Policy Research Working Paper 10642

Fertilizer Import Bans, Agricultural Exports, and Welfare

Evidence from Sri Lanka

Devaki Ghose
Eduardo Fraga
Ana Fernandes

WORLD BANK GROUP
Development Economics
Development Research Group
December 2023

Policy Research Working Paper 10642

Abstract

In May 2021, Sri Lanka's government imposed an abrupt ban on chemical fertilizer imports. This paper leverages this natural experiment to quantify the costs of a lack of access to fertilizer for agricultural production and trade in a developing economy heavily dependent on agriculture. Using high-frequency firm-level trade data, ground production records, newly developed remote sensing crop yield estimates, and event study designs, the analysis reveals significant declines in fertilizer imports, agricultural output, and exports of fertilizer-dependent crops. These findings underscore the importance of trade policy for chemical fertilizer, which is hard to substitute with organic or domestic alternatives in the short run. A quantitative spatial model of trade and agriculture shows the ban's average welfare effects were equivalent to a 4.35%income reduction, with disproportionate losses for farmers, estate workers, and the regions that cultivate fertilizer-intensive crops. The model also highlights the interaction of massive fertilizer subsidies, a domestic agricultural policy common in many countries, with trade policy: by nearly eliminating fertilizer use, an import ban scales down the subsidy program and its associated income transfers from non-farming to farming sectors, thus attenuating the welfare losses of mobile workers. The findings quantify the costs of lack of fertilizer access and the role of trade and industrial policy in determining such access.

This paper is a product of the Development Research Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The authors may be contacted at dghose@worldbank.org, epinheirofraga@worldbank.org, and afernandes@worldbank.org.

Fertilizer Import Bans, Agricultural Exports, and Welfare: Evidence from Sri Lanka*

Devaki Ghose† Eduardo Fraga‡ Ana Fernandes§

Originally published in the Policy Research Working Paper Series on *December 2023*.
This version is updated on *January 2025*.
To obtain the originally published version, please email prwp@worldbank.org.

JEL Codes: D58; F13; F14; O13; Q17.
Keywords: Trade, industrial policy, import ban, non-tariff measures,fertilizer, and agriculture.

*Eduardo is extremely grateful to Costas Arkolakis, Sam Kortum, Lorenzo Caliendo, and Ana Cecilia Fieler for their continuous support and guidance. We also thank Doug Gollin, Sebastian Sotelo, Heitor Pellegrina, Michael Peters, Nuno Limao, Lauren Falcao Bergquist, John Keyser, Gonzalo Varela, Athula Senaratne, Sheu Salau, Manoj Thibbotuwawa, Camille Reverdy, Daan van Soest, Kari Heerman, and Alejandro Forero for helpful discussions. We thank Mutlu Ozdogan for his work on the remote sensing estimates. The authors gratefully acknowledge support by the "Whole of Economy" Program, a Climate Support facility administered by the World Bank. We also thank the Umbrella Facility for Trade trust fund (financed by the governments of the Netherlands, Norway, Sweden, Switzerland, and the United Kingdom) and the World Bank's Research Support Budget for financial support. The findings, interpretations, and conclusions expressed in this paper are solely those of the authors and do not necessarily represent the views of the World Bank, its affiliated organizations, or the Executive Directors or the countries they represent. All errors are our responsibility.

†Corresponding author, World Bank, DECRG. Email: dghose@worldbank.org
‡London Business School. Email: efraga@london.edu
§World Bank, DECRG. Email: afernandes@worldbank.org

1 Introduction

Increased use of critical intermediate inputs such as fertilizers is associated with enhanced agricultural productivity, a necessary condition for structural transformation and growth (Bustos et al., 2016; Foster and Rosenzweig, 2008; Gollin et al., 2002; McArthur and McCord, 2017). This leads governments worldwide to subsidize fertilizers with the aim of increasing agricultural yields and fostering economic development. Yet, recent conflicts, such as Russia's invasion of Ukraine (Hebebrand and Glauber, 2023; Zereyesus et al., 2022), and trade policy activism in the form of import and export restrictions are forces leading to fertilizer price increases and shortages. Trade protectionism is particularly detrimental, as most countries rely on imports for their modern agricultural inputs, of which fertilizers are key (Farrokhi and Pellegrina, 2023).[1] More than 80% of countries are net importers of fertilizer, with imports accounting for more than 75% of their fertilizer supply.[2] Against this background, what are the economy-wide costs and distributional consequences of restricted fertilizer access in a developing country where agriculture is of central importance? And what is the role of trade policy in determining such access?

Answering these questions has been challenging, as most attempts to quantify the value of fertilizer rely on small-scale field experiments (Beaman et al., 2013; Duflo et al., 2011) that identify local effects precisely but do not inform on general equilibrium (GE) effects (Bergquist et al., 2023; Muralidharan and Niehaus, 2017). Analyzing the distributional effects of limited fertilizer access with a classic spatial trade and agriculture model is also challenging due to the need to account for trade policy and for certain key features of developing countries.[3] First, fertilizer markets in developing countries are often characterized by heavy government subsidies that act as a net transfer from non-farming to farming occupations, with important distributional consequences (Amaglobeli et al., 2024; Druilhe and Barreiro-Hurlé, 2012). Second, the share of income spent on food falls with income (non-homothetic preferences), a pattern seen worldwide. Third, inequality in land holdings across individuals is an important feature of developing countries (Bauluz et al., 2020). Land inequality leads to income inequality, which, combined with non-homothetic preferences, creates regional differences in food expenditure shares. These differences influence the relative sizes of the agricultural and manufacturing sectors. Since the manufacturing sector provides an employment alternative for workers during agricultural downturns, its size influences welfare.

This paper overcomes these challenges and quantifies the economy-wide costs of a lack of access to fertilizer by leveraging an unprecedented natural experiment and building a spatial model with several realistic features of developing countries that have been mostly missing from quantita-

[1] Around the world, two thirds of every dollar spent on modern agricultural inputs are paid to foreign suppliers (Farrokhi and Pellegrina, 2023).

[2] https://www.fao.org/faostat/en/#data/RFN

[3] Some recent examples of these models include Bergquist et al. (2023), Nath (forthcoming), and Sotelo (2020). See Related Literature below for more examples.

tive models of trade and agriculture: non-homothetic preferences, crop-specific fertilizer subsidies, land inequality, and an import ban modeled as a quantitative restriction (QR) as opposed to an *ad valorem* cost. The government of Sri Lanka, a country with a large agricultural sector and one of the heaviest fertilizer subsidy programs in the world, imposed an abrupt and unexpected ban on the import of all chemical fertilizers in May 2021, a critical moment in the growing season. The stated objective was to move the country towards organic agriculture, while the unstated objective was saving foreign exchange.[4] Since more than 90% of the chemical fertilizer in Sri Lanka was imported and the market for organic fertilizer was virtually non-existent, the import ban resulted in an economy-wide fertilizer shortage.[5] This provides an ideal setting for quantifying the value of fertilizer while fully accounting for GE effects, the role of trade policy, and the influence of subsidies by using a new quantitative spatial model (QSM) of trade and agriculture.

Our analysis proceeds in four steps. First, we document three stylized facts on the ban's impacts by combining a rich set of novel data sources. To examine the impact on imports, we rely on high-frequency granular customs data on the universe of import transactions in Sri Lanka and on import ban data digitized by hand from regulatory gazettes. The ban was imposed suddenly and later lifted due to social unrest. This motivates our use of the dynamic difference-in-differences (DID) method of De Chaisemartin and d'Haultfoeuille (forthcoming), which corrects the biases in standard DID estimates when the treatment ends during the study period. Our first stylized fact is an estimated reduction in fertilizer import values of more than 99% four months after the ban's introduction (remaining low thereafter), driven by a decline in import quantities. The shift to organic agriculture was minimal, with organic fertilizer's share never exceeding 4% of either fertilizer imports or agricultural use.

To examine the ban's impacts on agricultural production, we focus on Sri Lanka's primary food crop, rice, and a major export crop, tea, both heavily reliant on fertilizer. Using government yield data, our second stylized fact shows that rice yields in the Maha season declined by over 30% between 2021 and 2022. After adjusting for local weather fluctuations by leveraging weather data and our novel pixel-level remote sensing yield estimates from satellite imagery, we still find a decline of 24.4% for rice and 28.7% for tea in 2022 compared to historical norms. This yield drop is accompanied by a significant increase in rice imports, which were minimal before the ban.

To capture the ban's downstream impacts on agricultural exports, we combine customs export data with crop-specific fertilizer use data to create a firm-level exposure measure based on pre-ban crop export portfolios and the fertilizer intensities of crops. Defining treated firms as those above the third quartile of exposure, we use the dynamic DID approach of De Chaisemartin and

[4] Sri Lanka's fertilizer subsidies correspond to 78%-83% of the world price (Kishore et al., 2021), one of the highest rates in the world. In 2020, 53.6% of the government's expenditure on the agricultural sector was on fertilizer subsidies (Weerahewa et al., 2021).

[5] https://www.fao.org/4/ag120e/AG120E12.htm

d'Haultfoeuille (forthcoming) to estimate the ban's effect on firm-level exports. Our third stylized fact is a 35% decline in agricultural exports for highly exposed firms (relative to less exposed firms) four quarters after the ban, a result that is robust to seasonality corrections. While these findings demonstrate the ban's adverse impacts, they cannot quantify its full GE effects on welfare without a theoretical model that accounts for labor, goods markets, and tax rate adjustments. This is especially policy-relevant in contexts where government tools like import restrictions can limit overall fertilizer access, while fertilizer subsidies can influence which groups are most impacted.

Our second step is thus to develop a quantitative GE model of trade and agriculture that accounts for cross-crop heterogeneity in production technology (including fertilizer intensity) and geographic variations in population, productivity, and crop specialization. The model incorporates consequential real-world features that affect welfare during fertilizer shortages, such as non-homothetic preferences, land inequality, QRs on fertilizer imports, and a fertilizer subsidy program. Non-homothetic preferences generate a negative relationship between household income and agriculture's share of expenditure, while QRs restrict fertilizer imports, driving a wedge between international and domestic fertilizer prices. Non-homothetic preferences help the model accurately predict the sizes of the agricultural and manufacturing sectors, with the latter providing an employment alternative for workers during the ban. Subsidies act as a transfer from mobile workers to farmers, and a complete import ban reduces this transfer, benefiting mobile workers. In this way, the import ban, by itself one of the most distortive industrial policies (McDonald et al., 2024), alleviates some of the distortive effects of fertilizer subsidies, another industrial policy (albeit a domestic one).

In the third step, to bring the model to the data, we estimate the model's parameters using a diverse set of empirical strategies and a rich set of novel data sources. We use household survey data to estimate the income elasticity of food's expenditure share (Engel elasticity), a policy-relevant parameter widely used to analyze food security concerns in developing countries (Baquedano et al., 2021), but hard to estimate causally. We follow an instrumental variables (IV) approach using unexpected income shocks (lottery winnings and disaster relief). Survey data on crop prices and household expenditures also enables us to estimate the elasticity of substitution *across crops* through an IV approach. This approach leverages the effect of regional crop suitability, exogenously determined by regional geography and measured by the Global Agro-Ecological Zones (FAO-GAEZ) project, on local crop prices. Technology parameters in agricultural production functions are estimated using input cost shares, computed using national Input-Output (IO) tables and data on fertilizer requirements and subsidy rates. Finally, unobserved variables (e.g. productivities and taste shifters) are backed out through model "inversion" by solving the equilibrium system of equations while taking as given the values of observed variables (e.g. populations, wages, land endowments, crop outputs/prices) in 2019, a pre-ban year that we use as our baseline.

The estimated model is used to understand the aggregate and distributional consequences of the import ban and to evaluate the role of industrial policies such as fertilizer subsidies and import restrictions in shaping these effects. In the main counterfactual exercise, we represent the fertilizer ban as the imposition of a strict QR that drives fertilizer imports down to zero. The existence of a small domestic fertilizer production sector ensures that, even in face of a full import ban, fertilizer use in agriculture does not fall to zero. The model results show that the elimination of fertilizer imports led to a 96.5% reduction in fertilizer use in Sri Lankan agriculture. As a result, crop yields declined by between 12% (for cinnamon) and 54% (for maize), depending on each crop's reliance on fertilizer. Lower production induced a decrease of 596 million US dollars (USD) in agricultural exports, severely limiting the ban's unstated objective of saving foreign exchange through a reduction in the trade deficit. Resulting welfare losses for district-level representative agents were equivalent to a 4.35% income reduction, on average, but incidence was highly heterogeneous. Districts' representative farmers, whose source of income is tied to the agricultural sector, suffered losses equivalent to 9%-23% of their baseline incomes, depending on their geographic location. In contrast, mobile workers, who can switch sectors, saw much more limited losses, which never surpassed 12% even in the worst affected, most agricultural districts, where the employment "buffer" provided by manufacturing is smallest. Farmer losses varied across space and were strongly associated with regional specialization in relatively fertilizer-intensive crops.The model implies that the average income loss in fertilizer-intensive districts was 8.2 percentage points (p.p.) higher than in less fertilizer-intensive districts, a gap that is statistically indistinguishable from the corresponding gap that can be inferred from nightlights data (7.7 p.p.).

In the fourth and final part of the paper, we use the model to analyze the effects of smaller QRs on fertilizer imports, which are more likely policy choices than the drastic full ban implemented in Sri Lanka. The magnitude of the average welfare loss is increasing and convex in the strictness of the QRs, reflecting the presence of decreasing marginal returns to fertilizer use in agricultural production. In other words, a full ban on fertilizer imports is more than twice as harmful as a 50% reduction in the quantity of imported fertilizer. Finally, we run additional counterfactual exercises to demonstrate the importance of modeling real-world features of developing countries, such as non-homothetic preferences, land inequality, and fertilizer subsidies, when studying agricultural trade policies. Without non-homothetic preferences, the size of the agricultural sector is underpredicted, leading to an underestimation of the ban's negative effects. Without subsidies, there is no net income transfer scheme from workers to farmers, so the model overestimates ban-induced worker losses because it misses how the ban helps workers by scaling down this transfer scheme.

Related Literature Our quantification of the costs of lack of access to fertilizer in general equilibrium adds to an extensive literature that evaluates the effects of fertilizer (or, more broadly, of modern agricultural inputs) on yields, which mainly uses randomized control experiments (Beaman et al., 2013; Carter et al., 2021; Duflo et al., 2008, 2011). While small-scale field experiments are essential in assessing local effects, they do not inform on GE effects once the intervention is scaled up (Bergquist et al., 2023; Muralidharan and Niehaus, 2017). For example, while providing fertilizer to a few farmers may not affect the rural labor market, a large-scale change in fertilizer availability may shift rural labor demand enough to affect wages.[6] In contrast, our setting of a nationwide, sudden, and unprecedented fertilizer shortage allows us to estimate the costs of a lack of fertilizer for an entire country. Our paper is complementary to recent work by Bergquist et al. (2023), who develop a spatial model of agriculture and trade to quantify the effects of large-scale agricultural policies such as provision of agricultural inputs. In contrast, we analyze a unique trade policy experiment involving the near-complete removal of fertilizer, rather than the provision of subsidized inputs, whose take-up ultimately depends on the decision of the farmer (Duflo et al., 2011; Suri, 2011). Notably, an underappreciated feature of the global fertilizer market is that most countries rely almost entirely on imports to obtain their fertilizer (Brooks and Donovan, 2024). We provide a framework to model the impact of trade policy on fertilizer access and the associated costs of limited access. Our paper also complements Artuc et al. (2023), who examine the impact of conflict on global agricultural trade, including a simulation where Russia and Ukraine restrict fertilizer exports. While such trade policies affect more countries than in our case, each country suffers less due to the ability to substitute toward unaffected suppliers.

Our theoretical model draws on the growing literature on QSMs focused on agriculture (Aggarwal et al., 2022; Bergquist et al., 2023; Carleton et al., 2024; Costinot et al., 2016; Desmet and Rossi-Hansberg, 2015; Fajgelbaum and Redding, 2022; Farrokhi and Pellegrina, 2023; Nath, forthcoming; Nigai, 2016; Pellegrina, 2022; Sotelo, 2020; Tombe, 2015), extending these frameworks by incorporating trade policy through QRs that impact fertilizer access. It also introduces real-world features that influence welfare during fertilizer shortages, such as fertilizer subsidies, land inequality, and non-homothetic preferences. We use a specific class of non-homothetic preferences, namely Price-Independent Generalized Linear (PIGL) preferences, that is borrowed from the structural transformation literature (Boppart, 2014; Eckert and Peters, 2023). Like Eckert and Peters (2023), we assume a particular parametrization for the income distribution and derive closed-form expressions for aggregate demand and indirect utility, which we use to compute the equilibrium of a spatial model with multiple heterogeneous regions. However, our model includes two types of agent (workers and farmers) and assumes a log-normal income distribution, showing that the aggregation properties of PIGL preferences are not limited to Frechet distributions.

The natural experiment of Sri Lanka's fertilizer import ban relates our work to the literature

[6] A further limitation of these randomized trials is that they are typically conducted on closely supervised experimental plots, so returns may differ from those on real-world farms (Beaman et al., 2013; Duflo et al., 2011).

on the impacts of trade policy and, specifically, of non-tariff measures (NTMs).[7] Most studies have focused on a subset of NTMs – regulatory or anti-dumping measures – often bundling several measures and generally estimating their impacts on the trade flows of the products directly subject to the NTMs (Ederington and Ruta, 2016). Import prohibitions have been much less studied, with only a few case studies on bans of specific agricultural product.[8] Two recent studies that are more related to our work (despite focusing on import licenses rather than import bans) estimate the effects of Argentina's discretionary import approvals on firm-level imports and import prices (Atkin et al., 2024) and on downstream exporters (Bernini et al., 2024). Another recent study, Artuc et al. (2024), examines the impact of Nigerian import bans on prices and welfare; however, in this case, the bans were not effectively enforced. Our paper's contribution to this literature is twofold. First, by exploiting a natural experiment, this is the first paper to estimate the causal impacts of banning the imports of a crucial intermediate input on the trade flows of both the input itself and the downstream exports dependent on it. Second, we add to Atkin et al. (2024) and Dhingra et al. (2023), the only two papers (to our knowledge) that incorporate NTMs in a quantitative GE trade model. While much of the NTM literature models NTMs as an increase in multiplicative "iceberg" trade costs, these modeling choices are not always innocuous (Ghose et al., 2024). Whereas QRs create rents for certain domestic agents, price increases from iceberg trade costs are typically modeled as pure losses. We model NTMs as QRs with arbitrary intensity, implemented through import permits granted to middlemen, thus creating a gap between domestic and international prices while still preserving analytical tractability. Our findings are novel to the NTM literature: in the absence of fertilizer subsidies, our model predicts a decline in downstream agricultural exports that is larger than the initial NTM-driven decline in fertilizer imports, thus jeopardizing the NTM's presumed motivation of saving foreign exchange. Finally, our work also informs the burgeoning debate on the revival of industrial policy: our large welfare loss estimates support the literature's perception that trade policy measures, such as import bans, are among the most distortive measures that fall under the (admittedly contested) definition of industrial policy (McDonald et al., 2024).

The rest of the paper is organized as follows. Section 2 discusses the fertilizer import ban in Sri Lanka. Section 3 presents our rich set of data sources. Section 4 presents reduced-form evidence on the impacts of the ban on imports, production, and exports. Section 5 lays out a QSM of trade and agriculture. Section 6 explains how to estimate the model using Sri Lankan data. Section 7 presents our counterfactual exercises and the corresponding estimates of the ban's effects on production, trade, and welfare across Sri Lankan regions. Section 8 concludes.

[7] Import bans are a control measure aimed at restraining the quantity of goods that can be imported and are thus classified under chapter E of the non-tariff measures MAST classification (UNCTAD, 2019).

[8] See e.g., Felt et al. (2011) on Japan's ban of Taiwanese pork imports; Leroux and Maclaren (2011) on an Australian import ban on bananas, and Peterson and Ord (2008) on a US import ban on avocados from Mexico.

2 Background on the Fertilizer Import Ban in Sri Lanka

On May 6, 2021, Sri Lanka's government abruptly banned the imports of chemical fertilizers with the stated goal of transitioning the country's agricultural sector into organic farming. Although president Gotabaya Rajapaksa had mentioned his aim to provide Sri Lankans with food without harmful chemicals during his 2019 election campaign, the ban was sudden and unexpected.[9] The use of agrochemicals had been linked to chronic health problems and adverse environmental impacts across the country, which was used as a justification for the drastic NTM, but its most pressing rationale was to save the 250 million USD spent annually on fertilizer imports in the context of a large current account deficit that had been exacerbated by the adverse impact of Covid-19 on the tourism industry, one of the country's main sources of foreign exchange. In fact, to address these external challenges the government had already been imposing bans on imports of non-fertilizer products since April 2020 (Dissanayaka and Thibbotuwana, 2021).[10]

Proponents of organic farming in Sri Lanka were appalled at the president's decision to shift into organic farming in a single agricultural season without providing farmers with adequate support and training for the shift. The president of the Lanka Organic Agriculture Movement argued that "The president's committee of advisers to implement the new agricultural policy had no knowledge of organic farming" and that it takes time to produce results: "It needs two to three seasons to develop microbes that enhance soil quality. It is during this time that the farmers needed governmental support. But that support was missing."[11] Since domestic production and imports of organic fertilizer were insufficient, organic fertilizers were in short supply and became very expensive. Some farmers ended up using low-quality fertilizers obtained on the black market. More fundamentally, agricultural specialists argue that, in terms of providing necessary nutrients for crop growth, organic fertilizer can only be a supplement and not a substitute for chemical fertilizer (even if organic fertilizer can help reduce the demand for chemical fertilizer over the long-run).

Quantifying the economic consequences of the agrochemicals import ban is the focus of our paper. Preliminary evidence in FAO and WFP (2022) suggests the ban had dramatically adverse effects on agricultural production in the major 2021-2022 growing season, resulting in large increases in food prices and reduced food security. Farmer protests and general social upheaval ensued, eventually inducing the government to lift the import ban later in 2021. However, while the ban was short-lived, it was in place during a very unfortunate time, around the start of the Maha season, which is Sri Lanka's major rice growing season. Moreover, after the ban ended, international fertilizer prices reached their all-time highs due, in part, to Russia's invasion of Ukraine in

[9] https://www.reuters.com/markets/commodities/fertiliser-ban-decimates-sri-lankan-crops-government-popularity-ebbs-2022-03-03.

[10] The stated policy goal of such import bans "was ambiguous, as policy statements outlined various objectives ranging from minimising foreign exchange leakages to import substitution" (Wijesinghe et al., 2023).

[11] https://asia.nikkei.com/Spotlight/Sri-Lanka-crisis/Sri-Lanka-aims-for-food-security-after-ill-fated-fertilizer-ban.

February 2022. Given Sri Lanka's low foreign reserves and a large depreciation of the Sri Lankan rupee (LKR), the shortage of agrochemicals continued into 2022, with adverse consequences for agricultural production and food security.[12]

2.1 Organic and Inorganic Fertilizer

The purpose of fertilizer additions is to provide plants with nutrients (mostly nitrogen, phosphorus, potassium) that are not sufficiently naturally occurring in the soil. However, there are important differences between organic or inorganic (chemical) fertilizers. Organic fertilizers are slow-releasing and may require the correct amount of heat and moisture for nutrient extraction at the time of plant need. Since this need is met with almost the right timing and amounts, leaching is minimized.[13] In contrast, inorganic fertilizers provide nutrient inputs rapidly because they are in plant uptake-ready form, but this can come at the expense of leaching, especially if timed incorrectly.

