric. Econ.* 43.4, pp. 537–560.

Center for International Earth Science Information Network - CIESIN - Columbia University (2018). *Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11.* Palisades, NY.

Chinowsky, Paul et al. (Mar. 2019). “Impacts of climate change on operation of the US rail network”. In: *Transp. Policy* 75, pp. 183–191.

Copeland, Brian R, Joseph S Shapiro, and M Scott Taylor (Jan. 2022). “Globalization and the environment”. In: *Handbook of International Economics.* Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Vol. 5. Elsevier, pp. 61–146.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin (Mar. 2020). “Fast Poisson estimation with high-dimensional fixed effects”. In: *Stata J.* 20.1, pp. 95–115.

Costinot, Arnaud, Dave Donaldson, and Cory Smith (2016). “Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World”. In: *J. Polit. Econ.* 124.1, pp. 205–248.

Crippa, M et al. (Mar. 2021). “Food systems are responsible for a third of global anthropogenic GHG emissions”. en. In: *Nature Food* 2.3, pp. 198–209.

Dallmann, Ingrid (Jan. 2019). “Weather Variations and International Trade”. In: *Environ. Resour. Econ.* 72.1, pp. 155–206.

Dawson, Jackie et al. (Dec. 2020). “Treatment of International Economic Trade in Intergovernmental Panel on Climate Change (IPCC) Reports”. In: *Current Climate Change Reports* 6.4, pp. 155–165.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (July 2012). “Temperature Shocks and Economic Growth: Evidence from the Last Half Century”. In: *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

Dingel, Jonathan I, Kyle C Meng, and Solomon M Hsiang (Jan. 2019). “Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate”.

Domínguez-Iino, Tomás (Jan. 2021). “Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains”. In: *Working paper*.

Drupp, Moritz A et al. (Nov. 2018). “Discounting Disentangled”. In: *American Economic Journal: Economic Policy* 10.4, pp. 109–134.

Eaton, Jonathan and Samuel Kortum (2002). “Technology, Geography, and Trade”. In: *Econometrica* 70.5, pp. 1741–1779.

Edjabou, Louise Dyhr and Sinne Smed (Apr. 2013). “The effect of using consumption taxes on foods to promote climate friendly diets – The case of Denmark”. In: *Food Policy* 39, pp. 84–96.

FAIRR Initiative (2021). *Appetite for Disruption: The Last Serving*. Tech. rep.

FAO (Mar. 2021). *FAOSTAT Value of Agricultural Production*.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (Oct. 2015). “The Next Generation of the Penn World Table”. In: *Am. Econ. Rev.* 105.10, pp. 3150–3182.

Felder, Stefan and Thomas Rutherford (1993). “Unilateral CO2 Reductions and Carbon Leakage: The Consequences of International Trade in Oil and Basic Materials”. In: *J. Environ. Econ. Manage.* 25.2, pp. 162–176.

Feyrer, James (Oct. 2019). “Trade and Income—Exploiting Time Series in Geography”. In: *Am. Econ. J. Appl. Econ.* 11.4, pp. 1–35.

Food and Agriculture Organization of the United Nations (2021). *FAOSTAT*.

Frankel, Jeffrey A and David H Romer (June 1999). “Does Trade Cause Growth?” In: *Am. Econ. Rev.* 89.3, pp. 379–399.

Funke, Franziska et al. (2021). “Is meat too cheap? Towards optimal meat taxation”. en. In: *SSRN Electron. J.*

Garcıa-Muros, Xaquin et al. (Jan. 2017). "The distributional effects of carbon-based food taxes". In: *J. Clean. Prod.* 140, pp. 996–1006.

Gerber, P J et al. (2013). *Tackling climate change through livestock: a global assessment of emissions and mitigation opportunities.* en. Tech. rep.

Görg, Holger, Richard Kneller, and Balázs Muraközy (Nov. 2012). "What makes a successful export? Evidence from firm-product- level data". en. In: *Can. J. Econ.* 45.4, pp. 1332–1368.