Since the ban only applied to agrochemicals, one may ask if the decreased use of inorganic fertilizer may have been offset by an increased use of organic fertilizer. We show this offset was very limited by examining the use of manure, a typical organic fertilizer. The quantity of manure applied to the soil had been increasing steadily over the twenty years prior to the ban and, importantly, its increase in 2021 merely followed this long-run trend, without any visible "jump" that would occur if farmers resorted to increased use of organic fertilizer as a key coping strategy to deal with the ban (Appendix Figure A3a). Moreover, manure's share of total fertilizer use was not atypically high in 2021 compared to the previous six years and this share was never higher than 2.5% (Appendix Figure A3b).

Trade data also confirms that the collapse of chemical fertilizer imports, which decreased to less than 40 million USD in 2021 from a pre-ban baseline of 100-250 million USD, was hardly compensated by organic fertilizer: while the organic share of fertilizer imports did increase after the ban, it was never higher than 4% (Appendix Table A3).

3 Data

Our reduced-form and model-driven analyses rely on a rich assortment of mostly novel data. Table 1 presents our main data sources and variables, grouping them into several categories: trade policy; trade flows; agricultural production, cultivated area, and yields; agricultural prices; fertilizer use and subsidies; potential agricultural yields; labor share in value added; wages, land ownership,

[12] Over 2022, the rupee devalued by nearly 45% against the USD.

[13] Leaching is the removal of excess nutrients from the plant/soil system. It occurs when water removes water-soluble nutrients out of the soil, by runoff or drainage. Leaching is an environmental concern if chemical fertilizers find their way into water bodies.

occupation, and household characteristics; and night lights. Appendix A provides additional details on all data sources and variables used. Below, we briefly highlight a few important features of our data.

Table 1: Data Sources and Variables

Topics	Variables	Time Period	Source
Trade Policy	**[1]** List of Harmonized System 8-digit (HS8) codes of products subject to import bans, and bans' start and end dates	March 2020-October 2022	Sri Lanka's Extraordinary Gazettes on Imports and Exports, digitized by hand
Trade Flows	**[2]** Import values and weights at HS8 product-month-year level	January 2017-October 2022	S&P Global Market Intelligence's Panjiva data platform
	[3] Export values at exporting firm-HS8 product-quarter-year level	January 2017-October 2022	S&P Global Market Intelligence's Panjiva data platform
Agricultural production, cultivated area, yields	**[4]** Rice production, cultivated area, and yields at district-season level	Maha 2013 to Yala 2022	Sri Lanka's Department of Census and Statistics (DCS)
	[5] Production, cultivated area, and yields by crop at district-season level (maize, groundnuts, potatoes, onions, cinnamon, cloves, tea)	Maha 2020 to Yala 2022	DCS and Sri Lanka Tea Board's annual reports
	[6] Production by crop at national level	2022	Central Bank of Sri Lanka
	[7] Remote sensing estimates of rice yields and cultivated areas at pixel level	2000-2022	Landstat and Sentinel-2 satellite, computed by authors using Ozdogan et al. (2024) methodology
	[8] Remote sensing estimates of tea yields at pixel level	2014-2022	Planet Satellite Data, computed by authors
	[9] Remote sensing estimates of tea map at pixel level	2021	Planet Satellite Data, computed by authors
Agricultural prices	**[10]** Producer prices at crop-district level for rice, potatoes, cinnamon, cloves, onions, groundnuts, maize	2019	DCS
	[11] Producer prices at national level for tea	2019	Food and Agriculture Organization of the United Nations (FAO)
	[12] National Consumer Price Index	2021-2022	DCS
	[13] Crop prices at national level for cinnamon, cloves, maize, onions, potatoes, rice, tea	2022	Central Bank of Sri Lanka
Fertilizer use and subsidies	**[14]** Fertilizer requirements by crop	2022	Sri Lanka's National Fertilizer Secretariat (NFS)
	[15] Fertilizer prices	2022	NFS
	[16] Cultivated area at national level by crop	2022	NFS
	[17] Use of organic fertilizer (manure) at national level	2001-2021	FAO
	[18] Use of chemical fertilizers at national level	2002-2021	FAO
	[19] Fertilizer subsidy rates at crop level	2013-2022	Department of Agriculture's Socio Economics & Planning Centre
Potential agricultural yields	**[20]** Potential attainable yields at crop-district level (all crops rainfed except rice, which is irrigated; climate conditions of 1981-2010)		FAO-GAEZ
Labor share in value added	**[21]** Employee compensation and value added by agricultural sector	2000	Sri Lanka's IO table of the Institute of Policy Studies (Amarasinghe and Bandara, 2005)
Wages, land ownership, occupation, household characteristics	**[22]** Average wages at district level and total number of mobile workers, estate workers, and farmers at district level	2016, 2019	Sri Lanka's Household Income and Expenditure Survey (HIES), from DCS
	[23] Household size (number of members), income, expenditure patterns (food prices and quantities, fraction of income spent on food), "income by chance or *ad hoc* gains" (including lottery and disaster/relief payments)	2019	HIES
	[24] Land ownership and land distribution	2016	HIES
Night Lights	**[25]** Night light intensity at district-month level	January 2016-December 2022	Visible and Infrared Imaging Suite (VIIRS) Day Night Band (DNB), from Colorado School of Mines

Trade-related variables. The list of products targeted by import bans and their time coverage are hand-coded from Sri Lanka's Extraordinary Gazettes on Imports and Exports. The fertilizer import bans started in May 2021, with most bans lifted by July 2021 or by November 2021 at the latest. We also use novel high-frequency granular data covering the universe of Sri Lanka's export and import transactions, relying on machine learning and text analysis to clean firm names and assign unique firm identifiers to create a panel of firms involved in exports.

Agricultural production. We complement administrative data on crop-district-level yields with remote sensing estimates of rice and tea yields at a highly granular level. Remote sensing provides a spatially granular view of crop cultivation, allowing a detailed assessment of the heterogeneous effects of the fertilizer import ban within Sri Lanka. For rice, we apply an expert-based image classification algorithm to satellite data to identify rice-growing areas using the specific growing conditions for paddy rice in Sri Lanka (Ozdogan et al., 2024).[14] To estimate rice yields, a statistical model correlates district-level yields from government statistics with a satellite-derived vegetation index (green chlorophyll index) known to be sensitive to rice yields for pixels identified as rice-growing areas. A random forest-based machine learning model is added to incorporate additional environmental variables known to affect rice yields, resulting in pixel-level rice yield estimates across two decades (2000-2022).[15] For tea, district-level yield data is regressed on a district-level vegetation index derived from satellite data for pixels deemed to be tea plantations.[16]

Fertilizer subsidies. Crop-level fertilizer subsidy rates are defined as the "nominal rate of protection at farm gate", which is computed by comparing the actual (farm-gate) fertilizer price to the price that would be expected given international fertilizer prices and the cost of bringing fertilizer from port to farm. The fertilizer subsidy rate is 81% for rice and 45%-54% for four other crops of interest (potato, maize, tea, and onions).[17]

Potential agricultural yields. Our use of FAO-GAEZ data for crop-district-level potential attainable yields follows the literature on agriculture and trade (Bustos et al., 2016; Costinot et al., 2016; Farrokhi and Pellegrina, 2023; Sotelo, 2020).

Wages and household characteristics. Three types of information in our survey data are particularly valuable: (i) on estate workers who live and work in Sri Lanka's tea plantations, for whom changing jobs is nearly impossible in the short run as their housing is tied to the tea estate where they work (Chandrabose, 2015); (ii) on the fraction of income households spend on food; and (iii) on income earned "by chance", including lottery gains and disaster/relief payments.

[14] The presence of rice-growing areas is determined by factors related to climate, water availability, and topographic position, along with farmer decision-making and technical expertise. The methodology monitors rice-related vegetation and water management from space and distinguishes it from other land use types using spectral (color) and temporal information included in the satellite signal. Because water and vegetation indexes are more visible for rice than for other crops, crop type mapping and yield are better predicted for rice than for other crops using remote sensing tools (Dong et al., 2016).

[15] These remote sensing estimates, aggregated to the DS level, are shown to be highly consistent with government-reported measures of rice area (from survey-based crop-cutting experiments) and rice yields (from production statistics).

[16] Since tea is a multiyear permanent crop, planted areas do not change annually. Thus, we fix tea-growing areas to the 2021 tea map.

[17] High fertilizer subsidy rates, in the 40-50% range, are also reported in the World Food Programme's analyses of food security in Sri Lanka (FAO and WFP, 2022, 2023).

Summary Statistics. Fertilizer and non-fertilizer products whose imports were banned accounted for 15%-19% of total imports in 2017-2019 (Appendix Table A4). This share declined to 12% in 2020 and 10% in 2021, suggesting the bans were at least partially effective. Summary statistics for some key variables of interest, shown in Appendix A.12, provide evidence of the ban's effects: the average value of fertilizer imports per firm-year plummeted from 2 million USD to only 491 thousand USD from the pre-ban to the post-ban period, while the average value of agricultural exports per firm-year fell from 1.62 million USD to 1.53 million USD. Remotely sensed crop yields decrease substantially from the pre-ban (2021) to the post-ban (2022) period: average rice yields across divisional secretariats (DSs) fall from more than 4,000 kg/ha to less than 3,000 kg/ha, while average tea yields fall from 1,644 kg/ha to 1,404 kg/ha.

4 Stylized Facts on Fertilizer Imports, Agricultural Production and Exports

In this section, we present three stylized facts related to the fertilizer import ban's adverse effects on fertilizer imports, crop yields, and exports of agricultural goods that utilize fertilizer.

4.1 Fact 1: Fertilizer Imports Declined After the Ban

We start by examining whether the import ban was effective in restricting imports. To do this, we estimate dynamic DID regressions using import data at the product-month level, where a product is defined as a unique Harmonized System 8-digit (HS8) code. The treated group consists of fertilizer products whose imports were banned after April 2021 and the control group consists of products that were not banned (both non-banned fertilizers and other types of non-banned products). We compare fertilizer imports in the treatment group with their control group counterparts following the DID methods proposed by De Chaisemartin and d'Haultfoeuille (forthcoming), which allow units to switch in and out of treatment. This is key for our study as the import ban was lifted less than one year after it was first introduced.[18] Our dynamic DID estimator recovers the evolving effects of the ban over time on two outcome variables: a dummy indicating non-zero imports (i.e., extensive margin) and the log of import value (i.e., intensive margin) at the product-month level between January 2019 and September 2022. Estimation details are provided in Appendix B.1.

Figure 1 plots the DID results, showing a clear decline in imports right after the ban was imposed, with the average probability of nonzero imports across banned products dropping by 31 percentage points in the second month of the ban (Figure 1a). Along the intensive margin, negative

[18] See Callaway (2023) for a complete review of DID methods.

Figure 1: Dynamic Ban Effects on Imports

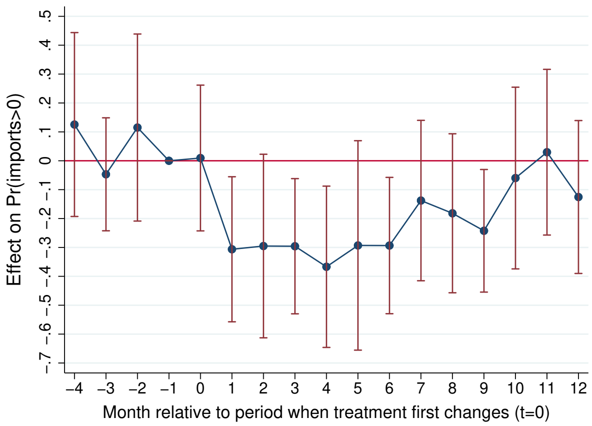

(a) Extensive Margin

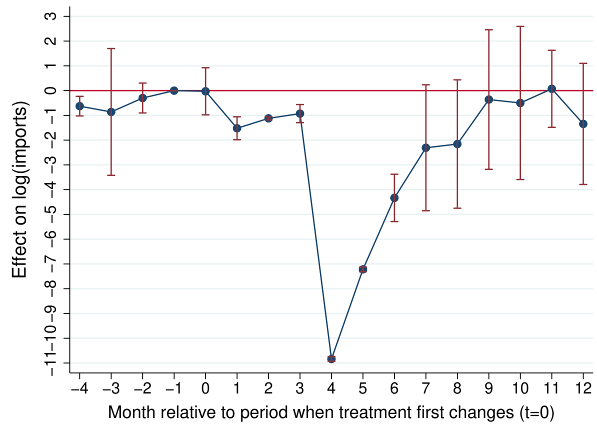

(b) Intensive Margin

Notes: The figures show De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of import bans on the extensive and intensive margins of imports, relative to the last month before the ban takes effect, using an HS8 product-month level panel. The sample covers the January 2019-September 2022 period (but only coefficients in January 2021-May 2022 are reported) and excludes non-fertilizer HS8 products whose imports were ever banned. The treatment variable is a set of dummies ban^{τ}_{ct} indicating whether fertilizer product c was first banned τ months before month t, and the not-yet-treated products serve as the control group, which includes both non-banned fertilizers and other products. The y-axis shows the effects of fertilizer import bans on the probability of nonzero imports (panel 1a) or on the log value of imports (panel 1b), where negative values correspond to a decrease in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.70 (panel 1a) and 0.0001 (panel 1b).

effects peaked four months after the ban was introduced, with imports falling by 99.998% (Figure 1b) and remaining low for the next several months.[19]

We also obtain DID estimates based on a smaller sample including only fertilizer products, thus comparing how the imports of banned versus non-banned fertilizers evolve over time after the ban is introduced. The results, albeit less precisely estimated, are qualitatively maintained (Appendix Figure A4). Due to complementarities in production, banning the imports of chemical fertilizers may also have negatively affected other fertilizers. On the intensive margin, a large difference in imports across banned and non-banned fertilizers emerged four months after the ban was introduced and persisted for several months thereafter.[20]

Finally, we establish the robustness of our results to concerns about seasonality and the concurrent Covid-19 pandemic shock. Our results are robust to using deseasonalized versions of our outcome variables (Appendix Figure A6).[21] The Covid-19 pandemic, still ongoing in 2021, would

[19]The coefficient for month $t = 4$ in Figure 1b equals -10.8. Given the log specification, the corresponding percentage effect on imports can be recovered through the expression: $(\exp(-10.8) - 1) \times 100\% \simeq -99.998\%$.

[20]As an additional check, we also obtain DID estimates using the log of imported product *quantities* as the outcome variable. The ban's effects on fertilizer imports remain very similar when imports are measured in kg instead of dollars (Appendix Figure A5). This is unsurprising, as Sri Lanka's economy is arguably too small to influence the prices of internationally traded commodities such as fertilizer.

[21]Both outcome variables (extensive and intensive margins) are deseasonalized by subtracting their HS8 product-

only bias our estimates if it affected imports of banned and non-banned fertilizers to different degrees, which we do not have a strong reason to expect. Still, we implement a falsification test by obtaining DID estimates with the treatment period redefined to start in March 2020 (when Covid-19 arrived) instead of May 2021. The results show that fertilizers who would eventually be banned in 2021 did not suffer a relative import decline in the early months of the pandemic (Appendix Figure A7).

4.2 Fact 2: Rice and Tea Production Declined Significantly After the Ban

We start by examining the production of Sri Lanka's primary food crop, rice, covering about 40% of total arable land. Figure 2a plots the year-to-year evolution of average rice yields in the main (Maha) growing season.[22] From a high of around 4,500 kilograms per hectare (kg/ha) in 2020, Maha rice yields declined precipitously to below 3,000 kg/ha in 2022. This represents a 31.6% decline in yield relative to the historical average of the previous nine years. The observed rice yield decline is very similar when we use our granular remote sensing estimates. This decline cannot be explained by observed weather variables that typically influence agricultural production (such as temperature, rainfall, sunlight, and vapor pressure): even after controlling for fluctuations in such variables, we still find a 24.4% yield decline relative to the historical average (Appendix C.4).[23] Figure 2b shows that rice imports were negligible before the import ban, as Sri Lanka had achieved self-sufficiency in rice production, but, starting in late 2021, rice imports surged up to January 2022, remaining elevated throughout the rest of the year. Even this surge was not enough to fully make up for lost production, though, with the availability of rice for domestic consumption falling by 22% between 2021 and 2022.[24]

Figure 3 presents rice yields at the DS level for the last Maha season before the ban (September 2020-March 2021) and the first Maha season after the introduction of the ban (September 2021-March 2022). Rice yields fell in almost all regions, with some spatial heterogeneity in the degree of yield reduction. The fall was more pronounced in Sri Lanka's northern and western regions, including major rice-producing districts such as Ampara, Anuradhapura, Batticaloa, and Polonnaruwa.

month-specific averages in 2017-2019, before the fertilizer ban.

[22] The Maha season is the main rice cultivation season in Sri Lanka, running from September to March of the following year (the second season, Yala, runs from May to August).

[23] Ozdogan et al. (2024) confirms that the notable decline in remote sensing rice yield estimates over the 2021-2022 period cannot be explained by variation in environmental variables and is directly correlated with the reduction in fertilizer availability ensuing from the import ban.

[24] We define the amount of rice available for domestic consumption in a given year as the sum of production and net imports. This amount fell from 5.2 billion kg in 2021 to 4 billion kg in 2022. Since we lack trade data for November and December 2022, we instead compute net imports for 2022 based on the period from November 2021 to October 2022. Correspondingly, net imports for 2021 are computed based on the period from November 2020 to October 2021. If there had been no change in net imports, rice availability in 2022 would have been 3.40 billion kg. Therefore, the import surge offset one third of the production losses.

Figure 2: Rice Yields and Imports in Sri Lanka over time

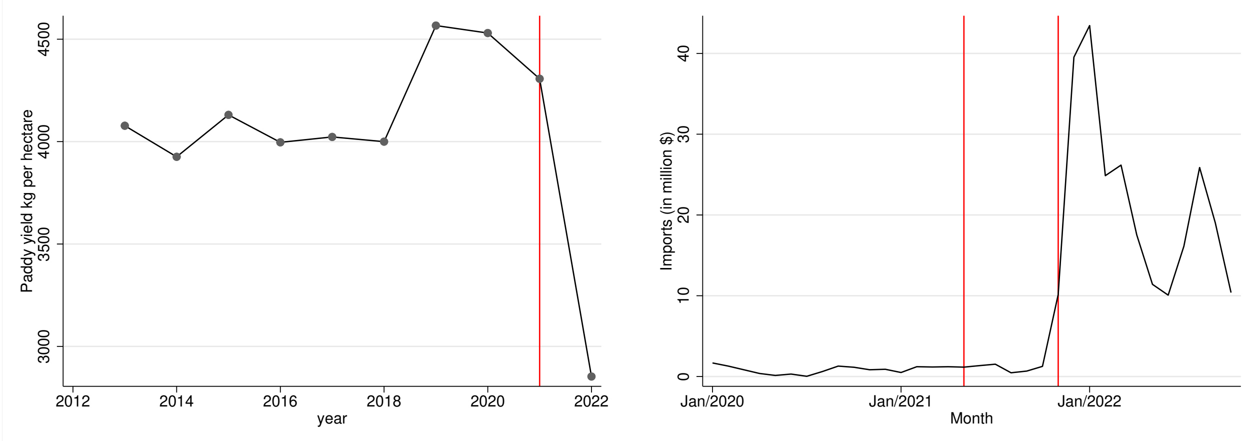

(a) Rice Yields (Maha season)

(b) Rice Imports

Notes: Panel 2a shows the rice yield (in kg/ha) for the Maha season between 2013 and 2022. The Maha season for rice is from September to March of the following year. The red line in 2021 marks the year fertilizer import bans were introduced. The data comes from the DCS. Panel 2b shows monthly rice imports, in millions (MM) of USD. The first red line in May 2021 marks the beginning of the ban, and the second red line in November 2021 marks the end of the ban. The data comes from the S&P Global Market Intelligence's Panjiva data platform.

Figure 3: Rice yields (Maha season) in Sri Lanka

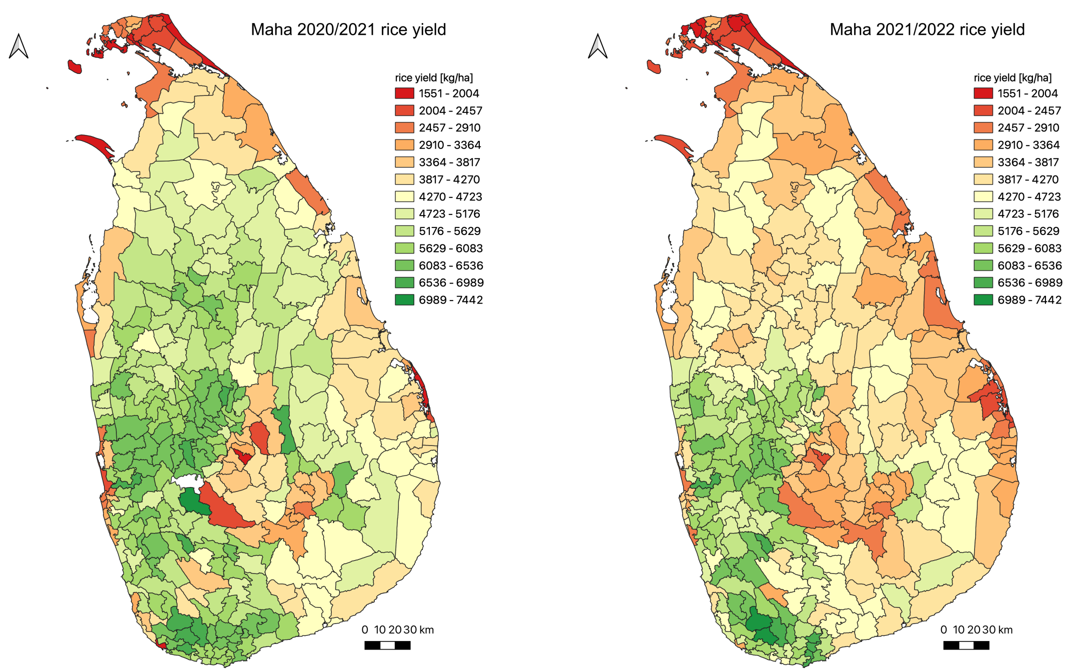

Notes: The left panel shows rice yields (in kg/ha) in the Maha 2020-2021 season, while the right panel shows the same variable but for the Maha 2021-2022 season. The Maha season for rice runs from September of a year to March of the following year. The data comes from the Ozdogan et al. (2024) remote sensing estimates of rice yields.

We also examine yields of Sri Lanka's major export crop, tea. Tea yields suffered a large drop of 28.7% in 2022 (after controlling for weather fluctuations) relative to historical averages, as seen in Appendix C.4. This evidence shows that the fertilizer ban harmed Sri Lankan agriculture substantially, consistent with the importance of fertilizer as an import determinant of agricultural productivity across the world (Donovan, 2021; Gollin et al., 2014).

4.3 Fact 3: Agricultural Exports Declined After the Fertilizer Import Ban

Facts 1 and 2 show that fertilizer imports and rice and tea production declined, and rice imports increased, with the fertilizer import ban. We now show the ban also had adverse downstream effects on agricultural exports.

To assess the ban's effects on the export performance of agricultural firms, we consider the sample of firms who exported in 2017-2019 (i.e. pre-ban) at least one of the 23 crops with fertilizer intensity data. To understand which firms were more exposed to the ban, we define a firm-level exposure measure capturing how much fertilizer is embodied in the crops each firm exports. This exposure measure is the weighted average of the firm's exported crops' fertilizer intensities, where the weights are given by the share of each crop in the firm's agricultural exports portfolio pre-ban.[25] We rank firms by fertilizer exposure and compute the third quartile of that measure. We define exporting firms to be treated if their fertilizer exposure is above the third quartile. As in Fact 1, we use the De Chaisemartin and d'Haultfoeuille (forthcoming) estimator to allow for the fact that the bans were switched on and off. Our dynamic DID estimator recovers the evolving effects of the ban over time on the log of firm export value at quarter level in the period between January 2019 and September 2022.[26] An advantage of our approach is that the firm's potential exposure to fertilizers is computed using information on its agricultural exports and on each product's fertilizer requirements. A disadvantage is that the control group may be imperfect, since all agricultural exporters might have been affected to some degree. If anything, this will work against our finding of any effect of the ban on exports. Appendix B.2 provides further details on the firm-level fertilizer exposure measure and the estimating equation.

Figure 4 presents the DID results, showing that, three quarters after the fertilizer ban was implemented (i.e. in Q1/2022), agricultural exports significantly declined for firms with a high pre-ban fertilizer exposure (relative to low-exposure firms). The magnitude of the estimated effects corre-

[25] Formally, the exposure measure for firm f is given by: $U_f = \sum_{v=1}^{23} FIC_v \times \left(X_{fv}^{2017-2019} / \sum_{v'} X_{fv'}^{2017-2019}\right)$, where v indexes crops, FIC_v is fertilizer intensity of crop v (in kg/ha), and $X_{fv}^{2017-2019}$ is the firm's total exports of crop v in the 2017-2019 period.

[26] The use of a quarterly, rather than a monthly time frequency as in our previous import regressions, has two justifications. First, fertilizer inventories might have been available which would take time to be depleted. Second, the agricultural cultivation cycle takes several months from fertilizer application to harvest and export. Our specification allows export effects to materialize some time after the ban was implemented.

sponds to a 35% reduction in firms' agricultural export value in that quarter (or, correspondingly, 8.6% of their annual export value).[27] The effect's timing is unsurprising, as Sri Lanka's major (Maha) growing season ends in Q1. We also show that our findings are robust to seasonality concerns: using a deseasonalized version of the outcome variable, we obtain a similar negative effect of the ban on downstream agricultural exports (Appendix Figure A8).

Figure 4: Dynamic Ban Effects on Firms' Agricultural Exports (High vs Low Fertilizer Intensity)

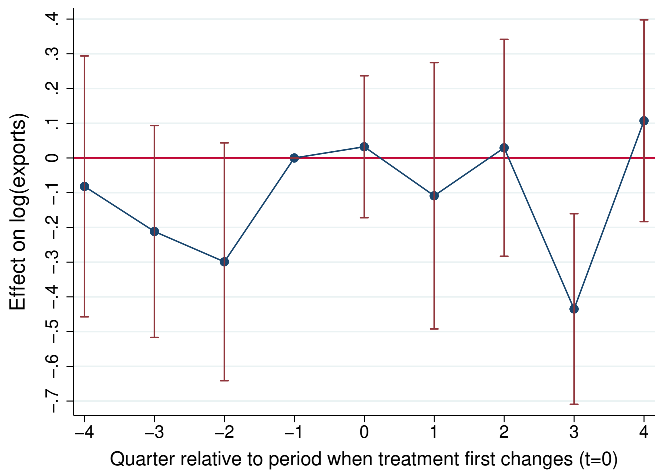

Notes: The chart shows De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of fertilizer import bans on firms' agricultural export value, relative to the last quarter before the ban takes effect, using a firm-year level export panel. The sample covers the Q1/2019-Q3/2022 period (but only coefficients in Q2/2020-Q2/2022 are reported) and includes all firms who exported at least one of 23 crops for which we have fertilizer intensity data in 2017-2019. Each firm's fertilizer exposure is computed as a weighted average of crops' fertilizer intensities with weights given by each crop's share in the firm's 2017-2019 exports. The treatment group consists of firms whose fertilizer exposure is above the sample's 75th percentile. Firms below the 75th percentile serve as the control group. The treatment period is defined as the set of all quarters when any fertilizer import bans were in place. The y-axis shows the effects of fertilizer import bans on log value of agricultural exports, where negative values correspond to a decrease in exports. The p-value of the joint significance test for the placebo estimators is 0.22.

While these three stylized facts demonstrate the fertilizer ban's impact on imports, agricultural production, and exports, it is impossible to quantify the GE effects of the resulting lack of access to fertilizer and the associated changes in welfare without a theoretical model.

5 A Quantitative Spatial Model of Trade and Agriculture

This section introduces a quantitative GE spatial model of trade and agriculture to analyze the ban's effects on relative prices, income, and welfare across regions and agent types in Sri Lanka.

[27] To compute annualized export losses, we use the fact that, in our sample of high-exposure (i.e. "treated") firms, the first quarter of 2019 accounted for 24.5% of that year's total exports. Assuming this same percentage would have held in 2022 in the absence of the ban, a 35% export loss in Q1/2022 translates into a 8.6% loss in 2022 as a whole. In any case, the DID nature of this exercise implies that both the 35% and 8.6% figures can be considered lower bounds for the ban's total effects, which also include GE effects. This distinction also explains the gap between our reduced-form export loss estimate (8.6%) and the corresponding estimate (22.6%) we will show in our model-driven counterfactual exercise (see Table 4 in Section 7.2).

The model highlights the policy relevance of government tools like import restrictions, which limit fertilizer access, and subsidies, which in turn shape the distributional impacts of limited fertilizer access. In section 7, we use the model to examine the welfare effects of the ban under counterfactual scenarios.

5.1 General Environment

In the model, the world is composed of R geographic regions. There are $I = R - 1$ regions within Sri Lanka, plus one additional region representing the rest of the world (RoW).[28] Each region i has a fixed population of size N_i which cannot move across regions.[29] The population is composed of three types of agents: N_i^e estate workers (who can only work in tea agriculture), N_i^m sectorally mobile workers (who can work in manufacturing or in any agricultural sector except tea), and N_i^F landowning farmers, with $N_i^e + N_i^m + N_i^F = N_i$. Additionally, there is a set of Sri Lankan middlemen who are in charge of domestic fertilizer markets (i.e. they own the domestic fertilizer endowment and also import foreign fertilizer to sell in Sri Lanka).[30]

There are two main economic sectors, agriculture (**A**) and manufacturing (**M**). Agriculture encompasses a finite set of K crops. Each crop (as well as the manufacturing good) is differentiated by regional origin (i.e., the Armington assumption). Furthermore, there is an additional minor sector, fertilizer (**f**), which is a single homogeneous good that is exclusively used as an intermediate input in agriculture. All markets are perfectly competitive.

5.2 Production

This section presents our assumptions concerning production and trade costs in each of the three sectors of the economy: agriculture, manufacturing, and fertilizer.

5.2.1 Agricultural Production

We model agricultural production similarly to Farrokhi and Pellegrina (2023). Each region i has an agricultural land mass of area L_{ik} that can only be used to grow crop k.[31] This land mass is

[28]When bringing the model to the data, the regions within Sri Lanka are its 25 districts (that aggregate across DSs).

[29]We argue this immobility assumption is reasonable because our empirical implementation of the model focuses on a relatively short time horizon of less than two years.

[30]For modelling convenience, we assume the population of middlemen has measure zero, so they can be omitted from calculations of average welfare.

[31]This assumption implies that land cannot be reallocated across crops, which is consistent with the short-term horizon of our empirical implementation of the model. Importantly, this assumption finds strong support in the Ozdogan et al. (2024) remote sensing estimates of cultivated rice areas in Sri Lanka, which indicate an extremely high correlation in the pixels classified as growing rice in each year relative to the previous year (for either Maha or Yala seasons). We also provide evidence of no post-ban reduction in cultivated rice area in Appendix C.4.

composed of a continuum of plots indexed by ω, each of size $l_{ik}(\omega)$, which can be used in combination with labor and fertilizer inputs to produce agricultural goods. Specifically, by combining a land portion of size $l_{ik}(\omega)$ within plot ω with labor input $n_{ik}(\omega)$ and fertilizer input $f_{ik}(\omega)$, it is possible to produce an amount $q_{ik}(\omega)$ of crop k according to the following Cobb-Douglas production function:

$$q_{ik}(\omega) = T_{ik}^{A}(n_{ik}(\omega))^{\gamma_k^n}(f_{ik}(\omega))^{\gamma_k^f}(l_{ik}(\omega))^{\gamma_k^l} \tag{1}$$

where T_{ik}^{A} is a land productivity parameter and (γ_k^n,γ_k^f,γ_k^l) are Cobb-Douglas coefficients.[32] These assumptions allow fertilizer dependence to vary across crops (since parameter γ_k^f depends on k), which was shown earlier to be the case. Furthermore, since productivity parameter T_{ik}^{A} varies across regions, they may differ in terms of absolute and comparative advantage within agriculture.

Individual production decisions can be aggregated to obtain crop production and land rents at the regional level. For a plot ω producing crop k in region i, the optimal unit production cost $c_{ik}(\omega)$ that results from solving its cost minimization problem can be written as:

$$c_{ik}(\omega) = \kappa_k (w_i^k)^{\gamma_k^n}(r_{ik}(\omega))^{\gamma_k^l}(p_i^{f,k})^{\gamma_k^f}\ (T_{ik}^{A})^{-1} \tag{2}$$

where w_i^k and $p_i^{f,k}$ are the relevant wage and fertilizer price (respectively) for crop-k agriculture, $r_{ik}(\omega)$ is the rental rate of plot ω, and κ_k is a composite parameter.[33] Perfect competition then implies that the unit cost $c_{ik}(\omega)$ must equal the local crop price (p_{ik}) in equilibrium.[34] This equality then allows us to solve for land rent $r_{ik}(\omega)$ as:

$$r_{ik}(\omega) = (T_{ik}^{A})^{\frac{1}{\gamma_k^l}} \underbrace{(p_{ik})^{\frac{1}{\gamma_k^l}}(w_i^k)^{\frac{-\gamma_k^n}{\gamma_k^l}}(p_i^{f,k})^{\frac{-\gamma_k^f}{\gamma_k^l}}(\kappa_k)^{\frac{-1}{\gamma_k^l}}}_{\equiv h_{ik}} \tag{3}$$

with composite variable h_{ik} "summarizing" the prices that are relevant for the plot. Aggregating over all plots that grow crop k in region i, we can write the total land rent R_{ik} earned by them as:

$$R_{ik} = (T_{ik}^{A})^{\frac{1}{\gamma_k^l}} h_{ik} L_{ik} \tag{4}$$

Due to the properties of the Cobb-Douglas production function, we also know that total payments

[32]We assume agricultural production functions have constant returns to scale, so $\gamma_k^n + \gamma_k^f + \gamma_k^l = 1$ for each crop k.

[33]$\kappa_k \equiv [(\gamma_k^n)^{\gamma_k^n}(\gamma_k^l)^{\gamma_k^l}(\gamma_k^f)^{\gamma_k^f}]^{-1}$. Tea agriculture employs estate workers, so its relevant wage w_i^{tea} equals w_i^e, the wage rate of estate workers in region i. For all other crops k other than tea, $w_i^k = w_i^m$, where w_i^m is the wage rate of mobile workers in region i. The relevant fertilizer price $p_i^{f,k}$ can vary across crops because different crops are subsidized by the government at different rates (see Section 5.4).

[34]The intuition for this equality is that multiple agricultural entrepreneurs compete for the right to farm plot ω by bidding up the land rent until the ensuing unit cost equals the crop price.

to land must equal a share γ_k^l of the agricultural revenue from crop k in region i:

$$R_{ik} = \gamma_k^l p_{ik} Q_{ik} \tag{5}$$

where Q_{ik} is the total physical production of crop k in region i. Combining equations (4) and (5), we can solve for production Q_{ik} as:[35]

$$Q_{ik} = (\gamma_k^l)^{-1} (T_{ik}^A)^{\frac{1}{\gamma_k^l}} h_{ik} L_{ik} (p_{ik})^{-1} \tag{6}$$

We assume that a crop k produced in region i can be either sold locally or exported to other regions. However, to export to region n the seller incurs a multiplicative "iceberg" trade cost $\tau_{ni,k}^A \geq 1$. Therefore, marginal-cost pricing implies that the crop price at the destination must be:

$$p_{ni,k} = \tau_{ni,k}^A p_{ik} \tag{7}$$

Finally, it is convenient to define an aggregate land rent variable R_i combining the land rent from all crops in region i:

$$R_i = \sum_{k=1}^{K} R_{ik} \tag{8}$$

5.2.2 Manufacturing Production

We assume that manufacturing production uses labor as its sole input. Manufacturing output q_i^M in region i is a simple linear function of regional manufacturing employment (n_i), with a multiplicative constant given by productivity parameter T_i^M.

As in agriculture, the manufacturing good produced in a region i can be either consumed locally or shipped to another region n, with the latter incurring an iceberg trade cost $\tau_{ni}^M \geq 1$. Given these assumptions, marginal-cost pricing (due to perfect competition) implies that the unit price of the manufacturing good produced in i and offered for sale in n must be:

$$p_{ni}^M = \frac{\tau_{ni}^M w_i^m}{T_i^M} \tag{9}$$

where w_i^m is the wage rate of mobile workers in region i.

[35] Intuitively, equation (6) shows that equilibrium regional crop production depends positively on land productivity, land endowment, and crop prices, and negatively on labor and fertilizer prices.

5.2.3 Fertilizer Endowments

As in Farrokhi and Pellegrina (2023), we assume that the quantity of fertilizer Q_i^f owned by agents who live in region i is simply given as a natural endowment F_i (rather than being produced). Moreover, fertilizer is a tradable homogeneous good, so *in the absence of import restrictions* its price must follow the law of one price. This implies that fertilizer price p_i^f in region i must equal the international fertilizer price (p_{RoW}^f).[36]

To model the import ban, we allow the Sri Lankan government to restrict fertilizer imports through the use of QRs. To implement these QRs, the government distributes import licenses to middlemen, allowing them to import up to $\bar{f}$ kg of foreign fertilizer.[37] If the QRs are binding, then the domestic fertilizer price across Sri Lankan regions ($\{p_i^f\}_{i=1}^I$) will be a number p_{LKA}^f strictly greater than the international price p_{RoW}^f, with the price difference being pocketed by the middlemen in the form of license rents QR^{rent}:[38]

$$QR^{rent} = (p_i^f - p_{RoW}^f)\bar{f} \tag{10}$$

5.3 Consumers

The N_n agents living in region n earn income and consume. Mobile workers inelastically supply one unit of labor to local firms and are perfectly mobile across sectors, thus earning the local mobile wage w_n^m independently of the sector in which they work.[39] Estate workers also supply one unit of labor inelastically, but they can only work in tea agriculture, earning the local estate wage w_n^e. Farmers, who own land, rent out this land for agricultural production, earning land rents. Finally, there is a measure-zero set of middlemen who own Sri Lanka's fertilizer endowment and the fertilizer import licenses issued by the government (if any), thus deriving their income from fertilizer sales and license rents.

5.3.1 Land Ownership

We allow for inequality in land ownership. Denoting the total extension of land owned by farmer h in region n as L_n^h, we assume that variable L_n^h follows a log-normal distribution with parameters

[36] Variable p_i^f is defined as the price of fertilizer *gross* of subsidies (see Section 5.4).

[37] Whether or not there are QRs in place, we assume that Sri Lanka is a net importer of fertilizer, which is consistent with the observed preponderance of imported chemical fertilizer in its agricultural sector. For example, in 2019 Sri Lanka imported 210 million USD in chemical fertilizer while only exporting 480,749 USD.

[38] Note that variable QR^{rent} is the *total*, country-wide license rent, which is then distributed equally among middlemen. We assume that Sri Lankan middlemen have the same spatial distribution as Sri Lankan mobile workers. As a consequence, regional license rents are given by: $QR_i^{rent} = QR^{rent} \times N_i^m / (\sum_{k=1}^I N_k^m)$.

[39] We relax the perfect mobility assumption in Appendix H by introducing frictions to cross-sector mobility.

μ_n and σ^2_{Ln}.[40] We also assume that, conditional on her total land size L^h_n, each farmer receives a random sample of plots. These assumptions imply the land rent R^h_n received by a farmer h in region n is proportional to her land size:

$$R^h_n = L^h_n \frac{R_n}{L_n}$$

where $L_n = \sum_k L_{nk}$ is total arable land in region n, and R_n is the aggregate land rent paid to all land in region n. Given the log-normal distribution of land size, it follows that land rent in region n also has a log-normal distribution, as follows:

$$R^h_n \sim \text{log-N}\big(\mu_n + \ln(R_n) - \ln(L_n), \sigma^2_{Ln}\big) \tag{11}$$

5.3.2 Consumption Across Sectors

Preferences for agricultural versus manufacturing goods are non-homothetic.[41] For all agents except middlemen, we represent this non-homotheticity using PIGL preferences (Eckert and Peters, 2023), in which an agent's indirect utility function is given by:[42]

$$V(y, P^A, P^M) = \frac{1}{\eta}\Big(\frac{y}{(P^A)^{\phi}(P^M)^{1-\phi}}\Big)^{\eta} - \nu \ln\Big(\frac{P^A}{P^M}\Big) \tag{12}$$

where y is income, P^A and P^M are prices of agricultural and manufacturing goods, respectively, and $\eta, \phi \in (0, 1)$ and ν are exogenous parameters.[43] This formulation implies that agriculture's share of expenditure (ξ^A) can be written as:

$$\xi^A(y, P^A, P^M) = \phi + \nu\Big(\frac{y}{(P^A)^{\phi}(P^M)^{1-\phi}}\Big)^{-\eta} \tag{13}$$

Because $\eta > 0$, equation (13) implies that agriculture's expenditure share decreases as the agent's income y increases, with an asymptote ϕ. Since the middleman population has measure zero, a middleman's income y is very high and thus his food share converges to ϕ. Therefore, to simplify, we assume middlemen have homothetic preferences with food share equal to ϕ.

Because agents (other than middlemen) have non-homothetic preferences, the aggregate agricultural expenditure in a region n (denoted X^A_n) depends on the distribution of income across agents within that region and can be written as:[44]

[40] This assumption is consistent with the empirical distribution of landholdings observed in Sri Lanka's household survey data (see Appendix C.3).

[41] This assumption is strongly supported by Sri Lanka's household survey data, in which the expenditure share of agricultural goods decreases with household income (see Appendix Figure A10).

[42] Similarly to Eckert and Peters (2023), we choose PIGL over other classes of non-homothetic preferences (e.g. Stone-Geary) due to its convenient aggregation properties.

[43] In our context, P^A and P^M are sectoral price indices that aggregate the prices of multiple goods, as defined below.

[44] See a proof in Appendix G.1. Empirical evidence for equation (14) is provided in Appendix C.5.

$$X_n^A = \phi(1-t)E_n + \nu(1-t)^{1-\eta}\left(\frac{N_n^m(w_n^m)^{1-\eta} + N_n^e(w_n^e)^{1-\eta} + N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}}}{(P_n^A)^{-\eta\phi}(P_n^M)^{-\eta(1-\phi)}}\right) \quad (14)$$

where P_n^A and P_n^M are price indices of agricultural and manufacturing goods (respectively) in region n (defined in Sections 5.3.3 and 5.3.4 below), t is a flat income tax rate (defined in Section 5.4 below), $r_n \equiv R_n/N_n^F$ is the average land rent earned by farmers in region n, and E_n is aggregate income in region n, which can be written as:

$$E_n = \underbrace{w_n N_n^m + w_e N_n^e}_{\text{wage income}} + \underbrace{R_n}_{\text{land rents}} + \underbrace{p_n^f F_n}_{\text{fertilizer sales}} + \underbrace{QR_i^{rent}}_{\text{license rents}} \quad (15)$$

The aggregate expenditure of region n on manufacturing goods (X_n^M) is then simply determined by subtracting aggregate agricultural expenditure from aggregate disposable income:

$$X_n^M = (1-t)E_n - X_n^A \quad (16)$$

5.3.3 Within-Agriculture Consumption

Consumers allocate their agricultural expenditure across multiple crops according to Constant Elasticity of Substitution (CES) preferences with elasticity σ_A. Defining P_n^A as the agricultural price index in region n and β_{nk}^A as the expenditure share of crop k *within total agricultural expenditure*, we can then write:

$$\beta_{nk}^A = \frac{b_k P_{nk}^{1-\sigma_A}}{(P_n^A)^{1-\sigma_A}}, \text{ for } k \in \{1, ..., K\}, \quad \text{and} \quad P_n^A = \left(\sum_k b_k P_{nk}^{1-\sigma_A}\right)^{\frac{1}{1-\sigma_A}} \quad (17)$$

where b_k is an exogenous preference shifter, and P_{nk} is the price index of crop k in region n (defined below). Furthermore, for each crop k, agents consume geographically differentiated varieties from all regions following CES preferences with (Armington) elasticity σ_K. Defining $\beta_{ni,k}^A$ as the share of origin i within the expenditures of region n on crop k, we can then write:

$$\beta_{ni,k}^A = \frac{b_{i,k}(p_{ik}\tau_{ni,k}^A)^{1-\sigma_K}}{P_{nk}^{1-\sigma_K}}, \text{ for } i \in \{1, ..., R\}, \quad \text{and} \quad P_{nk} = \left(\sum_i b_{i,k}(p_{ik}\tau_{ni,k}^A)^{1-\sigma_K}\right)^{\frac{1}{1-\sigma_K}} \quad (18)$$

where $b_{i,k}$ is an exogenous preference shifter, and p_{ik} is the local producer price of the variety of crop k that is produced in origin i.

5.3.4 Manufacturing Consumption

As in agriculture, each agent allocates manufacturing expenses across geographically differentiated manufacturing varieties according to CES preferences with (Armington) elasticity σ_M. Defining β_{ni}^M as the share of origin i *within the total manufacturing expenditures* of destination region n, and P_n^M as the manufacturing price index in region n, we can then write:

$$\beta_{ni}^M = \frac{(p_{ni}^M)^{1-\sigma_M}}{(P_n^M)^{1-\sigma_M}}, \text{ for } i \in \{1, ..., R\}, \quad \text{and} \quad P_n^M = \Big(\sum_i (p_{ni}^M)^{1-\sigma_M}\Big)^{\frac{1}{1-\sigma_M}}$$

where p_{ni}^M is the price in region n of the manufacturing variety from origin i. Substituting pricing equation (9) into these two expressions, we can rewrite manufacturing expenditure shares and price index as follows:

$$\beta_{ni}^M = \left(\frac{\tau_{ni}^M w_i}{T_i^M P_n^M}\right)^{1-\sigma_M}, \text{ for } i \in \{1, ..., R\}, \quad \text{and} \quad P_n^M = \Big(\sum_i (\tau_{ni}^M w_i / T_i^M)^{1-\sigma_M}\Big)^{\frac{1}{1-\sigma_M}} \tag{19}$$

5.4 Taxes and Subsidies

Motivated by the widely documented importance of fertilizer subsidies for Sri Lankan agriculture, we assume that the national government subsidizes fertilizer purchases by farmers.[45] For each crop k, a subsidy rate s_k is applied on the gross market price of fertilizer (p_i^f), so the relevant fertilizer price $p_i^{f,k}$ faced by a farmer who grows crop k in region i is given by:

$$p_i^{f,k} = \begin{cases} (1 - s_k) p_i^f, & \text{if } i \in \{1, ..., I\} \\ p_i^f, & \text{if } i = RoW \end{cases} \tag{20}$$

Subsidy costs are paid by the government, who funds itself through a national income tax: all agents (farmers, workers, and middlemen) in all domestic regions pay an income tax with flat rate t, with the rate chosen such that tax revenues exactly cover governmental expenditures. Therefore, given subsidy rates, regions' aggregate incomes ($\{E_i\}_i$), and agricultural production and producer prices ($\{Q_{ik}, p_{ik}\}_{i,k}$), it can be shown that the necessary tax rate t to fund the government is:

$$t = \Big(\sum_{i=1}^{I} E_i\Big)^{-1} \times \Big(\sum_k \frac{s_k \gamma_k^f}{1 - s_k} \sum_{i=1}^{I} p_{ik} Q_{ik}\Big) \tag{21}$$

[45] On Sri Lanka's fertilizer subsidy policy, see Section A.5 and FAO and WFP (2022, 2023).