Gouel, Christophe and David Laborde (Nov. 2018). *The Crucial Role of International Trade in Adaptation to Climate Change.* Tech. rep. 25221.

Grubb, Michael et al. (Mar. 2021). "Induced innovation in energy technologies and systems: a review of evidence and potential implications for CO2 mitigation". en. In: *Environ. Res. Lett.* 16.4, p. 043007.

Hallegatte, Stephane, Jun Rentschler, and Julie Rozenburg (2020). *Adaptation Principles : A Guide for Designing Strategies for Climate Change Adaptation and Resilience.* Tech. rep. Washington, D.C.: The World Bank.

Head, Keith and Thierry Mayer (Jan. 2014). "Chapter 3 - Gravity Equations: Workhorse,Toolkit, and Cookbook". In: *Handbook of International Economics.* Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Vol. 4. Elsevier, pp. 131–195.

Hedenus, Fredrik, Stefan Wirsenius, and Daniel J A Johansson (May 2014). "The importance of reduced meat and dairy consumption for meeting stringent climate change targets". en. In: *Clim. Change* 124.1-2, pp. 79–91.

Heid, Benedikt, Mario Larch, and Yoto Yotov (Oct. 2017). *Estimating the Effects of Non-discriminatory Trade Policies within Structural Gravity Models.* Tech. rep. 2017-10.

Henderson, B et al. (Mar. 2018). "The power and pain of market-based carbon policies: a global application to greenhouse gases from ruminant livestock production". In: *Mitigation and Adaptation Strategies for Global Change* 23.3, pp. 349–369.

Hersbach, Hans et al. (July 2020). "The ERA5 global reanalysis". en. In: *Quart. J. Roy. Meteor. Soc.* 146.730, pp. 1999–2049.

Jaravel, Xavier (May 2019). "The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector". In: *Q. J. Econ.* 134.2, pp. 715–783.

Jens Horbach, Christian Rammer, and Klaus Rennings (2011). *Determinants of Eco-innovations by Type of Environmental Impact The Role of Regulatory Push/Pull, Technology Push and Market Pull.* Tech. rep. 11-027. ZEW.

Jones, Benjamin F and Benjamin A Olken (May 2010). "Climate Shocks and Exports". In: *Am. Econ. Rev.* 100.2, pp. 454–459.

Kahn, Matthew E et al. (Dec. 2021). "Long-term macroeconomic effects of climate change: A cross-country analysis". In: *Energy Econ.* 104, p. 105624.

Kalkuhl, Matthias and Leonie Wenz (Sept. 2020). "The impact of climate conditions on economic production. Evidence from a global panel of regions". In: *J. Environ. Econ. Manage.* 103, p. 102360.

Kehlbacher, Ariane et al. (July 2016). "The distributional and nutritional impacts and mitigation potential of emission-based food taxes in the UK". In: *Clim. Change* 137.1, pp. 121–141.

Kolstad, Charles D and Frances C Moore (Jan. 2020). "Estimating the Economic Impacts of Climate Change Using Weather Observations". In: *Review of Environmental Economics and Policy* 14.1, pp. 1–24.

Krugman, Paul R (Dec. 1987). "Is Free Trade Passe?" In: *J. Econ. Perspect.* 1.2, pp. 131–144.

Larch, Mario and Joschka Wanner (2017). "Carbon tariffs: An analysis of the trade, welfare, and emission effects". In: *J. Int. Econ.* 109.C, pp. 195–213.

Mbow, C. et al. (2019). "Food Security". In: *Climate Change and Land: an IPCC special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems.* Ed. by M. van Diemen et al.

Moore, Frances C et al. (Nov. 2017). "New science of climate change impacts on agriculture implies higher social cost of carbon". en. In: *Nat. Commun.* 8.1, p. 1607.

Nath, Ishan B (June 2020). "The food problem and the aggregate productivity consequences of climate change". In: *NBER Working Paper Series* Working Paper 27297.