5.5 Equilibrium

We close the model by stating market clearing conditions in the markets for fertilizer, crops, and labor. First, fertilizer marketing clearing depends on whether QRs are in place. In the absence of QRs, market clearing can be expressed as a single equation for the whole world:

$$\underbrace{F^G \equiv \sum_{i=1}^{R} F_i}_{\text{global fertilizer supply}} = \sum_{i=1}^{R} \underbrace{\sum_k \frac{\gamma_k^f}{p_i^{f,k}} p_{ik} Q_{ik}}_{\text{fertilizer demand in } i} \tag{22}$$

where $p_i^{f,k}$ is given by equation (20), fertilizer price p_i^f in each region i equals the international fertilizer price p_{RoW}^f, and agricultural production Q_{ik} is given by equation (6). Therefore, with no QRs, there is a single fertilizer price across the whole world. However, with QRs, the international price and the price within Sri Lanka (p_{LKA}^f) may diverge. In that case, fertilizer market clearing consists of two sets of equations, one for Sri Lanka and one for the RoW:

$$\underbrace{\bar{f} + \sum_{i=1}^{I} F_i}_{\text{fertilizer available for use in Sri Lanka}} = \sum_{i=1}^{I} \underbrace{\sum_k \frac{1}{p_{LKA}^f} \frac{\gamma_k^f p_{ik} Q_{ik}}{(1 - s_k)}}_{\text{fertilizer demand in } i} \tag{23}$$

$$\underbrace{F_{RoW} - \bar{f}}_{\text{fertilizer available for use in the RoW}} = \frac{1}{p_{RoW}^f} \underbrace{\sum_k \gamma_k^f p_{RoW,k} Q_{RoW,k}}_{\text{fertilizer expenditure in the RoW}} \tag{24}$$

Second, the market clearing conditions of every agricultural crop produced by every region can be written as the following set of equations:

$$\underbrace{Q_{ik}}_{\text{local production of } k} = \sum_n \frac{\tau_{ni,k}^A}{p_{ni,k}} \underbrace{X_n^A \beta_{nk}^A \beta_{ni,k}^A}_{\text{exports of } k \text{ from } i \text{ to } n}, \forall\, (i,k) \tag{25}$$

where bilateral prices ($p_{ni,k}$) are given by equation (7), aggregate agricultural expenditures X_n^A are given by equation (14), and crop and origin expenditure shares (β_{nk}^A and $\beta_{ni,k}^A$, respectively) are given by equations (17) and (18), respectively.

Third, the labor market clearing condition for mobile workers in each region i is given by:

$$\underbrace{N_i^m}_{\text{mobile labor force}} = \frac{1}{w_i^m} \Bigg(\sum_{k \neq \text{tea}} \underbrace{\gamma_k^n p_{ik} Q_{ik}}_{\text{wage bill of crop-}k\text{ farms}} + \underbrace{\sum_n X_n^M \beta_{ni}^M}_{\text{manufacturing wage bill}} \Bigg) \tag{26}$$

with manufacturing expenditure X_n^M in region n given by equation (16) and origin expenditure shares β_{ni}^{M} given by equation (19). Demand for mobile labor arises from two sources: agriculture (excluding tea) and manufacturing. Intuitively, agricultural demand for labor depends on each crop's total revenue but also on its labor intensity (represented by Cobb-Douglas coefficient γ_k^n). Since manufacturing uses labor as its sole input, its full revenue is paid to labor as wages.

Fourth, the labor market clearing condition for estate workers in each region i is simply:

$$\underbrace{N_i^e}_{\text{estate labor force}} = \frac{1}{w_i^e} \underbrace{\gamma_{tea}^n p_{i,\text{tea}} Q_{i,\text{tea}}}_{\text{wage bill of tea farms}} \tag{27}$$

Given the four market clearing conditions described above, we can now formally define the model's equilibrium as follows:

Definition 1 (Equilibrium) *Given the model's parameters* (R, K, η, ν, ϕ, $\boldsymbol{\sigma}$, $\boldsymbol{\gamma}$) *and exogenous variables* ($\boldsymbol{F}$, $\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{T}$, $\boldsymbol{\tau}$, $\boldsymbol{b}$, $\bar{f}$, $\boldsymbol{s}$) *, an* ***equilibrium*** *is a set of endogenous variables* (t, $\boldsymbol{p^A}$, $\boldsymbol{p^f}$, $\boldsymbol{p^-}$, $\boldsymbol{P}$, $\boldsymbol{w}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$, $\boldsymbol{Q}$, $\boldsymbol{QR}^{rent}$) *satisfying equations 3, 5, 6, 7, 8, 14, 15, 16, 17, 18, 19, 20, 21, 25, 26, 27. Additionally, if QRs are* ***not*** *in place, equation 22 must hold, with* $p_i^f = p_{RoW}^f$ *for all regions* $i \in \{1, ..., R\}$*. Alternatively, if QRs* ***are*** *in place, equations 23 and 24 must hold, with* $p_i^f = p_{LKA}^f$ *for all domestic regions* $i \in \{1, ..., I\}$*.*[46]

5.6 Welfare

Equilibrium variables can also be used to express the welfare of various groups of agents. For mobile workers in region i, whose only source of income is wages, welfare V_i^m can be written as:

$$V_i^m = \frac{1}{\eta}\Big(\frac{(1-t)w_i^m}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big) \tag{28}$$

Note that the welfare of a mobile worker does not depend on whether she works for agriculture or manufacturing, nor on the specific crop within agriculture in which she is employed. This follows from the regional wage being equalized across sectors and crops due to frictionless worker mobility across sectors. Similarly, the welfare V_i^e of estate workers in region i is given by:

$$V_i^e = \frac{1}{\eta}\Big(\frac{(1-t)w_i^e}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big) \tag{29}$$

[46] Variables in bold are collections of region- and/or crop-level terms, as follows: $\boldsymbol{\sigma} = (\sigma_A, \sigma_M, \sigma_K)$, $\boldsymbol{\gamma} = \{\gamma_k^n, \gamma_k^l, \gamma_k^f\}_{k=1}^K$, $\boldsymbol{N} = \{N_i, N_i^m, N_i^e, N_i^F\}_{i=1}^R$, $\boldsymbol{L} = \{\{L_{ik}\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{T} = \{T_i^M, \{T_{ik}^A\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{\tau} = \{\{\tau_{ni}^M, \{\tau_{ni,k}^A\}_{k=1}^K\}_{i=1}^R\}_{n=1}^R$, $\boldsymbol{b} = \{\{b_{nk}, \{b_{ni,k}\}_{i=1}^R\}_{k=1}^K\}_{n=1}^R$, $\boldsymbol{s} = \{s_k\}_{k=1}^K$, $\boldsymbol{p}^A = \{\{p_{ik}\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{p}^f = \{p_i^f, \{p_i^{f,k}\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{p}^- = \{\{h_{ik}, \{p_{ni,k}\}_{n=1}^R\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{P} = \{P_i^A, \{P_{ik}\}_{k=1}^K, P_i^M\}_{i=1}^R$, $\boldsymbol{w} = \{w_i^m, w_i^e\}_{i=1}^R$, $\boldsymbol{R} = \{R_i, \{R_{ik}\}_{k=1}^K\}_{i=1}^R$, $\boldsymbol{E} = \{E_i\}_{i=1}^R$, $\boldsymbol{X} = \{X_i^A, X_i^M\}_{i=1}^R$, $\boldsymbol{\beta} = \{\{\beta_{ik}^A, \{\beta_{in,k}^A\}_{n=1}^R\}_{k=1}^K, \{\beta_{in}^M\}_{n=1}^R\}_{i=1}^R$, $\boldsymbol{Q} = \{\{Q_{ik}\}_{k=1}^K\}_{i=1}^R$, $\mathbf{QR}^{rent} = (QR^{rent}, \{QR_i^{rent}\}_{i=1}^I)$.

The welfare of a farmer h depends on the amount of land rent R^h she earns, which in turn depends on the amount of land she owns. Averaging across all farmers of a region i, we show in Appendix G.2 that average farmer welfare V_i^F can be expressed as:

$$V_i^F = \frac{1}{\eta}\Big(\frac{1}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} e^{\eta[\mu_i+\ln((1-t)R_i/L_i)]+\eta^2\frac{\sigma_L^2}{2}} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big) \tag{30}$$

Finally, we aggregate welfare across the different types of agent using the law of iterated expectations, and write the average welfare V_i^{avg} in region i as:[47]

$$V_i^{avg} = \frac{N_i^m}{N_i}V_i^m + \frac{N_i^e}{N_i}V_i^e + \frac{N_i^F}{N_i}V_i^F \tag{31}$$

6 Estimation

To use our quantitative model to study the impacts of the fertilizer import ban on the Sri Lankan economy, we start by estimating a pre-ban baseline of the model in 2019. In this section, we describe the steps for such estimation: first parameter estimation, then model "inversion" to back out implied unobserved variables.

6.1 Step 1: Estimating Parameters

6.1.1 Preference Parameters

This group of parameters consist of PIGL non-homotheticity parameters (ϕ, η, ν) and elasticities of substitution (σ_A, σ_M, σ_K). PIGL parameter ϕ represents the asymptotic expenditure share of agriculture as income grows large (while keeping price indices P_n^A and P_n^M fixed), as implied by equation (13). Therefore, we simply set ϕ to 1.05%, the lowest reported food expenditure share in our 2019 household data. Parameter ν is set to 0.12, following Eckert and Peters (2023).

To estimate the Engel elasticity η, we take logs of equation (13), add household subscripts, collect relevant terms into a region fixed effect, and add a household size control variable and an error term, thus obtaining the following equation, which is estimated on 2019 data from Sri Lanka's Household Income and Expenditure Survey (HIES):

$$\ln(\xi_{nh}^A - \phi) = -\eta \ln(y_{nh}) + \theta hhsize_{nh} + \omega_n + \epsilon_{nh} \tag{32}$$

where h indexes households, n indexes regions, ξ_{nh}^A is the share of income the household spends on food, y_{nh} is household income, $hhsize_{nh}$ is household size, ω_n is a region fixed effect, and ϵ_{nh} is

[47] The average welfare expression does not include terms relating to the middlemen due to their measure zero.

an error term.[48] To allay concerns about the potential endogeneity of household income (e.g. due to seasonality in prices and employment opportunities), we follow an IV approach. We instrument household income y_{nh} with two arguably exogenous sources of income: income obtained from lottery (and other *ad hoc* gains) and income obtained from disasters and other relief payments.[49] To our knowledge our paper is the first to estimate the Engel elasticity off exogenous income variation generated by a lottery.[50] The resulting IV estimate for η, shown in Table 2, is 0.656.[51]

Table 2: Preference Parameters

Parameter	Value	Standard Error	Source	Description
ϕ	0.0105	.	From HIES data	Asymptotic agricultural share
ν	0.12	.	Eckert and Peters (2023)	PIGL parameter
η	0.656	0.129	Estimated on HIES data	Engel elasticity
σ_M	2.528	.	Feenstra et al (2018)	EoS across origins (manufacturing)
σ_A	1.714	0.153	Estimated on HIES data	EoS across crops
σ_K	2.377	.	Match crop price increase	EoS across origins (within crop)

Notes: for each preference parameter in the model, the table displays the value we impose for it, the standard error of its estimate (when applicable), the source from which it is taken or estimated, and a brief description of the parameter.

To estimate the elasticity of substitution across crops (σ_A), we take logs of equation (17), add household subscripts, collect relevant terms into crop- and household-specific fixed effects, and add an error term, obtaining the following equation, which is estimated on 2019 HIES data:

$$\ln(\beta^A_{nhk}) = -(\sigma_A - 1)\ln(P_{nhk}) + \omega_k + \omega_h + \epsilon^A_{nhk} \tag{33}$$

where k indexes crops, β^A_{nhk} is the share of crop k within the household's total food expenditures, P_{nhk} is crop price, ω_k and ω_h are crop and household fixed effects (respectively), and ϵ^A_{nhk} is an error term. To address concerns about the potential endogeneity of crop prices (e.g. due to unobserved region-specific tastes that affect both prices and expenditure shares), we follow the IV approach in Sotelo (2020), instrumenting price variable $\ln(P_{nhk})$ with a dummy variable D^{GAEZ}_{nk}, which takes value one if the potential yield of crop k in region n is nonzero, according to the FAO-GAEZ data (see Appendix A.6). Thus, our identification of the cross-crop elasticity of substitution σ_A comes from variation in crop prices due to regional shifters of agricultural supply driven by climactic and geological endowments, and thus arguably orthogonal to unobserved factors that may potentially affect consumer behavior. Our estimate, $\hat{\sigma}_A = 1.714$, is lower than in Sotelo (2020) (2.39) but

[48] Here a region is one of Sri Lanka's 25 districts since DS are not recorded in the household data.
[49] Both exogenous sources of income refer to the previous 12 months.
[50] Other papers (Boppart, 2014; Comin et al., 2021; Eckert and Peters, 2023; Fan et al., 2023) have estimated the parameters that govern non-homothetic preferences using household data but not using such IV approach.
[51] See Appendix D.1 for details on the 2SLS estimation results of (η, σ_A), including first stage results.

higher than in much of the literature.[52]

We calibrate the value of the Armington elasticity of substitution across different origins of the same crop to $\hat{\sigma}_K = 2.377$ so that the average increase in agricultural producer prices generated by the model-driven counterfactual exercise of Section 7.1 matches the corresponding increase of 22.65% that can be observed in the data between January 2021 and January 2022 (see Appendix Figure A11).[53] The Armington elasticity of substitution within manufacturing is set to $\sigma_M = 2.528$, which is consistent with the lower range of estimates from Feenstra et al. (2018).[54]

6.1.2 Trade Costs

Agricultural and manufacturing firms can trade their goods (both domestically and internationally) by incurring "iceberg" trade costs ($\{\tau^M_{ni}\}_{n,i}$, $\{\tau^A_{ni,k}\}_{n,i,k}$): they represent how many units of a good must be shipped from origin i so that one unit will arrive at destination n. To estimate these trade cost matrices, we start by assuming the following parametrizations:

$$\tau^M_{ni} = \exp(\alpha \times distkm_{ni}) \tag{34}$$

$$\tau^A_{ni,k} = \exp(\alpha \times distkm_{ni}) \times (B_k)^{\mathbb{1}\{i=RoW \text{ or } n=RoW\} \times \mathbb{1}\{i \neq n\}} \tag{35}$$

where $\alpha > 0$ is a parameter governing the semi-elasticity of trade costs with respect to distance, $distkm_{ni}$ is the (geodesic) straight-line distance in kilometers between the geographic centroids of regions n and i, and $B_k \geq 1$ is an extra (multiplicative) cost of trading agricultural crop k internationally.[55]

To estimate semi-elasticity α, we use the fact that the cost of transporting tea for 100 kilometers within Sri Lanka was $p_{T,100} = 1,496$ LKR per metric ton (LKR/MT) in 2019, while the farm-gate price of tea was $p_{tea} = 455,473$ LKR/MT. The corresponding iceberg trade cost can then be written as $\tau^A_{ni,tea} = 1 + \frac{p_{T,100}}{p_{tea}}$, which, using equation (35), implies a semi-elasticity of $\hat{\alpha} = 3.3 \times 10^{-5}$.

For each crop k, we estimate international trade cost B_k as the gap between the crop's interna-

[52] For example, the estimated elasticity in Behrman and Deolalikar (1989) is 1.2.

[53] Intuitively, a higher value of σ_K blunts the ban's effect on crop prices in the model because it implies that each crop's international variety is a good substitute for its Sri Lankan varieties, thus keeping domestic prices in check through international competition.

[54] See row "Apparel Manufacturing" in Table 3 of Feenstra et al. (2018). We chose this particular value because apparel manufacturing is an important industry within Sri Lankan manufacturing.

[55] When district l in Sri Lanka trades internationally, we assume the relevant distance is the distance between l and the Colombo district, where Sri Lanka's main port is located. Formally, we set $dist_{l,RoW} = dist_{l,Colombo}$ and $dist_{RoW,l} = dist_{Colombo,l}$. Cost B_k can be interpreted as a "border-crossing" cost (e.g. a tariff) that is paid to trade crop k internationally *in addition to* the transportation cost. We assume the manufacturing sector faces no such additional cost. Unlike international trade costs B_k, the distance semi-elasticity α is the same for all crops (and for manufacturing). This assumption is necessary because data on transportation costs is only available for one crop (tea).

tional price and its domestic price at the port of Colombo:

$$\hat{B}_k = \max\left\{\frac{p_{RoW,k}}{p_{Colombo,k}}, \frac{p_{Colombo,k}}{p_{RoW,k}}\right\} \tag{36}$$

where $p_{RoW,k}$ is the 2019 price of crop k in the RoW, and $p_{Colombo,k}$ is its 2019 "benchmark price" at the port of Colombo. This estimation approach selects between the ratio of international price to domestic price and its reciprocal the one that is larger than unity.[56] Thus, we implicitly assume that an international price that is larger (*smaller*) than its domestic price reveals a crop's status as an export (*import*) crop. The magnitude of such price gap reveals the severity of the barriers faced by the crop's exporters (*importers*). Our estimated $\hat{B}_k$ are displayed in Table 3.

6.1.3 Production Function Parameters

Since our paper focuses on agriculture, it is important to find context-appropriate values for the parameters $\{\gamma_k^f, \gamma_k^l, \gamma_k^n\}_k$ that govern Cobb-Douglas agricultural production function (1). Leveraging 2022 NFS data on fertilizer intensity, we estimate the fertilizer coefficient $\{\gamma_k^f\}_k$ of each crop k as:

$$\gamma_k^f = \frac{M_k^{f,LKA}}{p_k^{LKA} Q_k^{LKA}}, \quad k \in \{1, ..., K\} \tag{37}$$

where $M_k^{f,LKA}$ is annual fertilizer expenditure by Sri Lankan farmers for production of crop k (in USD), p_k^{LKA} is the price of crop k in Sri Lanka (in USD/kg), and Q_k^{LKA} is Sri Lanka's annual production of crop k (in kg).[57] The resulting estimates are displayed on Table 3. The remaining Cobb-Douglas coefficients (γ_k^l, γ_k^n) must add to $(1-\gamma_k^f)$ for each crop k. We apportion this residual value to land and labor using data from Sri Lanka's IO matrix on the labor share of value added:

$$\gamma_k^n = \frac{CEMP_{s(k)}}{VA_{s(k)}} \times (1 - \gamma_k^f) \tag{38}$$

where $CEMP_s$ is total employee compensation in sector s, VA_s is total value added (at factor cost prices) in sector s, and $s(k)$ is the sector to which crop k belongs. The land coefficient γ_k^l is then computed as: $\gamma_k^l = 1 - \gamma_k^f - \gamma_k^n$.

6.1.4 Land Distribution Parameters

The model's parameters that govern land distributions ($\{\mu_i, \sigma_{Li}^2\}_i$) are estimated by leveraging the properties of the log-normal. We use 2016 HIES data to estimate $\hat{\sigma}_{Li}$ as the (weighted) standard

[56] As ratios $\frac{p_{RoW,k}}{p_{Colombo,k}}$ and $\frac{p_{Colombo,k}}{p_{RoW,k}}$ are each other's reciprocals, only one of them can be larger than unity.

[57] Fertilizer use $M_k^{f,LKA}$ is computed at post-subsidy fertilizer prices and includes the three main types of fertilizer used: urea, muriate of potash (MOP), and triple superphosphate (TSP).

Table 3: Other Parameters

	Cobb-Douglas Coefficients			International trade costs
Crop	Fertilizer (γ_k^f)	Labor (γ_k^n)	Land (γ_k^l)	B_k
Cinnamon	0.032	0.429	0.539	2.36
Cloves	0.043	0.424	0.533	2.36
Groundnuts	0.066	0.414	0.520	2.36
Maize	0.177	0.365	0.458	1.07
Onions	0.075	0.410	0.515	2.14
Potatoes	0.129	0.394	0.477	2.39
Rice	0.110	0.469	0.421	1.43
Tea	0.109	0.487	0.404	9.40

Notes: the table displays our estimates of the model parameters that govern agricultural production technology ($\{\gamma_k^f, \gamma_k^n, \gamma_k^l\}_{k=1}^K$) and agricultural international trade costs ($\{B_k\}_{k=1}^K$). See Sections 6.1.2 and 6.1.3 for descriptions of our estimation methodologies.

deviation of the logarithm of landholdings across farmers in district i, and $\hat{\mu}_i$ as: $\hat{\mu}_i = \ln(L_i/N_i^F) - \hat{\sigma}_{Li}^2/2$, where L_i is total arable land (in ha) and N_i^F the number of farmers in district i.[58]

6.2 Step 2: Inversion to Recover Unobservables

Equipped with estimates for the model's parameters, we proceed to recover the remaining unobserved variables needed to fully characterize the 2019 baseline economy. We use the process commonly known as model inversion, which can be informally described as plugging observed variables into the model's equilibrium system of equations and then solving the system to back out unobserved variables.[59] More formally:

Definition 2 (Inversion) *Given the model's parameters, observed exogenous variables ($\boldsymbol{N}^m$, $\boldsymbol{N}^F$, $\boldsymbol{L}$, $\boldsymbol{\tau}$, $\boldsymbol{s}$), and observed endogenous variables ($\boldsymbol{w}$, $\boldsymbol{Q}$, $\mathbf{p}^{\boldsymbol{A}}$, $\mathbf{p}^{\boldsymbol{f}}$), an* ***inversion*** *is a set of unobserved exogenous variables ($\boldsymbol{T}$, $\boldsymbol{b}$, $\boldsymbol{N}^e$, $\boldsymbol{F}$) and unobserved endogenous variables (t, $\mathbf{p}^-$, $\boldsymbol{P}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$) such that the resulting set of endogenous variables is an equilibrium of the economy with the resulting set of exogenous variables and with no QRs.*

Intuitively, the inversion process picks values for unobserved variables to rationalize the data. For example, consider equation (6), which expresses equilibrium agricultural production (Q_{ik}) as a function of productivity (T_{ik}^A), prices, and land mass. Suppose that, at a given moment, the observed Q_{ik} is "too high" in the sense that the equation's left-hand side is larger than its right-hand side. The inversion process will then increase productivity T_{ik}^A enough to balance the equation.