Newell, Richard G, Brian C Prest, and Steven E Sexton (July 2021). "The GDP-Temperature relationship: Implications for climate change damages". In: *J. Environ. Econ. Manage.* 108, p. 102445.

Osberghaus, Daniel (Oct. 2019). "The Effects of Natural Disasters and Weather Variations on International Trade and Financial Flows: a Review of the Empirical Literature". In: *Economics of Disasters and Climate Change* 3.3, pp. 305–325.

Poore, J and T Nemecek (June 2018). "Reducing food's environmental impacts through producers and consumers". en. In: *Science* 360.6392, pp. 987–992.

Popp, David (Mar. 2019). "Environmental Policy and Innovation: A Decade of Research".

Popp, David, Tamara Hafner, and Nick Johnstone (Sept. 2007). *Policy vs. Consumer Pressure: Innovation and Diffusion of Alternative Bleaching Technologies in the Pulp Industry.* Tech. rep. 13439.

Ramsey, F P (1928). "A Mathematical Theory of Saving". In: *Econ. J. Nepal* 38.152, pp. 543–559.

Sager, Lutz (Apr. 2019). "The global consumer incidence of carbon pricing: Evidence from trade". In: *GRI Working Paper Series.*

Säll, Sarah (Jan. 2018). “Environmental food taxes and inequalities: Simulation of a meat tax in Sweden”. In: *Food Policy* 74, pp. 147–153.

Säll, Sarah and Ing-Marie Gren (Aug. 2015). “Effects of an environmental tax on meat and dairy consumption in Sweden”. In: *Food Policy* 55, pp. 41–53.

Sandström, Vilma et al. (Dec. 2018). “The role of trade in the greenhouse gas footprints of EU diets”. In: *Global Food Security* 19, pp. 48–55.

Shapiro, Joseph S (Nov. 2016). “Trade Costs, CO_2, and the Environment”. In: *American Economic Journal: Economic Policy* 8.4, pp. 220–254.

Silva, J M C Santos and Silvana Tenreyro (Nov. 2006). “The log of gravity”. en. In: *Rev. Econ. Stat.* 88.4, pp. 641–658.

Simon, Dora (2020). “To Beef, or Not To Beef: Trade, Meat, and the Environment”.

Sousa, José de, Thierry Mayer, and Soledad Zignago (2012). “Market access in global and regional trade”. In: *Reg. Sci. Urban Econ.* 42.6, pp. 1037–1052.

Springmann, Marco et al. (Nov. 2016). “Mitigation potential and global health impacts from emissions pricing of food commodities”. en. In: *Nat. Clim. Chang.* 7.1, pp. 69–74.

Springmann, Marco et al. (Oct. 2018). “Options for keeping the food system within environmental limits”. en. In: *Nature* 562.7728, pp. 519–525.

The World Bank (2020). *Worldwide Governance Indicators.*

— (2022). *World Bank Country and Lending Groups - Historical Classification.*

UNIDO (2020). *INDSTAT 2 Industrial Statistics Database at 2-digit level of ISIC Revision 3.* Vienna.

United Nations (2021). *UN Commodity Trade Statistics Database.*

Veugelers, Reinhilde (Dec. 2012). “Which policy instruments to induce clean innovating?” In: *Res. Policy* 41.10, pp. 1770–1778.

Willmott, C J and K Matsuura (2018). *Terrestrial Precipitation: 1900-2017 Gridded Monthly Time Series.*

Wirsenius, Stefan, Fredrik Hedenus, and Kristina Mohlin (2011). “Greenhouse gas taxes on animal food products: rationale, tax scheme and climate mitigation effects”. In: *Clim. Change* 108.1, pp. 159–184.

Yotov, Yoto V et al. (2016). *An advanced guide to trade policy analysis: The structural gravity model.* Tech. rep. WTO.

Zech, Konstantin M and Uwe A Schneider (Mar. 2019). “Carbon leakage and limited efficiency of greenhouse gas taxes on food products”. In: *J. Clean. Prod.* 213, pp. 99–103.