[58] Landholdings are defined as the total amount of land owned by a household (excluding land with housing built on top) and are measured in ha. The standard deviation is computed using only households with positive landholdings. We use 2016 HIES data because landholding data was not available in the more recent 2019 HIES. See Appendix C.3 for more details on the estimation of the land distribution parameters.

[59] Alternatively, inversion consists of recovering unobserved from observed variables using the model's structure.

Appendix E provides details on the inversion procedure. We use data from 2019, which we choose as our baseline year because it precedes the 2021 fertilizer ban and other unusual events, such as the Covid-19 pandemic, and also due to household data availability.[60]

7 Counterfactuals

In this section, we use the estimated model to investigate the GE effects of Sri Lanka's fertilizer import ban (and of other potential alternative policies). Through counterfactual exercises, we demonstrate how government policies like import restrictions limit fertilizer access and how fertilizer subsidies in turn mediate the distributional impacts of such limitations. While the sudden and unexpected ban aids in the identification of structural parameters, less extreme forms of fertilizer shortage help quantify the value of fertilizer in more realistic situations such as the global fertilizer shortages driven by Russia's invasion of Ukraine. The counterfactuals also highlight the importance of realistic model features such as non-homothetic preferences in shaping welfare outcomes.

In each counterfactual, we start from our estimated baseline economy and then compute a counterfactual equilibrium in which the government imposes QRs to imports (represented by a maximum allowed amount $\bar{f}$ of imported fertilizer) while all other exogenous variables ($\boldsymbol{F}$, $\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{T}$, $\boldsymbol{\tau}$, $\boldsymbol{b}$, $\boldsymbol{s}$) and parameters (η, ν, ϕ, $\boldsymbol{\sigma}$, $\boldsymbol{\gamma}$) are kept at their original levels. We then compare the values of the endogenous variables in the counterfactual equilibrium to their original values, with the difference being interpreted as the causal effect of the QR in general equilibrium. Since the bulk of the fertilizer used by Sri Lanka is initially imported, the effects of a full ban can also be interpreted as the economy-wide cost of losing access to this critical intermediate input.

7.1 Setting Counterfactual Parameters

Three practical considerations arise when implementing a counterfactual. First, one price must be normalized, as the equilibrium system is homogeneous of degree zero in the set of prices. We set the average wage in the RoW (w_{RoW}) to its actual value in the pre-ban period.[61]

Second, we must choose a specific value for parameter $\bar{f}$, the amount of fertilizer imports allowed by the government under the QR regime. In our main counterfactual, we focus on a full ban of fertilizer imports, setting $\bar{f}$ to zero, which is consistent with both Sri Lanka's official ban policy and our reduced-form estimated ban effects in Section 4.1.

[60] We include in the analysis the eight crops ($K = 8$) for which there is data on fertilizer intensity and district-level producer prices, namely: cinnamon, cloves, groundnut, maize, onions, potatoes, rice, and tea.

[61] We use the world's income per capita in 2019 of 11,330.5 USD, equivalent to approximately 4.08 million LKR (using the 2022 exchange rate). For convenience of exposition, all results are displayed in USD.

Third, we need a welfare metric to express the ban's welfare impacts. While it is possible to compute changes in utility between the baseline and the counterfactual equilibrium using equation (12), it is difficult to interpret the meaning of such metrics.[62] We instead use Equivalent Variation (EV) metrics, defined as the hypothetical change in income that would replicate the welfare effects of the ban under the set of prices of the baseline equilibrium. For each district, we compute welfare effects for multiple types of agents: mobile workers, estate workers, three types of farmer (median, average, and representative), and the representative agent.[63] Appendix F provides detailed definitions of EVs and agent types.

7.2 Baseline Results: The Effect of the Fertilizer Import Ban

Table 4 presents the model-generated impacts of the fertilizer import ban on outcomes measured at the national level. The primary effect of the ban is the full elimination of fertilizer imports, with a resulting increase of 6.41 USD/kg (or 2,068%) in gross fertilizer prices in Sri Lanka between the baseline and the counterfactual equilibrium. This leads to a massive decrease of 4 billion kg (or 96.5%) in fertilizer use in Sri Lankan agriculture. The ensuing decline in the government's expenses with fertilizer subsidies allows the tax rate to be cut by 2.4 percentage points (or 27%).

Lack of fertilizer causes declines in Sri Lankan crop yields, which ultimately result in a reduction in downstream agricultural exports to the tune of 596 million USD. From a trade balance perspective, these export losses represent a 48% offset of the 1.25 billion USD reduction in fertilizer imports, indicating that the ban's implicit goal of saving foreign exchange may have been severely limited by the large decline in crop exports. This finding would have been missed by a naive inspection of fertilizer import values that failed to account for such downstream effects. Furthermore, as will be shown in Section 7.5, this offset factor would be even higher, and possibly even higher than 100%, in the absence of Sri Lanka's large fertilizer subsidies.[64]

To better understand the ban's effects on agriculture, Table 5 shows crop-by-crop results. As a result of the ban, physical yields fell by 12 to 54%, with considerable heterogeneity across crops. Similarly, crop exports fell across the board but with substantial variation across crops. On one extreme, the most fertilizer-intensive crop, maize, saw yields fall 53.7% and exports 36.1%. Less

[62] The interpretation of utility as an "aggregated quantity", which is often used for CES preferences, is not valid for our non-homothetic utility function. Moreover, it is not even guaranteed that our indirect utility is positive in either equilibrium. Therefore, a decrease in utility could, in principle, be larger than 100%, among other issues.

[63] Workers of a given district are assumed to be identical, so the welfare effect of the ban is the same for all workers. The median (average) farmer of each district is defined as a farmer whose landholding size is equal to the median (average) in that district. The representative farmer of a district is the farmer whose utility is equal to V_i^F, the district's average farmer utility. Similarly, the representative agent is defined as a hypothetical agent whose income provides him with a level of utility equal to V_i^{avg}, the district's average utility.

[64] The intuition for this finding is that subsidies drive farmers to above-optimal fertilizer use. In the absence of subsidies, the cost and marginal revenue product of fertilizer are more closely aligned, so a marginal decrease in fertilizer use leads to reductions of roughly similar magnitude in fertilizer imports and agricultural exports.

Table 4: Ban Effects (Country-Level)

Variable	% Effect	Effect	Description
f_{LKA}	-96.5%	-4,043 MM kg	Fertilizer use (kg)
$M^f_{q,LKA}$	-100%	-4,043 MM kg	Fertilizer imports (kg)
M^f_{LKA}	-100%	-1,253 MM USD	Fertilizer imports (USD)
p^f_{LKA}	+2,068%	+6.41	Fertilizer price pre-subsidy (USD/kg)
$p^{f,sub}_{LKA}$	+2,072%	+1.36	(Average) fertilizer price post-subsidy (USD/kg)
EXP^A_{LKA}	-22.6%	-596 MM USD	Agricultural exports (USD)
$NEXP^A_{LKA}$	-27.9%	-601 MM USD	Net agricultural exports (USD)
t	-27.3%	-2.4 p.p.	Tax rate (percentage points)

Notes: The table shows absolute and relative (%) changes in several variables between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. Fertilizer use (f_{LKA}), in MM kg, is the total fertilizer demand by Sri Lanka's agricultural sector. Imported fertilizer quantity ($M^f_{q,LKA}$), in MM kg, is the total quantity of fertilizer that Sri Lanka buys from the RoW, while imported fertilizer value (M^f_{LKA}), in MM USD, is the total revenue that foreign sellers obtain from such sales. The fertilizer price in Sri Lanka (p^f_{LKA}), in USD/kg, is the price earned by fertilizer sellers. The (average) post-subsidy fertilizer price ($p^{f,sub}_{LKA}$) at which fertilizer can be bought by Sri Lankan farmers is the ratio of farmers' total fertilizer expenditures to their total fertilizer use. Gross agricultural exports (EXP^A_{LKA}), in MM USD, are the total agricultural sales of Sri Lanka to the RoW, while net agricultural exports ($NEXP^A_{LKA}$) are gross agricultural exports minus gross agricultural imports. The tax rate (t) is a flat income tax levied by Sri Lanka's government on all domestic agents.

Table 5: Ban Effects (Crop-Level)

Crop (k)	Fertilizer intensity (γ^f_k)	Production		Gross Exports	
		in MM kg	in %	in MM USD	in %
Cinnamon	0.032	-2.8	-12.2%	-15.9	-7.3%
Clove	0.043	-1.1	-16.2%	-3.7	-9.8%
Groundnuts	0.066	-8.0	-21.9%	-4.7	-13.5%
Maize	0.177	-168.3	-53.7%	-30.8	-36.1%
Onions	0.075	-9.2	-26.8%	-4.6	-16.8%
Potatoes	0.129	-15.0	-42.9%	-7.2	-27.9%
Rice	0.110	-1,940	-37.9%	-505	-24.2%
Tea	0.109	-84.0	-30.1%	-24.4	-19.0%

Notes: For each crop, the table shows its fertilizer intensity and the change in Sri Lanka's crop production and gross exports between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. The fertilizer intensity of each crop k is defined as the fertilizer coefficient γ^f_k in its Cobb-Douglas production function. Production of crop k, in MM kg, is computed by adding district-level production across all Sri Lankan districts ($\sum_{i=1}^I Q_{ik}$). Gross exports of crop k, in MM USD, are the total crop-k sales of Sri Lanka to the RoW, which is computed as: $\sum_{i=1}^I X^A_{RoW}\beta^A_{RoW,k}\beta^A_{RoW,i,k}$.

Figure 5: Ban Effects on Exports, by crop

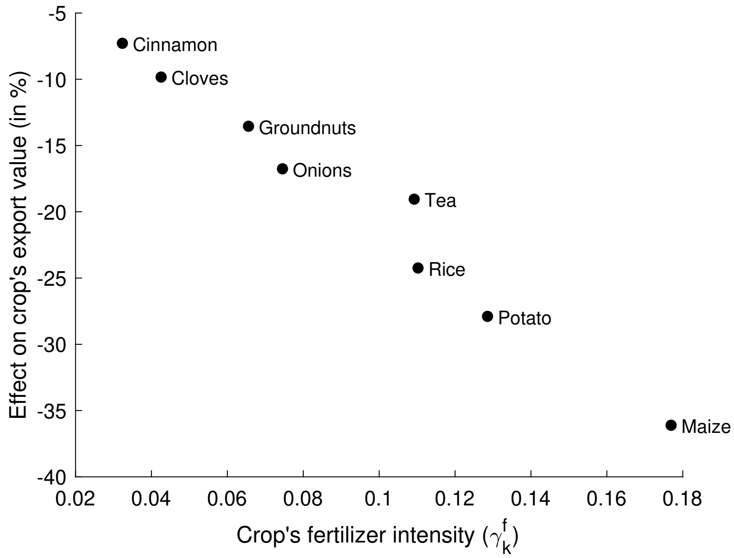

Notes: The figure plots the fertilizer intensity of each crop against the % change in the value of the crop's exports from Sri Lanka to the RoW between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. The fertilizer intensity of crop k is defined as the fertilizer coefficient γ_k^f in its Cobb-Douglas production function. The export value of crop k is the total crop-k sales of Sri Lanka to the RoW, which is computed as: $\sum_{i=1}^{I} X^A_{RoW}\beta^A_{RoW,k}\beta^A_{RoW,i,k}$. The correlation coefficient between fertilizer intensity and % change in export value is -0.986.

fertilizer-intensive crops like cinnamon saw decreases of only 12.2% in yields and 7.3% in exports. Rice yields decline by 37.9%, which is comparable to the 31.6% drop observed in the 2022 data relative to the historical average (see Section 4.2). Further illustration of the ban's effects on agriculture is provided by Figure 5, which plots each crop's fertilizer intensity coefficient γ_k^f against the ban's effect on its export value. There is a clear strong negative correlation of -0.986 between the two variables, with the three most fertilizer-intensive crops (rice, potatoes, and maize) suffering the largest percentage losses in exports by far.

A crucial benefit of using a quantitative GE model is the ability to provide estimates of the ban's welfare effects. Column (1) of Table 6 displays average effects on the welfare of mobile workers, estate workers, farmers, and the representative agent.[65] The average welfare loss for the representative agent is equivalent to a 4.35% decrease in income. However, there is much heterogeneity around that average: while welfare effects are generally negative, the sign and magnitude of welfare changes vary considerably by agent type and geography. Farmers and estate workers are harmed disproportionately because they are attached to the agricultural sector, in contrast to mobile workers, who have the option of switching employment to the manufacturing sector even if they initially work in agriculture.[66] This leads to a composition effect in which the welfare of the representative agent decreases more in regions with a relatively high fraction of farmers, as shown

[65] Appendix Table A8 displays the full set of welfare changes for all agent types and districts.

[66] In some districts, mobile workers even have a gain in welfare (see Appendix Table A8) because their loss from lower agricultural productivity is so small that it is more than offset by their gain from reduced tax rates (see Table 4).

Table 6: Ban Effects on Welfare, by agent type

	Equivalent Variation (EV)					
	Main CF	Homothetic preferences	No subsidies	No subsidies + Homothetic preferences	Halving (QR)	Halving (Iceberg)
Agent type	(1)	(2)	(3)	(4)	(5)	(6)
Mobile Worker	-0.67%	-0.46%	-1.52%	-0.82%	+0.13%	-0.25%
Estate Worker	-16.58%	-17.71%	-9.84%	-9.96%	-3.19%	-3.54%
Representative Farmer	-19.51%	-21.01%	-11.43%	-11.78%	-3.9%	-4.24%
Representative Agent	-4.35%	-2.79%	-3.37%	-1.98%	-0.64%	-1.01%

Notes: for each agent type, Columns (1)-(4) show the average % welfare change between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. The (unweighted) average is computed across all Sri Lankan districts. The representative farmer of a district i is defined as a farmer whose utility is equal to the district's average farmer utility (V_i^F). Similarly, the representative agent of a district i is defined as an hypothetical agent whose (net-of-tax) income is just enough to provide her with the district's average utility (V_i^{avg}). Welfare changes are defined as EVs: starting from the agent's baseline (net-of-tax) income, the EV is the % change in (net-of-tax) income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels. Column (1) refers to our main counterfactual exercise (see Section 7.1), while Columns (2), (3), and (4) refer to alternative counterfactual exercises in which we assume (respectively) homothetic preferences, no fertilizer subsidies, and both homotheticity and no-subsidies (see Section 7.5). Columns (5) and (6) show average % welfare changes in two counterfactual exercises in which the quantity of fertilizer imported into Sri Lanka is reduced by 50% either via QRs (Column 5) or via increases in iceberg trade costs (Column 6).

in Figure 6a. For example, in districts where farmers make up less than 20% of the population, the representative agent obtains a (small) welfare gain. In contrast, where farmers make up more than half of the population, the damages are equivalent to income losses in the range of 4%-13%.

Spatial heterogeneity in welfare impacts is further illustrated by Figure 7. For farmers (Figure 7a), this heterogeneity is explained by geographic differences in the relative importance of specific agricultural crops. Figures 6b and 6c show that welfare losses for farmers and the representative agent were worse in districts whose agricultural sector specializes in relatively fertilizer-intensive crops.[67] For example, in districts with a 6% fertilizer cost share, the harm to the representative farmer was roughly equivalent to a 10% income loss, whereas the corresponding harm in districts with a 11%-12% fertilizer cost share was as high as 23%.

For mobile workers, spatial heterogeneity in welfare losses (Figure 7b) is explained by a different type of geographic variation, namely in the size of the agricultural sector.[68] Figure 6d shows that the welfare of mobile workers decreased by more in districts with a high agricultural share of employment. Thus, although mobile workers suffered less than farmers everywhere, those living in heavily agricultural areas still faced severe losses equivalent to as much as a 12% income reduction, as those areas did not have a vibrant manufacturing sector that could serve as a buffer and absorb the excess labor exiting from the adversely affected agricultural sector.

One potential concern about our counterfactual exercise is that it assumes perfect mobility of

[67] To measure district-level fertilizer intensity, we compute fertilizer's cost share across the whole agricultural sector of each district i at baseline as: $(\sum_k \gamma_k^f p_{ik} Q_{ik})/(\sum_k p_{ik} Q_{ik})$.

[68] The fact that spatial heterogeneity in farmer and worker welfare effects are explained by different factors is also suggested by the substantial differences between Figures 7a and 7b in the spatial pattern of welfare changes.

Figure 6: Ban Effects on Welfare, by district

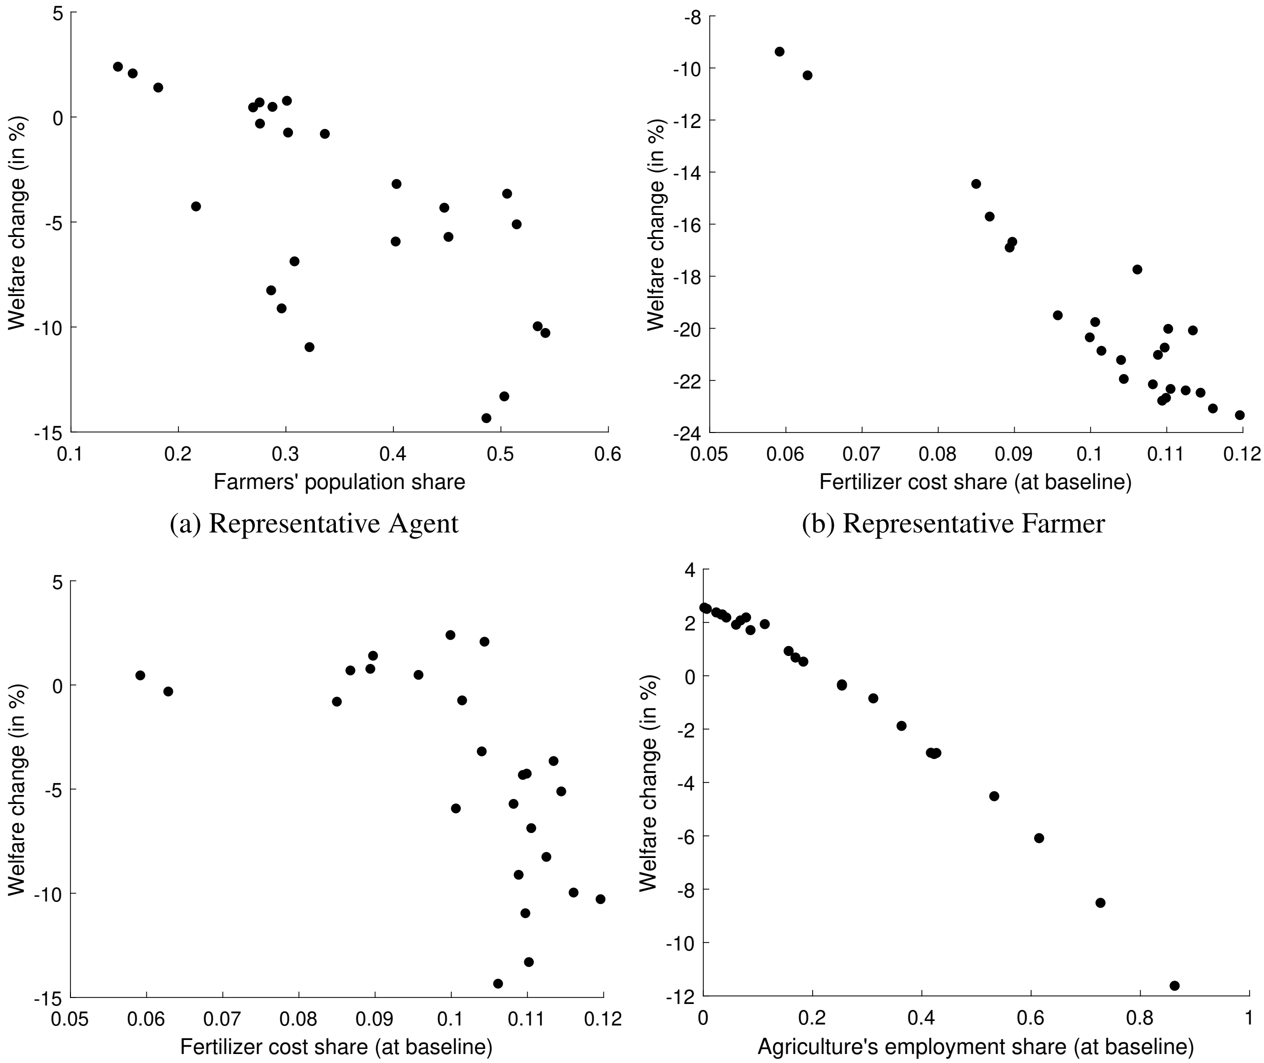

(a) Representative Agent

(b) Representative Farmer

(c) Representative Agent

(d) Mobile Worker

Notes: The figure displays scatter plots in which eat dot represents one of Sri Lanka's districts, the vertical axes show the % change in the welfare of (various types of) agents between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs, and the horizontal axes show various baseline variables that help explain the welfare changes. Panels 6a and 6c show welfare changes for the representative agent, while panels 6b and 6d show welfare changes for the representative farmer and the mobile worker (respectively). Welfare % changes are always defined as an EV: starting from the agent's baseline (net-of-tax) income, the EV is the % change in (net-of-tax) income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels. The (district-level) baseline variables in the horizontal axes are the following: in panel 6a, the farmer population share, defined as the number of farmers in the district divided by its total population; in panels 6b and 6c, the fertilizer cost share, defined as farmers' total fertilizer expenditures divided by total agricultural revenue; in panel 6d, the agricultural employment share, defined as the fraction of mobile workers who work in agriculture. The correlation coefficient between the two variables in each panel are -0.670 (panel 6a), -0.951 (panel 6b), -0.611 (p 6c), and -0.991 (panel 6d).

Figure 7: Spatial Heterogeneity of Ban Effects on Welfare

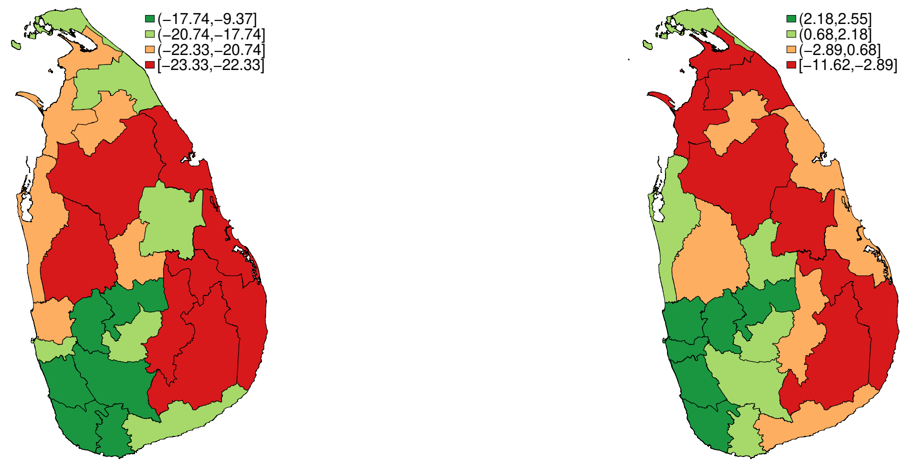

(a) Representative Farmer (b) Mobile Worker

Notes: The figures show the % change in the welfare of each district's representative farmer (panel 7a) and mobile worker (panel 7b) between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. Welfare % change is defined as an EV: starting from an agent's baseline (net-of-tax) income, the EV is the % change in (net-of-tax) income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels.

workers between manufacturing and agriculture (and across agricultural crops). In Appendix H, we extend the model to account for potential frictions to intersectoral labor reallocation. This extended model predicts similar welfare effects of the ban regardless of the severity of frictions, suggesting that frictionless intersectoral reallocation is not a first-order mechanism driving our main results.

7.3 The Effects of Smaller Shocks

Our estimates show Sri Lanka's fertilizer import ban produced very negative economic consequences. Even if countries do not implement policies as drastic as a full ban, many impose QRs on imports. To investigate the consequences of such smaller policy "shocks", we recompute our counterfactual exercise multiple times, each time imposing a different level of QRs to fertilizer imports ($\bar{f}$). The results are displayed in Figure 8.

Unsurprisingly, Figure 8a shows that the magnitude of the average welfare loss of the representative agent increases with QRs strictness (measured by the imposed % reduction in fertilizer import quantities). Moreover, welfare losses increase at increasing rates due to the presence of decreasing marginal returns to fertilizer use in agricultural production. Figure 8b repeats the analysis but plotting changes in domestic fertilizer prices (rather than in imported quantities) on the horizontal axis. As prices rise, welfare losses increase at decreasing rates.[69] The contrast between

[69]There are two reasons for the decreasing marginal effects of fertilizer prices. First, the agricultural profit function is convex in the prices of all inputs, including fertilizer (intuitively, an initial fertilizer price increase leads farms to use less fertilizer so, when prices increase again later, this new price increase is applied on a smaller "base" quantity).

Figure 8: Relationship Between Shock Size and Welfare Effects

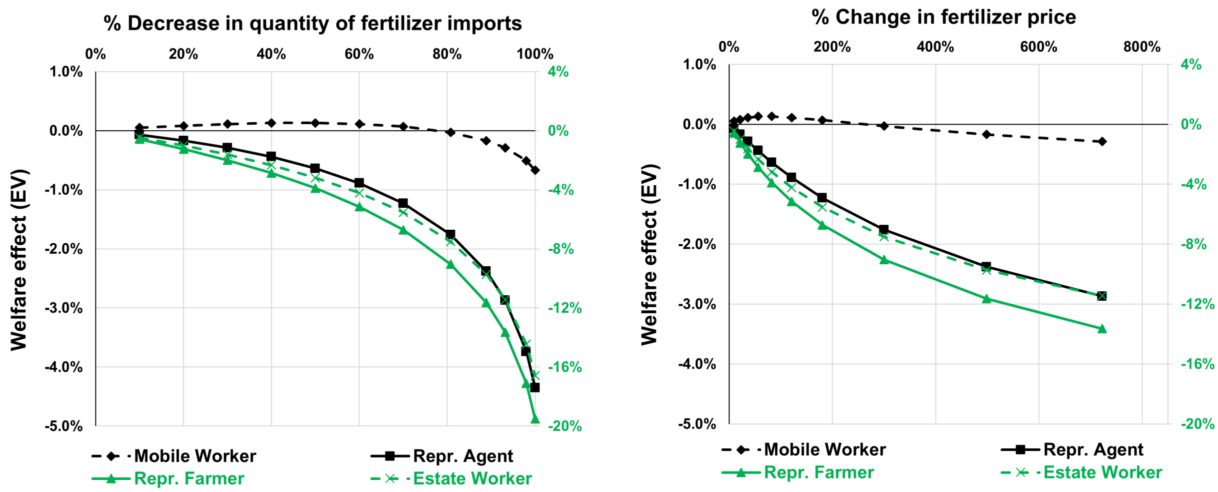

(a) Welfare and Fertilizer Quantities

(b) Welfare and Fertilizer Prices

Notes: The figures show the changes in welfare, imported fertilizer quantities, and fertilizer prices (gross of subsidies) in a series of alternative model-generated counterfactual exercises, each of which imposes a different level of QRs against the importation of fertilizer into Sri Lanka. For each counterfactual exercise, the horizontal axis in Panel 8a shows the % decline in the quantity of imported fertilizer caused by the QRs, while the horizontal axis in Panel 8b shows the resulting % increase in fertilizer prices (gross of subsidies). In both panels, the vertical axis shows welfare changes for four types of agents (mobile worker, estate worker, representative farmer, and representative agent), as measured by (the cross-district average of) the EV faced by that agent type as a result of the QRs. The EV is defined as the percent change in (net-of-tax) income that would lead to the same welfare loss as the QRs if all prices were kept constant at their baseline levels.

the increasing marginal effects of Figure 8a and the decreasing marginal effects of Figure 8b is explained by the fact that the % change in fertilizer prices is not only increasing but also convex in QR strictness.[70]

7.4 Comparison with Literature

We compare our findings with the existing fertilizer literature in Table 7. Panel A displays results from a branch of the literature that focuses on small-scale interventions, usually (but not always) consisting of randomized control trials (RCTs) that provide agricultural inputs to a treated group of farmers or regions and compare their outcomes to those of a control group. The first row reproduces our main findings: the ban led to a 96.5% decrease in fertilizer use in Sri Lanka (Table 4), with a resulting 38% decrease in rice yields (Table 5). Since the papers in this literature deal with increases (rather than decreases) in fertilizer use, the second row “mirrors” our main findings, recasting them

Second, a given fertilizer price hike makes agriculture less profitable relative to manufacturing, leading the latter sector to expand at the expense of the former. So, when a second price hike comes, it hits an agricultural sector that is now smaller than it was before the first hike.

[70]This convexity can be observed visually by noticing that, while the horizontal distance between pairs of consecutive points is approximately constant in Figure 8a, it increases as we move from the left to the right of Figure 8b. Intuitively, a marginal tightening of QRs raises fertilizer price by more when these QRs were already tight to begin with.

Table 7: Comparison with Fertilizer Literature

	Panel A: Small-Scale Interventions		
	Δ Fertilizer Quantity	Δ Yield	Ratio: (2) to (1)
Paper	(1)	(2)	(3)
This paper	-96.5%	-38%	0.39
This paper (mirrored)	+2,757%	+61%	0.022
This paper (partial equilibrium)	.	.	0.10
Beaman et al. (2013)	+150%	+30%	0.20
Bold et al. (2017)	+110%	+25%	0.23
Carter et al. (2021)	+78%	+26%	0.33
Suri (2011)	+18%	+8%	0.44
Duflo et al. (2008)	+100%	+20%	0.20
	Panel B: General Equilibrium Models		
	Δ Fertilizer Price	Δ Welfare	Ratio: (2) to (1)
Paper	(1)	(2)	(3)
This paper	+2,068%	-4.4%	-0.002
This paper (smaller shock)	+300%	-1.8%	-0.006
Bergquist et al. (2023)	+300%	-3.5%	-0.012
Sotelo (2020)	-6.3%	.	-0.300
Farrokhi and Pellegrina (2023)	+67 %	-1.8%	-0.027

Notes: For each listed paper, Panel A shows the % changes in the quantity of fertilizer used by farmers and in crop yields in the paper's main counterfactual exercise. In the first two rows (i.e. for this paper), "yield" refers to rice yields. The second row is a "mirror" of the first row in the sense that it shows the changes in fertilizer quantities and yields that would be needed to undo the changes reported in the first row. The third row shows the "partial-equilibrium" elasticity of yield to fertilizer, computed as the average of coefficient γ_k^f across crops (weighted by each crop's national revenue at baseline). For Beaman et al. (2013), the change in Column (2) refers to a change in revenues; we assume there were no price changes so that a revenue change can be interpreted as a yield changes. Panel B shows, for each listed paper, the % changes in fertilizer prices and in welfare in the paper's main counterfactual exercise. In the first two rows (which correspond to this paper), the fertilizer price is gross of subsidies. The second row shows the results of an additional counterfactual exercise in which the imposed QRs are relatively lax, so that fertilizer prices (gross of subsidies) increase by only 300% (see Section 7.3). In both panels, Column (3) shows the ratio between the variables in Column (2) and (1).

in terms of an increase in fertilizer use, to make them more comparable to the literature.[71] For each row, Column (3) presents the ratio of the % change in yields to the % change in fertilizer use, which can thus be interpreted as an elasticity. While this elasticity ranges from 0.20 to 0.44 in other papers, it is much lower (0.022) in our analysis. A possible explanation for this difference is the presence of GE effects: a country-level reduction in fertilizer use (such as the one caused by the import ban) reduces labor demand in agriculture, thus reducing wages and making it cheaper to hire more workers. When farms then hire more workers, this "undoes" part of the initial negative effects on yields. This wage mechanism is usually absent from the RCT literature due to its focus on small-scale interventions that do not affect the wider economy. We confirm the importance of such GE effects by leveraging our agricultural production functions (estimated in Section 6.1.3) to compute a "partial-equilibrium" effect of fertilizer on crop yields.[72] The resulting elasticity, displayed in the

[71]Decreasing a variable by 96.5% is equivalent to dividing it by 28.57. So, to undo (or "mirror") this decrease, we must multiply by 28.57, which is equivalent to an increase of 2,757%. By the same logic, a decrease of 38% is mirrored by an increase of 61%.

[72]To compute the partial equilibrium elasticity of crop yields with respect to fertilizer quantity, we start from the fact that, for a specific crop k, this elasticity is simply the Cobb-Douglas coefficient γ_k^f (see equation 1). We then average

third row of Panel A, is almost five times higher than our corresponding GE elasticity (0.022), and in the ballpark of the partial equilibrium elasticities in the literature.

Panel B of Table 7 compares our results to a different branch of the fertilizer literature that uses GE trade models to assess the welfare effects of changes in technology or policy relating to agricultural inputs. The first row reproduces our main findings (Table 4): the ban increased fertilizer prices by 2,068% and decreased welfare (measured by the average EV of the representative agent) by 4.4%. Since the magnitude of our shock far surpasses anything examined in other papers, in the second row we show results of a "smaller shock" counterfactual (see Section 7.3) in which fertilizer prices rise by only 300%, which is closer to other papers' shock sizes and thus more comparable. In addition to the change in fertilizer prices, Panel B also displays the change in welfare and the ratio between the two variables (i.e. the elasticity) for each paper. Again, our ratio of -0.006 has a lower magnitude than all other papers. This may be partially due to the context under study, since Bergquist et al. (2023) study Ugandan agriculture and Farrokhi and Pellegrina (2023) study a mix of countries from multiple continents. As African countries tend to use less fertilizer than South Asian countries (such as Sri Lanka), their production and welfare may benefit more from marginal increases in use, given decreasing returns to fertilizer penetration. Another explanation for the difference between our findings and the literature is in the nature of the treatment: Bergquist et al. (2023) subsidize not only fertilizer but also the use of special seeds, and Farrokhi and Pellegrina (2023) feature an additional adjustment channel in which farmers shift their practices from traditional to modern once fertilizer becomes sufficiently accessible.

7.5 Influence of Non-Homothetic Preferences and Subsidies

Our model incorporates non-homothetic preferences and fertilizer subsidies, both of which are notable features of Sri Lanka's agricultural economy.[73] In this section, we show that neglecting to include these important features in the model would substantially alter its predictions regarding the welfare effects of the fertilizer import ban, thus leading to biased estimates.

To show the importance of non-homothetic preferences, we perform an exercise in two steps. First, starting from the (non-homothetic) baseline equilibrium estimated in Section 6, we change the key PIGL preference parameters to $\eta = 1$ and $\nu = 0$ (under which PIGL collapses to homothetic, Cobb-Douglas preferences) and compute a "homothetic baseline equilibrium" by applying Definition 3 with no QRs. Second, we apply Definition 3 again to recompute the homothetic equilibrium

this γ_k^f coefficient across crops, weighted by each crop's share of total agricultural revenue in Sri Lanka. Formally, the partial-equilibrium elasticity is: $\sum_k \gamma_k^f \times \left(\frac{\sum_{i=1}^I p_{ik} Q_{ik}}{\sum_l \sum_{i=1}^I p_{il} Q_{il}}\right)$, where $p_{ik} Q_{ik}$ is the revenue earned by crop-k agriculture in district i at baseline.

[73] For direct evidence of non-homothetic preferences, see Appendix C.2.

but now with QRs ($\bar{f} = 0$), as in Section 7.[74] Therefore, this exercise amounts to an investigation of the predictions our model would make if it did not feature non-homothetic preferences.

Comparing Columns (1) and (2) of Table 6, we see that the predicted effect of the ban on the welfare of mobile workers is less negative in a homothetic economy than in the non-homothetic economy. This finding is explained by the fact that agriculture comprises 24.9% of employment (on average) in the non-homothetic baseline equilibrium but only 15% in the homothetic baseline equilibrium.[75] In other words, the baseline agricultural (resp. manufacturing) sector is smaller (resp. larger) in the homothetic case. A large manufacturing sector acts as buffer that attenuates the ban's negative effects on mobile workers, as explained in Section 7.2. That is why mobile workers suffer less in the homothetic case.

We then run a similar exercise to demonstrate the importance of fertilizer subsidies. Starting from the baseline equilibrium estimated in Section 6, we set subsidy rates ($\{s_k\}_{k=1}^{K}$) to zero, compute the resulting baseline equilibrium, and then impose QRs and recompute this equilibrium. Column (3) of Table 6 shows that the absence of subsidies intensifies the welfare losses of mobile workers while ameliorating the corresponding losses for tea estate workers and the representative farmer. The intuition for this finding comes from the fact that subsidies are financed through a flat tax on the income of all agents. In practice, such scheme corresponds to a net income transfer from mobile workers to farmers and estate workers. By drastically reducing the use of fertilizer in the country, the ban also reduces the total amount of subsidies and, consequently, the scale of the transfer scheme, thus helping mobile workers and hurting farmers and tea workers. On the other hand, in a world without subsidies this transfer scheme does not exist in the first place, so the "subsidy scale-down" channel that we have just described is absent from the ban's welfare effects. Hence the worse (resp. milder) welfare losses of mobile workers (resp. farmers and estate works) in the case without subsidies.[76]

It is also interesting to note that, in this no-subsidy exercise, the ban decreases fertilizer imports by 197 million USD and total agricultural exports by 278 million USD. That is, foreign exchange "savings" through lower imports are *more than completely offset* by a corresponding reduction in agricultural exports, in contrast to a mere 48% "offset factor" in our main estimates (see Section 7.2).[77] In the absence of fertilizer subsidies, there is a closer correspondence between the import price of fertilizer and its marginal revenue product. As a result, the reduction in fertilizer use

[74] While this exercise changes the value of the two PIGL parameters ν and η, the remaining parameters and unobservable variables are kept at their original values estimated in Section 6.

[75] The ILO estimates that Sri Lanka had 25% of its total employment in agriculture in 2019. Thus, the non-homothetic model matches the ILO data much more closely than the homothetic model.

[76] Note that this also implies a smaller gap in welfare losses between workers and farmers relative to the case with subsidies (Column 1 of Table 6), leading us to underestimate how unequal the incidence of the ban is across occupations.

[77] The 48% offset factor is computed by comparing the decrease in agricultural exports (596 million USD) to the decrease in fertilizer imports (1.253 billion USD) in our main estimates, which are displayed in Table 4.

induced by the ban decreases agricultural export revenues and fertilizer imports by roughly similar amounts.

Finally, Column (4) of Table 6 shows the results of a combined exercise with both homothetic preferences *and* subsidies set to zero. Across all the four cases shown in the Table, this is the one with the mildest welfare losses for the representative agent. On the other extreme, welfare losses are the worst in the main counterfactual exercise (Column 1), which features both fertilizer subsidies and non-homothetic preferences.[78] Therefore, the more features we neglect to include in the model, the more severely we will underestimate the ban's negative effects.

7.6 Influence of Modeling Trade Policy: Iceberg Trade Costs versus Quantitative Restriction

In our model, trade policy is represented by QRs. This differs from much of the NTM literature, which models NTMs as an increase in multiplicative "iceberg" trade costs.[79] These modeling choices are often not innocuous: while QRs raise prices and create rents that can be captured by certain domestic agents, the price increases resulting from iceberg trade costs are typically modeled as losses, with no one benefiting from them. Therefore, using iceberg trade costs to model what is really a QR tends to lead to overestimation of the welfare losses induced by the policy.

To illustrate this point, we perform two additional counterfactual exercises in which we use different means to reduce the quantity of imported fertilizer by 50%. In the first exercise, this is achieved via QRs, while in the second exercise it is done by increasing the iceberg trade cost of transporting fertilizer from the RoW to Sri Lanka.[80] Results are displayed in Columns (5) and (6) of Table 6. As expected, for all types of agents, welfare losses are less severe under QRs than under iceberg cost increases, with the average welfare effect on mobile workers even switching from negative to positive in the former case.[81]

[78]Intuitively, the model with homothetic preferences (Column 2) underestimates the mobile worker's true losses (Column 1) while estimating farmer and estate worker losses reasonably accurately. Then, removing subsidies from the model (Column 4) radically decreases the magnitude of farmer and estate worker losses relative to Column (2), thus also reducing the magnitude of the representative agent's (i.e. the "average") losses.

[79]One exception is Atkin et al. (2024), who explicitly model a discretionary import licensing system.

[80]Note that, regardless of the trade policy instrument, there are no rents to be captured in the extreme cases (such as Sri Lanka's ban) in which imports are driven down to zero. That is why we perform the comparison exercise between the two instruments using a less restrictive (50%) reduction in the quantity of imported fertilizer.

[81]The mechanism by which welfare losses are attenuated in the QR case is related to the tax rate. In both cases, reduction in fertilizer imports leads to a reduction in fertilizer use in Sri Lanka, which scales down the fertilizer subsidy program, lowering the tax rate that is needed to fund such program. However, the tax rate drops by a greater magnitude in the QR case (0.77 percentage points) than in the iceberg cost case (0.4 percentage points).

7.7 Model Validation: Evidence from Night Lights

Our model-driven counterfactual exercise replicates the stylized facts documented in Section 4, with large decreases in fertilizer imports, crop yields, and agricultural exports. To further validate our model, we now show that the counterfactual's predictions of the ban's effects on regional incomes also align with the income decreases observed in the data, as proxied by night light intensity.

A well-established literature (Henderson et al., 2012) has confirmed the validity of night lights as a reliable way to indirectly assess the spatial distribution of per capita income when direct income metrics with the desired coverage or frequency are unavailable. The elasticity of per capita income to the average night light intensity of a geographical unit is often estimated to be around 0.3. We combine this elasticity with the observed reduction in average night light intensity in each Sri Lankan district between April 2021 (i.e. the last month before the ban was implemented) and January 2022 to estimate districts' income losses in that period.[82] Estimated income losses are displayed in Figure 9a. The decreases in district incomes predicted by the counterfactual exercise are displayed in Figure 9b. A visual comparison of the two figures suggests that the model performs well in terms of replicating the spatial pattern of income losses during the post-ban period. For example, both maps show that the north of Sri Lanka was particularly severely affected relative to the west. Moreover, Figure 9c shows that these northern regions are also the ones most economically dependent on fertilizer, thus explaining why they were so affected by the ban.[83]

To more formally assess the model's ability to explain the observed post-ban changes in night lights, we run a dynamic event study analysis at the district-month level, assigning districts to the treatment or control group depending on whether they were above or below the median of fertilizer importance at baseline (see Figure 9c). The outcome variable is (the log of) average night light intensity and the treatment period is defined as starting in May 2021 (i.e. the month when the ban was implemented). The results, displayed in Appendix Figure A15, show that night lights decreased by 26.7 log points in treated (i.e. relatively fertilizer-dependent) districts relative to control (i.e. less fertilizer-dependent) districts eight months into the ban, corresponding to an income loss of 7.7%. Using the same assignment of districts into treated and control groups, the model-predicted average income losses (Figure 9b) are 8.2 percentage points worse in treated districts than in control districts. Therefore, by comparing the model's predictions to the night lights evidence, we conclude that the model can explain the whole differential income loss of fertilizer-dependent districts observed in the data during the relevant time period.

Finally, we take advantage of our night lights data to further confirm that our main results are

[82] For details on our night lights data, see Appendix A.9.

[83] To measure the economic importance of fertilizer for a district, we use the district's baseline ratio of fertilizer expenditure to aggregate income, as predicted by our model. This metric captures both the relative importance of the agricultural sector in the district's economy and the agricultural sector's degree of specialization in fertilizer-intensive crops.

Figure 9: Fertilizer Importance and Changes in Night Lights, by district

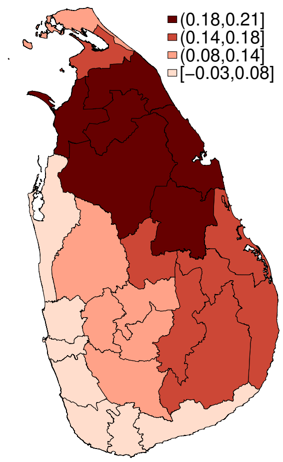

(a) Fall in log income implied by change in night lights (April 2021 - January 2022)

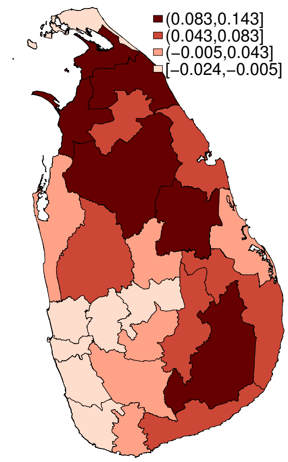

(b) Fall in log income predicted by model's counterfactual

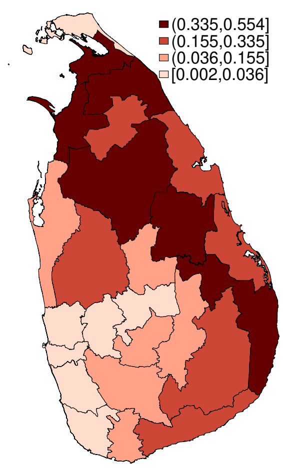

(c) Fertilizer's importance in local economy (at baseline)

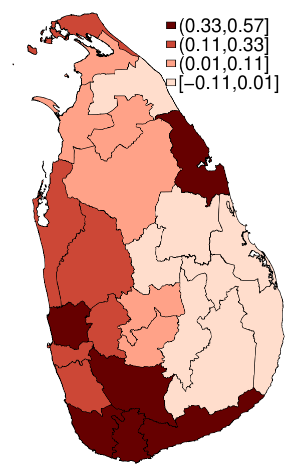

(d) Fall in log income implied by change in night lights (February 2020 - November 2020)

Notes: Panel 9a shows the reduction in the log of district income between April 2021 and January 2022, as predicted by the reduction in districts' night light intensity in the same period. Panel 9d displays the same variable as Panel 9a but for the period between February and November 2020. Panel 9c shows each district's baseline fertilizer importance index, defined as the ratio of the value of fertilizer inputs (gross of subsidies) used by its agricultural sector to the district's aggregate income in our estimated baseline model. Panel 9b shows the decline in (the log of) net-of-tax income in each district predicted by our model-generated counterfactual exercise (see Section 7.2), where this income decline is defined as the EV of the district's representative agent.

not driven by the potentially confounding influence of the Covid-19 pandemic (see Section 4.1). Figure 9d shows that the spatial pattern of income losses (proxied by night lights) in the period following the pandemic's arrival in March 2020 was very different from the pattern of losses in the post-ban period (Figure 9a), indicating that potentially lingering effects of the pandemic cannot explain the negative economic effects that we attribute to the fertilizer ban.

8 Conclusion

This paper leverages an unprecedented natural experiment, an abrupt and unexpected import ban that disrupted the supply of chemical fertilizers to Sri Lanka, to quantify the costs for agricultural production and trade of a lack of access to fertilizer in the context of a developing economy where agriculture is of central importance. Relying on novel high-frequency firm-level trade data, detailed agricultural ground production data, crop yield estimates from state-of-the-art remote sensing techniques, and cutting-edge event study designs, we show that the fertilizer import ban led to dramatic declines in fertilizer imports, agricultural production, and exports of fertilizer-intensive agricultural products.

To better understand the mechanisms and welfare implications of the ban's effects, we propose and estimate a QSM of agriculture and trade with QRs on imports that reflects key realistic features of many developing countries, such as high fertilizer subsidies, land inequality, and non-homothetic preferences. We show that these features are important when considering the aggregate and distributional impacts of trade policy. For example, heavy fertilizer subsidies, a domestic agricultural policy common in many countries, interact with trade policy: an import ban nearly eliminates fertilizer use, scaling down the income transfers from non-farming to farming sectors that are engineered through subsidies, and thus attenuating the relative welfare losses of mobile workers. Non-homothetic preferences, by determining the share of income spent on food and thus the size of the agricultural and manufacturing sectors, determine the outside option that workers face when considering whether to move away from the ban-affected sector, thus affecting welfare. The model implies that the ban's welfare effects vary across Sri Lankan regions, depending on the region's propensity to grow fertilizer-intensive crops. We find that average losses due to the lack of fertilizer were equivalent to a 4.35% income decline, with stronger incidence on farmers and estate workers (whose income is tied to agriculture) relative to mobile workers. The ban lowered crop yields by 12%-54%, leading to losses of 596 million USD in foregone agricultural exports, which offsets around half of the foreign exchange savings from reduced fertilizer imports. Our findings quantify the GE costs for agriculture of a lack of fertilizer, an estimate that helps inform fertilizer-related policy (such as fertilizer subsidies).

References

Aggarwal, S., Giera, B., Jeong, D., Robinson, J. and Spearot, A. (2022). Market access, trade costs, and technology adoption: Evidence from northern Tanzania, *Review of Economics and Statistics* pp. 1–45.

Amaglobeli, D., Benson, T. and Mogues, T. (2024). Agricultural producer subsidies: Navigating challenges and policy considerations.

Amarasinghe, D. and Bandara, J. S. (2005). Input output tables for Sri Lanka – 2000, *Research Studies: Macroeconomic Policy Series No. 16*, Institute of Policy Studies, Colombo, Sri Lanka.

Artuc, E., Chaudhuri, S. and McLaren, J. (2010). Trade shocks and labor adjustment: A structural empirical approach, *American Economic Review* **100**(3): 1008–1045.

Artuc, E., Falcone, G., Porto, G. and Rijkers, B. (2024). Protectionism, evasion and household welfare: Evidence from nigeria's import bans, *Available at SSRN 4931914* .

Artuc, E., Porto, G. and Rijkers, B. (2023). Crops, conflict and climate change, *Working paper*.

Atkin, D., Blaum, J., Fajgelbaum, P. and Ospital, A. (2024). Trade barriers and market power: Evidence from Argentina's discretionary import restrictions, *NBER working paper 32037*.

Baquedano, F., Zereyesus, Y. A., Valdes, C. and Ajewole, K. (2021). International food security assessment, 2021–32, *GFA-32*, US Department of Agriculture, Economics Research Service.

Bauluz, L., Govind, Y. and Novokmet, F. (2020). Global land inequality.

Beaman, L., Karlan, D., Thuysbaert, B. and Udry, C. (2013). Profitability of fertilizer: Experimental evidence from female rice farmers in Mali, *American Economic Review* **103**(3): 381–386.

Behrman, J. R. and Deolalikar, A. (1989). Is variety the spice of life? Implications for calorie intake, *The Review of Economics and Statistics* pp. 666–672.

Bergquist, L. F., Faber, B., Fally, T., Hoelzlein, M., Miguel, E. and Rodríguez-Clare, A. (2023). Scaling agricultural policy interventions, *Working paper*.

Bernini, F., García-Lembergman, E. and Juarez, L. (2024). The Consequences of Non-Tariff Trade Barriers: Theory and Evidence from Import Licenses in Argentina (Preliminary Results), *Discussion Paper No IDB-DP-1063*, Inter-American Development Bank, Department of Research and Chief Economist.

Bold, T., Kaizzi, K. C., Svensson, J. and Yanagizawa-Drott, D. (2017). Lemon technologies and adoption: Measurement, theory, and evidence from agricultural markets in Uganda, *The Quarterly Journal of Economics* **132**(3): 1055–1100.

Boppart, T. (2014). Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences, *Econometrica* **82**(6): 2167–2196.

Brooks, W. and Donovan, K. (2024). Global conflict and local agricultural policy in rwanda.

Bustos, P., Caprettini, B. and Ponticelli, J. (2016). Agricultural productivity and structural transformation: Evidence from Brazil, *American Economic Review* **106**(6): 1320–65.

Caliendo, L., Dvorkin, M. and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the China trade shock, *Econometrica* **87**(3): 741–835.

Callaway, B. (2023). Difference-in-differences for policy evaluation, *in* K. F. Zimmermann (ed.), *Handbook of Labor, Human Resources and Population Economics*, p. 1–61.

Carleton, T., Crews, L. and Nath, I. (2024). Agriculture, trade, and the spatial efficiency of global water use, *Working paper*.

Carter, M., Laajaj, R. and Yang, D. (2021). Subsidies and the African Green Revolution: Direct effects and social network spillovers of randomized input subsidies in Mozambique, *American Economic Journal: Applied Economics* **13**(2): 206–229.

Chandrabose, A. S. (2015). Outgoing labour and its impact on the tea plantation sector in Sri Lanka, *Proceedings of 5th International Symposium on Emerging Trends and Challenges in Multidisciplinary Research*, South Eastern University of Sri Lanka.

Comin, D., Lashkari, D. and Mestieri, M. (2021). Structural change with long-run income and price effects, *Econometrica* **89**(1): 311–374.

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

De Chaisemartin, C. and d'Haultfoeuille, X. (forthcoming). Difference-in-differences estimators of intertemporal treatment effects, *The Review of Economics and Statistics* .

Desmet, K. and Rossi-Hansberg, E. (2015). On the spatial economic impact of global warming, *Journal of Urban Economics* **88**: 16–37.

Dhingra, S., Freeman, R. and Huang, H. (2023). The impact of non-tariff barriers on trade and welfare, *Economica* **90**: 140–177.

Dissanayaka, N. and Thibbotuwana, M. (2021). Sri Lanka's Agri-food Trade: Structure, Opportunities, Challenges Impacts of COVID-19, *Food Security Policy Research, Capacity, and Influence (PRCI) Research Paper 6*, Michigan State University.

Dong, J., Xiao, X., Menarguez, M. A., Zhang, G., Qin, Y., Thau, D., Biradar, C. and Moore III, B. (2016). Mapping paddy rice planting area in northeastern Asia with Landsat 8 images, phenology-based algorithm and Google Earth Engine, *Remote Sensing of Environment* **185**: 142–154.

Donovan, K. (2021). The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity, *The Review of Economic Studies* **88**(5): 2275–2307.

Druilhe, Z. and Barreiro-Hurlé, J. (2012). Fertilizer subsidies in sub-saharan africa.

Duflo, E., Kremer, M. and Robinson, J. (2008). How high are rates of return to fertilizer? Evidence from field experiments in Kenya, *American Economic Review* **98**(2): 482–488.

Duflo, E., Kremer, M. and Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya, *American Economic Review* **101**(6): 2350–2390.

Eckert, F. and Peters, M. (2023). Spatial structural change, *Working paper*.

Ederington, J. and Ruta, M. (2016). Nontariff measures and the world trading system, *in* K. Bagwell and R. W. Staiger (eds), *Handbook of Commercial Policy*, Elsevier, chapter 5, pp. 211–277.

Fajgelbaum, P. and Redding, S. J. (2022). Trade, structural transformation, and development: Evidence from argentina 1869–1914, *Journal of Political Economy* **130**(5): 1249–1318.

Fan, T., Peters, M. and Zilibotti, F. (2023). Growing Like India– The Unequal Effects of Service-led Growth, *Econometrica* **91**(4): 1457–1494.

FAO and WFP (2022). *Special report - FAO/WFP crop and food security assessment mission (CFSAM) to the Democratic Socialist Republic of Sri Lanka*, United Nations, Rome.

FAO and WFP (2023). *Special report - FAO/WFP crop and food security assessment mission (CFSAM) to the Democratic Socialist Republic of Sri Lanka*, United Nations, Rome.

Farrokhi, F. and Pellegrina, H. S. (2023). Trade, technology, and agricultural productivity, *Journal of Political Economy* .

Feenstra, R. C., Luck, P., Obstfeld, M. and Russ, K. N. (2018). In search of the Armington elasticity, *The Review of Economics and Statistics* **100**(1): 135–150.

Felt, M.-H., Gervais, J.-P. and Larue, B. (2011). Market power and import bans: The case of Japanese pork imports, *Agribusiness* **27**: 47–61.

Fernandes, A., Freund, C. and Pierola, D. (2016). Exporter behavior, country size and stage of development: Evidence from the exporter dynamics database, *Journal of Development Economics* **119**(C): 121–137.

Flaaen, A., Haberkorn, F., Lewis, L., Monken, A., Pierce, J., Rhodes, R. and Yi, M. (2023). Bill of lading data in international trade research with an application to the COVID-19 pandemic, *Review of International Economics* **31**(3): 1146–1172.

Foster, A. D. and Rosenzweig, M. R. (2008). Economic development and the decline of agricultural employment, *in* T. P. Schultz and J. A. Strauss (eds), *Handbook of Development Economics*, Vol. 4, Elsevier, chapter 47, p. 3051–83.

Ghose, D., Khan, S. and Montfaucon, A. (2024). Unpacking non-tariff measures: Beyond the tip of the iceberg, *Technical Note, World Bank* .

Gollin, D., Lagakos, D. and Waugh, M. E. (2014). Agricultural productivity differences across countries, *American Economic Review* **104**(5): 165–170.

Gollin, D., Parente, S. and Rogerson, R. (2002). The role of agriculture in development, *American Economic Review* **92**(2): 160–64.

Hebebrand, C. and Glauber, J. W. (2023). The Russia-Ukraine war after a year: Impacts on fertilizer production, prices, and trade flows, *The Russia-Ukraine conflict and global food security*, International Food Policy Research Institute (IFPRI), chapter 8, pp. 43–47.

Henderson, J. V., Storeygard, A. and Weil, D. N. (2012). Measuring economic growth from outer space, *American Economic Review* **102**(2): 994–1028.

Kishore, A., Alvi, M. and Krupnik, T. J. (2021). Development of balanced nutrient management innovations in south asia: perspectives from bangladesh, india, nepal, and sri lanka, *Global Food Security* **28**: 100464.

Leroux, A. and Maclaren, D. (2011). The optimal time to remove quarantine bans under uncertainty: the case of Australian bananas, *Economic Record* **276**: 140–152.

McArthur, J. W. and McCord, G. C. (2017). Fertilizing growth: Agricultural inputs and their effects in economic development, *Journal of development economics* **127**: 133–152.

McDonald, M. B. J., Ruta, M. and Van Heuvelen, M. E. (2024). *Industrial policy: Trade policy and World Trade Organization considerations in IMF surveillance*, Vol. 2, International Monetary Fund.

Muralidharan, K. and Niehaus, P. (2017). Experimentation at scale, *Journal of Economic Perspectives* **31**(4): 103–124.

Nath, I. (forthcoming). Climate change, the food problem, and the challenge of adaptation through sectoral reallocation, *Journal of Political Economy* .

Nigai, S. (2016). On measuring the welfare gains from trade under consumer heterogeneity, *The Economic Journal* **126**: 1193–1237.

Ozdogan, M., Wang, S., Ghose, D., Fraga, E., Fernandes, A. and Varela, G. (2024). Field-scale rice area and yield mapping in Sri Lanka with optical remote sensing and limited training data, *Working paper*.

Pellegrina, H. S. (2022). Trade, productivity, and the spatial organization of agriculture: Evidence from Brazil, *Journal of Development Economics* **156**(102816): 1–19.

Peterson, E. and Ord, D. (2008). Avocado pests and avocado trade, *American Journal of Agricultural Economics* **90**(2): 321–335.

Sotelo, S. (2020). Domestic trade frictions and agriculture, *Journal of Political Economy* **128**(7): 2690–2738.

Suri, T. (2011). Selection and comparative advantage in technology adoption, *Econometrica* **79**(1): 159–209.

Tombe, T. (2015). The missing food problem: Trade, agriculture, and international productivity differences, *American Economic Journal: Macroeconomics , 7(3):* **7**(3): 226—-258.

UNCTAD (2019). *International Classification of Non-tariff Measures - 2019 version*, United Nations, Geneva, Switzerland.

Weerahewa, J., Senaratne, A. and Babu, S. (2021). Reforming fertilizer import policies for sustainable intensification of agricultural systems in sri lanka: Is there a policy failure?

Wijesinghe, A., Yogarajah, C. and Rathnayake, N. (2023). Import controls in Sri Lanka: Political preference and incentive distortions, *International Economics Series: No. 15*, Institute of Policy Studies of Sri Lanka, Colombo, Sri Lanka.

Zereyesus, Y. A., Cardell, L., Valdes, C., Ajewole, K., Zeng, W., Beckman, J., Ivanic, M., Hashad, R. N., Jelliffe, J. and Kee, J. (2022). International food security assessment, 2022–32, *USDA Miscellaneous 329074*, US Department of Agriculture.

Appendices

A Data Appendix

In this section, we list and describe the multiple data sources we use in this paper, as well as the necessary data cleaning and adjustment procedures we perform. Where relevant, we show descriptive statistics based on the data. We cover not only the data sources mentioned in Section 3, but also other ancillary data sources.

A.1 Trade Policy

We obtained Sri Lanka's Extraordinary Gazettes on Imports and Exports (Control Regulations) from March 2020 to October 2022, which allow us to identify the HS8 products that were subject to import bans.[84] We digitized the gazettes and manually inspected all entries to create a database with HS8 products and start and end dates of import bans. Appendix Figure A1 shows the timing of the fertilizer import bans, while Appendix Table A1 shows the list of HS8 product codes corresponding to fertilizers subject to the import ban.

Figure A1: Timing of Fertilizer Import Bans

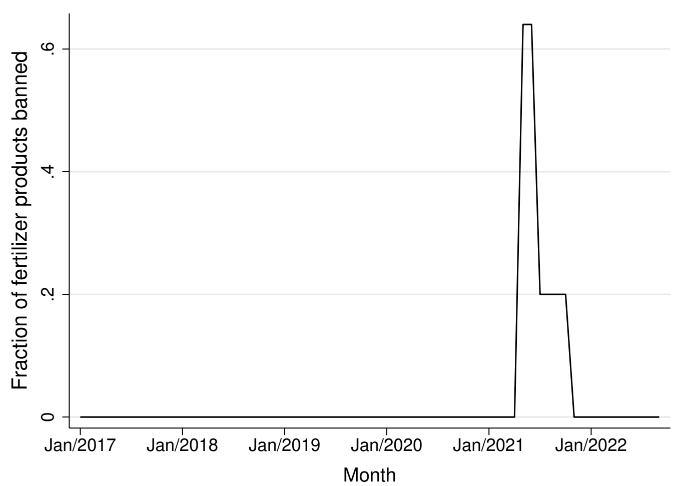

Notes: For each month between January 2017 and September 2022, the chart shows the fraction of HS8 fertilizer products that were subject to an import ban in that month. The denominator of the fraction is 25, which is the number of fertilizer products that were imported by Sri Lanka at any point between 2017 and 2022. A fertilizer product is defined as any HS8 product that is listed in Chapter 31 of the HS system. The data comes from Sri Lanka's Extraordinary Gazettes.

[84] The Extraordinary Gazettes were accessed at: http://www.imexport.gov.lk/index.php/en/downloads/gazette.html

Table A1: List of Fertilizer Products Subject to Import Bans

HS8	Description
31024000	Mixtures of ammonium nitrate with calcium carbonate or other inorganic non-fertilizing substances
31026000	Double salts and mixtures of calcium nitrate and ammonium nitrate
31028000	Mixtures of urea and ammonium nitrate in aqueous or ammoniacal solution
31029000	Other, including mixtures not specified in the foregoing subheadings
31031100	Superphosphates containing by weight 35% or more of diphosphorus pentaoxide (P2O5)
31031900	Other
31039000	Other
31049000	Other
31051000	Goods of this Chapter in tablets or similar forms or in packages of a gross weight not exceeding 10 kg
31052000	Mineral or chemical fertilizers containing the three fertilizing elements nitrogen, phosphorus and potassium
31053000	Diammonium hydrogenorthophosphate (diammonium phosphate)
31054000	Ammonium dihydrogenorthophosphate (monoammonium phosphate) and mixtures thereof with diammonium hydrogenorthophosphate (diammonium phosphate)
31055100	Containing nitrates and phosphates
31055900	Other
31056000	Mineral or chemical fertilizers containing the two fertilizing elements phosphorus and potassium
31059000	Other

Notes: The table displays the list of HS8 fertilizer products that were subject to import bans at some point in 2021. An HS8 product is classified as fertilizer if it is listed in Chapter 31 of the HS system. The data is from Sri Lanka's Extraordinary Gazettes.

A.2 Trade Flows

We use novel high-frequency granular trade data obtained from S&P Global Market Intelligence's Panjiva data platform covering the universe of Sri Lanka's export and import transactions from January 2017 to October 2022 based on bills of lading (BoL). The data includes the name and address of the shipper (exporter) and the consignee (importer) of each shipment as well as HS8 product codes. Exporting and importing firms are identified by their names, as BoL data does not require inclusion of tax identification numbers.[85] We use machine learning and text analysis techniques to clean firm names and assign to each name a unique firm identifier, thus constructing a panel of Sri Lankan firms engaged in exports or imports. Appendix A.2.1 below provides details on our fuzzy matching algorithms.

The data is subject to a series of cleaning procedures as in Fernandes et al. (2016). Specifically, we drop observations for which trade value is zero or missing, product codes are missing or not in the HS 2017 revision list, partner country is missing or is Sri Lanka (which may indicate trans-

[85]While Panjiva adds a firm identifier to the BoL data that Flaaen et al. (2023) use for the United States, we found such identifier to be extremely unreliable for Sri Lankan firms. The same firm appearing with a different name spelling has a different identifier, and a given firm making an export transaction has a different identifier than the same firm making an import transaction. On any given day, the same firm transacting more than once could also have different Panjiva identifiers.

actions with Special Economic Zones, but this cannot be systematically assessed). Data quality is assessed, as in Fernandes et al. (2016), by a comparison of total exports and total imports obtained from aggregating the transaction-level data at year level with total yearly exports and imports for Sri Lanka from World Integrated Trade Solution (WITS) and from the Central Bank of Sri Lanka. The ratios to WITS in 2017-2021 are in the 97%-100% range for exports and the 83%-94% range for imports. To obtain our final datasets, we drop exports and imports of oil (HS Chapter 27), which are generally not well captured in customs datasets. To focus on true commercial export and import transactions by firms, we also drop transactions of currency paper notes (HS 490700) conducted only by the Central Bank of Sri Lanka, of large machinery entering and exiting Sri Lanka temporarily for construction projects such as dredgers (vessels) (HS 890510, HS 890190), and of arms and ammunition (HS Chapter 93).

We aggregate transaction-level values (in USD) and net weights (in kg) up to the HS8 product-month-year level for our import analysis by summing across firms and source countries. We aggregate transaction-level values (in USD) up to the firm-HS8 product-year level for our export analysis by summing across destination countries. For our export analysis, we focus on agricultural products, which are listed in HS Chapters 06-24.

A.2.1 Cleaning Firm Names

The names of Sri Lankan firms engaged in export or import transactions are noisy. Names appear with multiple spellings or with spelling errors, and some entries are a number or an address instead of a proper names. In brief, the steps for the cleaning of names are as follows:

1. Eliminate from the list of firm names those entries that are not proper names (numbers, addresses, and the expressions 'NIL', 'N/A', or 'To the order').

2. Remove from firm names a list of stop words, geographic stop words, and symbols that are not letters nor numbers. The list of stop words was selected based on text analysis indicating the most common words appearing in firm names, such as 'corporation', 'international', 'limited', and 'llc'. We also remove prefixes and suffixes (such as 'co', 'sa', and 'pvc') from firm names.

3. Based on the list of pre-processed firm names obtained in Step 2, we use N-gram similarity (using cosine distance metrics) to identify potential pairs of similar firm names (using a lower similarity threshold of 0.6, where 1 means two names are identical).

4. Potential pairs of similar firm names identified in Step 3 go through an algorithm to more precisely determine if they are really similar based on Levenshtein distance (a measure of similarity between two strings). Distance cutoffs defining what is a "small" distance between a pair of names are determined through machine learning algorithms based on subsamples of firm name pairs whose similarity was determined through manual inspection.

5. The pairs of similar firm names identified in Step 4 are then sorted alphabetically to identify neighbor names (i.e., names that appear in consecutive rows after sorting) and correct any potential similarities that may have been missed in Step 4.

6. The algorithms in Step 4 generate “big groups” of firm names that are considered similar because they share several common words. To break up these “big groups”, firm names are sorted alphabetically within each group and, if two consecutive firm names do not have sufficiently similar pre-processed names (as captured by a cutoff in the Levenshtein distance), the group is “broken” at that place, resulting in two separate groups of similar firm names.

7. After similar firm names are identified from the steps above and assigned temporary unique numeric identifiers, firm names are then sorted alphabetically and based on firm address. If two consecutive names have similar addresses and similar names and did not have the same unique numeric identifier, they are joined in a new final unique numeric identifier.

The steps above are applied to clean the names of Sri Lankan firms. A large number of transactions, especially in the import data, are made by firms whose names indicate they are courier companies. Since these transactions are made for a third party, frequently individuals, and since their total trade value is negligible, we dropped all observations whose names are of courier firms.[86] Finally, we also drop from the sample observations for which the firm name corresponds to an individual making personal imports under a permit.[87] Such observations account for a very small total trade value.

A.3 Agricultural Production, Cultivated Area, and Yields

We rely on both administrative data and remote sensing estimates.

A.3.1 Administrative Data

For rice, we obtain data on production and cultivated area in each district from the ‘Paddy Statistics’ produced by the DCS for each growing season from 2012-2013 Maha to 2022 Yala.[88] Crucially, this time frame includes the 2019-2020 Maha and 2020 Yala seasons, which we use as the pre-ban baseline to estimate the model in Section 6.2.

For non-rice, non-tea crops (maize, groundnuts, potatoes, onions, cinnamon, and cloves), we obtain data on production and cultivated area at the divisional secretariat (DS) level from the DCS for each growing season from Maha 2019-2020 to Yala 2022, which we aggregate up to the district level. For tea, the DCS data includes information on cultivated area but not on production, so we

[86] Namely: “dhl express parcel delivery”, “d h l”, “fits aviation”, “cargo logistics worldwide aramex courier freight” “non commercial xpress by air freight” “Skynet shipping setrans” “sky international”.

[87] Such firm names have the structure “Person Name+Institution+PERMIT”. For example: “KADAWARAGE C V WING COMMANDER (DENTAL DOCTOR) SRI LANKA AIR FORCE MIN.OF DEFENCE PERMIT”.

[88] Paddy Statistics were accessed at: http://www.statistics.gov.lk/Agriculture/StaticalInformation/Paddy_Statistics.

obtain district-level tea production data from the Sri Lanka Tea Board's annual reports for 2019-2022.

Production is measured in kilograms (kg) and cultivated area is measured in hectares (ha). For crops with harvests in both Yala and Maha seasons (maize, rice, and groundnuts), we compute total production and total cultivated area in each year and district by summing across the two seasons. For each crop, we obtain yields (in kg/ha) as the ratio between production and cultivated area.

To estimate fertilizer coefficients in our model's Cobb-Douglas agricultural production functions (see Section 6.1.3), we use national production data for each crop in 2022, in metric tons (MT), from the Central Bank of Sri Lanka's Economic and Social Statistics 2023 report.[89] We use national production data for cinnamon, cloves, groundnuts, maize, onions, potato, rice, and tea.

A.3.2 Remote Sensing Estimates

We produce novel remote sensing estimates of rice and tea yields and of cultivated rice areas at a highly granular level. For rice, the methodology described in Ozdogan et al. (2024) relies on an expert-based image classification algorithm on satellite observations from Landstat and Sentinel-2 enhanced to isolate the rice signal based on the premise that paddy rice has specific growing conditions in Sri Lanka. The presence of rice-growing areas is determined by factors related to climate, water availability, topographic position, and farmer decision-making and technical expertise. The methodology monitors rice-related vegetation and water management from space and distinguishes it from other land use types using spectral (color) and temporal information included in the satellite signal. Since water and vegetation indexes are more visible for rice than for other crops, remote sensing tools can more accurately map crop type and predict yields for the former than for the latter (Dong et al., 2016). To estimate rice yields, we use a statistical model that correlates district-level rice yields from government statistics (described in Appendix A.3.1) with the satellite-derived vegetation index (green chlorophyll index), known to be sensitive to rice yields for the pixels identified as rice areas. A random forest-based machine learning model is then added to incorporate additional environmental variables known to affect rice yields, resulting in pixel-level rice yield estimates across two decades (2000-2022). These remote sensing estimates, aggregated to the DS level, are shown to be highly consistent with government-reported measures of rice area (from survey-based crop-cutting experiments) and rice yield (from production statistics).

For tea, we regress district-level yield data from Sri Lanka Tea Board reports (see Appendix A.3.1) on a district-level vegetation index derived from Planet Satellite Data for pixels deemed to be tea plantations for the period 2014-2022. Since the Tea Board reports only production, not yields, we divide total production by planted area in each district to obtain yields (in kg/ha). In

[89] The report is available at: https://www.cbsl.gov.lk/en/publications/other-publications/statistical-publications/economic-and-social-statistics-of-sri-lanka/ess-2023.

the satellite analysis, tea cultivation areas are held fixed to the 2021 tea map, which is applied to every year. The underlying assumption is that tea planted areas do not change on an annual basis because tea is a multiyear permanent crop. We also generate a derived confidence index reflecting the precision of the satellite-derived tea maps, which we use as statistical weights in our tea yield regressions in Appendix C.4. To obtain the confidence index, we first derive a tea probability map using an ensemble random forest image classification algorithm applied to tea data. In this map, any pixel with less than 50% probability is identified as a non-tea planted pixel. We remove all non-tea planted pixels, leaving only pixels with probabilities between 51% and 100%. We then average these probabilities across the pixels within each DS and round the resulting value to obtain the confidence index (ranging from 50 to 100) for each DS.

A.4 Agricultural Prices

To estimate our model, we use data on producer prices for each crop (except tea) and district from the 2016-2019 edition of DCS' Bulletin of Selected Retail and Producer Prices.[90] We extract average producer prices (in LKR/kg) in 2019 for rice, potatoes, cinnamon, cloves, onions, groundnuts, and maize. For onion and rice, whose prices are available for multiple varieties, we use the prices of red onions and red rice (raw). Whenever the producer price of a crop is missing for a given district, we impute it using the crop's average price across other Sri Lankan districts. Producer prices are missing for 45% of the district-crop pairs with non-zero agricultural production. Nevertheless, the imputation procedure is unlikely to create biases since producer prices are similar across space: the coefficient of variation of (non-missing) producer prices across districts averages 0.032, ranging from 0.009 to 0.058 across crops. For tea, we use producer price data (in LKR/MT) for 2019 from the Food and Agriculture Organization of the United Nations (FAO) for Sri Lanka as a whole, which we assume to hold in each individual tea-producing district.

The 2018-2021 and 2019-2022 editions of DCS' Bulletin of Selected Retail and Producer Prices are used to compute crop-level nominal producer price increases between January 2021 and January 2022. By subtracting from each crop's nominal price increase the aggregate inflation implied by the National Consumer Price Index during the same period, we obtain each crop's real price increase. This information is used as an input in the calibration of the cross-origin elasticity of substitution of agricultural crops (see Section 6.1.1).

To estimate the fertilizer coefficients of our model's Cobb-Douglas agricultural production functions (see Section 6.1.3), we use data on national crop prices in 2022 (in LKR/kg) from the Central Bank of Sri Lanka's Economic and Social Statistics 2023 report (except for groundnuts, whose national price comes from the 2019-2022 edition of the Bulletin of Selected Retail and

[90] The Bulletins of Prices can be found at: http://www.statistics.gov.lk/InflationAndPrices/StaticalInformation/Bulletins#gsc.tab=0

Table A2: Fertilizer Use and Intensity in Sri Lanka, by crop

	Fertilizer Type					
Crop	**Urea**	**TSP**	**MOP**	**Others**	**Total**	**Fertilizer Intensity (MT/ha)**
Chili pepper	3800	3800	3800	0	11400	0.300
Cinnamon	15114		7558	0	22672	0.662
Cloves	1359		2039	226	3624	0.454
Cocoa	178		222	44	444	0.243
Coconut	20147	2865	40295	15000	78307	0.173
Cowpea	1008	1550	1203	0	3761	0.243
Fruits	48359	35353	80067	0	163779	1.279
Green gram	1495	2300	1725	0	5520	0.240
Horse gram	3	6	2	0	11	0.073
Kurakkan	2345	605	908	0	3858	0.351
Maize	32300	14000	7000	0	53300	0.381
Nutmeg	484		727	80	1291	0.454
Onion	2926	1500	1126	0	5552	0.370
Paddy	247000	61000	74000	0	382000	0.294
Peanuts	1203	1850	1388	0	4441	0.240
Pepper	14562		10922	3640	29124	0.710
Potatoes	2310	1890	1750	0	5950	0.850
Rubber	6767		7151	4391	18309	0.146
Sesame	1595	1740	870	0	4205	0.290
Soybeans	.	.	.	.	.	0.067
Tea	96050		36864	59206	192120	0.959
Undu	1755	2700	2025	0	6480	0.240
Vegetables	22744	21297	14920	0	58961	0.640

Notes: for each crop, the table displays total fertilizer use in Sri Lanka (in MT) for four types of fertilizer: urea, triple superphosphate (TSP), muriate of potash (MOP), and other fertilizers. The last column shows fertilizer intensity (in MT/ha), which is computed by dividing the crop's total fertilizer use (in MT) by its cultivated area (in ha). The data comes from Sri Lanka's NFS, except for soybeans, whose fertilizer intensity data comes from the International Fertilizer Development Center and the International Fertilizer Association (specifically, from their 'Asia' estimates).

Producer Prices). We proxy the national prices of maize, onions, potato, rice, and tea with their (national) unit costs of production.[91] Since unit costs of production are not available for cinnamon and cloves, we use their (national) FOB export prices instead.

A.5 Fertilizer Use and Subsidies

To compute chemical fertilizer use in the production of various agricultural crops, we use data provided by Sri Lanka's NFS. Fertilizer requirements for 2022, displayed in Appendix Table A2, are reported as MT of fertilizer required per ha of cultivated land for each of 23 crops. As explained in Section 4.3, this fertilizer intensity metric is used in our main analysis to build a firm-level fertil-

[91] Due to our model's perfect competition assumption, a crop's unit cost of production equals its (farm-gate) producer price, which justifies using the former (which is available in the Economic and Social Statistics report) as a proxy for the latter. For tea and rice, unit costs are reported for the whole island, but for maize, onions, and potatoes they are reported for individual districts. We use data from Puttlam for onions and from Badulla for maize and potatoes, as these districts are major producers of their respective crops.

izer exposure measure (see equation 40). Sri Lanka's NFS data also includes estimates of fertilizer prices and of each crop's total cultivated area, which we combine with fertilizer requirements data to compute country-level fertilizer expenditure for each crop in 2022. Section 6.1.3 discusses how such fertilizer expenditures are used to estimate technological parameters.

NFS data on soybeans' fertilizer intensity was deemed unreliable by agriculture specialists, so we complement the NFS data with data from the International Fertilizer Development Center and the International Fertilizer Association. Specifically, we use the 'Asia' entries in these databases to impute the fertilizer intensity of soybeans that appears in Appendix Table A2. Appendix Figure A2 shows that this alternative fertilizer intensity metric has a close linear relationship with our original metric from the NFS (with the exception of soybeans, which are a clear outlier in the graph), thus also serving as a validation for the latter.

Figure A2: Fertilizer Intensity: Sri Lanka vs Asia

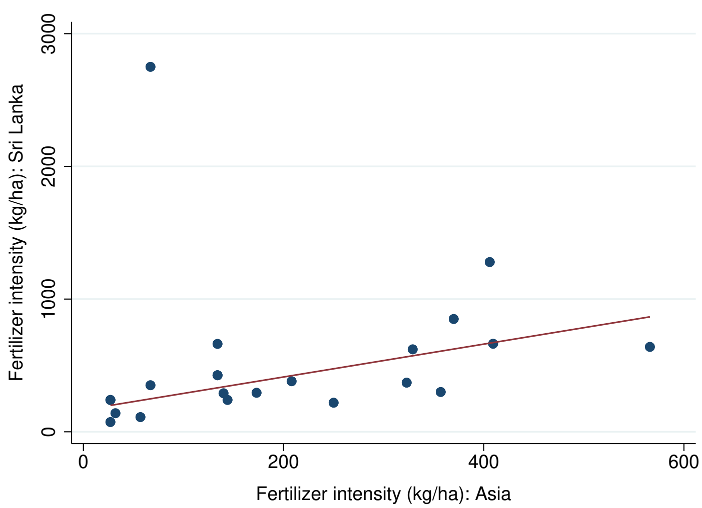

Notes: The figure shows the fertilizer intensity of 22 crops (see Appendix Table A2 for the full list), measured in kg/ha, from two alternative data sources: Sri Lanka's NFS (on the vertical axis) and the databases of the International Fertilizer Development Center and the International Fertilizer Association (on the horizontal axis). For the latter source, we use the estimates that correspond to Asia. The fertilizer intensity metrics add up the use of urea, TSP, and MOP. The outlier point on the top left corresponds to soybeans. The red line is the regression line when we exclude the outlier point from the sample, with a corresponding correlation coefficient of 0.677.

We use FAO data on the use of manure and chemical fertilizers in Sri Lankan agriculture in the period 2001-2021, namely the total nitrogen content (by weight) of manure applied to the soil and the total use (by weight) of 23 different types of chemical fertilizer. The data also includes the amount of manure that is lost through leaching or volatilization, allowing us to compute the net amount of manure that is applied to the soil.

We obtain data on fertilizer subsidies provided to farmers by Sri Lanka's government in the 2013-2022 period from the Cost of Cultivation Reports produced by the Department of Agriculture's Socio Economics & Planning Centre. For each crop, the fertilizer subsidy rate (in %) is based on the “nominal rate of protection at farm gate”, which is defined as the difference between actual and expected farm-gate fertilizer prices, with expected prices computed based on international prices and the cost of transportation from port to farm. Since there is no available subsidy data for cinnamon, cloves, and groundnuts, we impute their fertilizer subsidy rates with the average rate across potato, maize, tea, and onions.

A.6 Potential Crop Yields

To estimate the elasticity of substitution across crops of our model’s demand system, we use an IV approach that leverages variation in agricultural productivity across space due to exogenous geographic factors (such as climate and geology). To do so, we use data on district-level potential attainable yields for each of seven crops, which we obtain from the FAO-GAEZ project. The crops are: coconut, groundnut, maize, onion, rice, tea, and white potato. Unfortunately, cinnamon potential yields are not available in GAEZ. We extract potential yield data (in kg/ha) at the grid cell level and aggregate up to the district level. GAEZ offers different yield estimates depending on the irrigation system (rainfed or irrigated) and climate conditions (e.g. historical conditions, various climate change scenarios). To reflect the actual technologies commonly used in Sri Lankan agriculture, we choose the rainfed option for all crops except rice, which is irrigated. We choose the climate conditions of 1981-2010 without climate change. Section 6.1.1 describes how this data is used in estimation.

A.7 Labor Share in Value Added

To estimate labor and land coefficients in our model’s Cobb-Douglas agricultural production functions, we use data from Sri Lanka’s IO table for 2000, produced by the Institute of Policy Studies (Amarasinghe and Bandara, 2005). We extract information on employee compensation and value added (at factor cost prices) for a variety of agricultural sectors: “paddy”, “minor export crops”, “highland crops”, “tea growing” (separately for high, medium, and low elevations), and “potatoes”. Section 6.1.3 describes how this data is used in estimation.

A.8 Wages, Household Characteristics, Occupational Distribution, and Land Distribution

We use the 2016 and 2019 editions of the HIES for data on wages, employment, land ownership, land distribution, income, and consumption patterns. The HIES is a yearlong survey conducted in consecutive monthly rounds, with a nationally representative sample enumerated in each round to capture seasonal and regional pattern variation. The HIES provides household-level statistical weights. All averages, standard deviations, and counts we compute using HIES data are weighted using these weights.

We extract wage information from the 2019 HIES, which lists household members who "were paid employees during last four weeks/last calendar month". We compute each district's average wages as the (weighted) average of compensation in the main occupation across the district's workers (excluding workers with zero compensation), with compensation defined as the sum of "wages/salaries" and "tips, commissions, overtime pay etc" and measured in LKR. We compute average wages separately for estate and non-estate (mobile) workers. Estate workers are workers who live and work in Sri Lanka's tea plantations and for which changing jobs is nearly impossible in the short run, as their housing is tied to the tea estate they work in (Chandrabose, 2015). Average wages are multiplied by 12 so that they are expressed in annual terms.

We use information from the 2016 HIES to count the number of mobile workers, estate workers, and farmers in each district (making using of the statistical weights to obtain total counts). We define a "worker" as a person with a main occupation (in agriculture, manufacturing, or services) that paid her above-zero employment compensation in the previous month. We count the total number of such workers in each district and the fraction among these who are estate workers. We then subtract the number of estate workers from the total number of workers to obtain the number of mobile (i.e. non "tea-estate") workers. To compute the number of farmers, we first count the number of households in each district with positive landholdings. We only count land owned by the household (as opposed to leased, rented, etc.) with no housing built on top of it. We then estimate the number of farmers per district by multiplying the number of farming households by the average number of economically active agents per household. To estimate the average number of economically active agents per household, we divide the number of workers by the number of working households in the district, where a working household is defined as a household with at least one worker. The implicit assumption is that working households and farmer households are similar in terms of household size, structure, and labor force participation. Finally, we scale up the resulting numbers of workers and farmers by Sri Lanka's population growth between 2016 and 2019 (2.8%) to compute numbers for 2019, our baseline year for model estimation (see Section 6).

For parameter estimation in Section 6.1.1, we also use 2019 HIES data on household size (number of members, regardless of age), income, and expenditure patterns, in particular the fraction of income spent on food. Values and quantities of crops purchased by each household are available, allowing us to compute the prices faced by the household by dividing values by quantities. We also use information on "income by chance or *ad hoc* gains", including lottery gains and disaster/relief payments, for our IV estimation in Section 6.1.1.

A.9 Night Lights

For a model validation exercise in Section 7.7, we use data on night light intensity between January 2016 and December 2022 sourced from the Visible and Infrared Imaging Suite (VIIRS) Day Night Band (DNB) on board of satellites, prepared by the Colorado School of Mines. The key variable of interest is night light intensity at the year-month level for each 0.5 km x 0.5 km grid cell in Sri Lanka. For each month and district, we average the data across all grid cells that compose that district.[92]

A.10 Weather Variables

Since a wide array of weather variables have been shown to be important drivers of crop yields for many crops grown in Sri Lanka, it is important to control for weather in our analysis of the effects of the fertilizer import ban on yields (see Appendix C.4). Below, we describe our sources of weather data.

For rainfall data, we use the Climate Hazards Group InfraRed Precipitation with Station (CHIRPS) data, a quasi-global gridded rainfall time series data set based on a combination of high-resolution satellite imagery and in-situ station data covering latitudes between 50°S and 50°N, all longitudes, and all years from 1981 to the near-present.[93] For each agricultural growing season (Maha and Yala) and year between 1999 and 2022, we obtain three variables from CHIRPS: total rainfall, in millimeters (mm) per season; rainy days (defined as the number of days in the season receiving more than 2mm of rain); and dry spells (defined as the maximum number of consecutive days in the season receiving 1mm or less of rain). The Maha season is defined as starting on September 1 and ending on March 31. The Yala season is defined as starting on May 1 and ending on August 31. We aggregate the data up from the grid-cell level to the level of Sri Lankan districts.

For other weather variables, we use the ERA5-Land dataset, which provides high-resolution data on the evolution of the Earth's land surface between 1950 and the present.[94] For each growing season and year, we obtain four variables from ERA5-Land: average incoming shortwave solar radiation, in watts per square meter (W/m2); total vapor pressure deficit, in hectopascals; and average minimum and maximum temperatures, in degrees Celsius. Vapor pressure deficit (VPD) is the difference between saturated vapor pressure (temperature dependent) and the actual vapor pressure. Higher VPD values indicate higher air dryness and therefore often negatively correlate with yields. VPD is computed daily and summed over the season. Minimum and maximum temperatures are defined as the mean of daily minimum and maximum temperatures for the entire season. As with rainfall, we aggregate grid-cell level data on these four variables up to the district level.

[92] VIIRS data can be accessed at: https://eogdata.mines.edu/products/vnl/.
[93] CHIRPS data can be accessed at: https://www.chc.ucsb.edu/data/chirps.
[94] ERA5-Land data can be accessed at: https://www.ecmwf.int/en/era5-land.

A.11 Data on the Rest of the World

Since the rest of the world (RoW) is one of the regions in our quantitative model, we need data on its wages, production, cultivated area, prices, population, and employment so that we are able to bring the model to the data. Wages in the RoW are assumed to be equal to the world's GDP per capita (in current USD) in 2019 (our baseline year for model estimation), which was 11,330.5 USD according to the World Bank's World Development Indicators.[95]

We set the number of workers in the RoW to the size of the global labor force, which was 3.45 billion in 2019 according to data from the International Labour Organization (ILO). We then estimate the number of farmers in the RoW in two steps. First, we multiply the number of workers in the RoW by the agricultural share of global employment (26%, also from ILO data), thus obtaining the number of agricultural workers in the RoW.[96] Second, we multiply this number by the ratio of farmers to agricultural workers (0.29) which we compute using data from the Brazilian Agricultural Census as a benchmark. To calculate this ratio, we use "number of establishments directed by owner or her business partner" as a proxy for the number of farmers; and "number of occupied people in agricultural establishments" as a proxy for the number of agricultural workers.[97] The resulting estimate for the number of farmers in the RoW is approximately 260 million.

Our data on agricultural production, cultivated area, and producer prices in the RoW comes from the FAO.[98] We obtain data for each country's production quantities (in MT), harvested area (in ha), and producer price (in USD/MT) of seven crops in 2019, then sum across countries to obtain production and cultivated area in the RoW. The crops we use are listed in the FAO data with the following names: "Cinnamon and cinnamon-tree flowers, raw", "Cloves (whole stems), raw", "Groundnuts, excluding shelled", "Maize (corn)", "Potatoes", "Rice", and "Onions and shallots, green". Each crop's producer price in the RoW is set to the median price across all available countries. All variables that are measured in MT are converted to kg.

Finally, it is sometimes necessary in our analysis to convert currency denominations between the USD and the LKR. We use the exchange rate $E_{2019} = 176$ LKR/USD for 2019 and $E_{2022} = 360$ LKR/USD for 2022, both of which were obtained from Google Finance.[99]

[95] The data portal can be accessed at: https://databank.worldbank.org/indicator/NY.GDP.PCAP.CD/1ff4a498/Popular-Indicators

[96] ILO data is available on the World Bank's data portal at: https://data.worldbank.org/indicator/SL.TLF.TOTL.IN and https://data.worldbank.org/indicator/SL.AGR.EMPL.ZS

[97] These aggregated statistics from Brazil's 2006 Agricultural Census are available in Tables 2724 and 265 of the online portal SIDRA, which can be accessed at: https://sidra.ibge.gov.br/

[98] This data can be accessed through the FAOSTAT online portal at: https://www.fao.org/faostat/en/data/QCL and https://www.fao.org/faostat/en/data/PP

[99] These two exchange rates are from June 28 2019 and June 24 2022, respectively.

A.12 Descriptive Statistics

Appendix Table A3 shows the share of fertilizer imports that are organic, while Appendix Figure A3 shows the use of organic fertilizer in agriculture in Sri Lanka. Appendix Table A4 shows the share of imports subject to import bans from 2017 to 2021.

Table A3: Fertilizer Imports, Total versus Organic

Period	Total Fertilizer Imports (in MM USD)	% Organic
2016	128	0.078%
2017	98	0.128%
2018	249	0.148%
2019	211	0.072%
2020	244	0.111%
2021 (no-ban months)	23	0.639%
2021 (ban months)	11	3.763%

Notes: for each year, the table shows total fertilizer imports by Sri Lanka in millions (MM) of USD and the percentage of these imports consisting of organic fertilizer. The data comes from the S&P Global Market Intelligence's Panjiva data platform. A fertilizer is any HS8 product listed in Chapter 31 of the HS system. Organic fertilizers are classified as "animal or vegetable" (as opposed to "mineral or chemical"), and are listed under heading 3101 of the HS system. We break 2021 into two subperiods: the "ban months" (May-November) when the fertilizer import ban was in place, and the "no-ban months" (January-April and December) when the ban was not in place.

Figure A3: Use of Manure in Agriculture

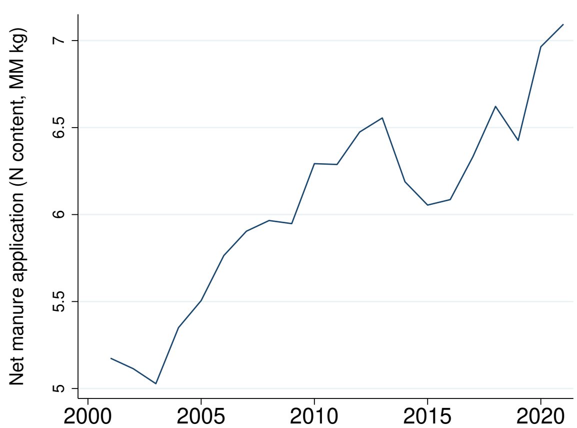

(a) Use of Manure

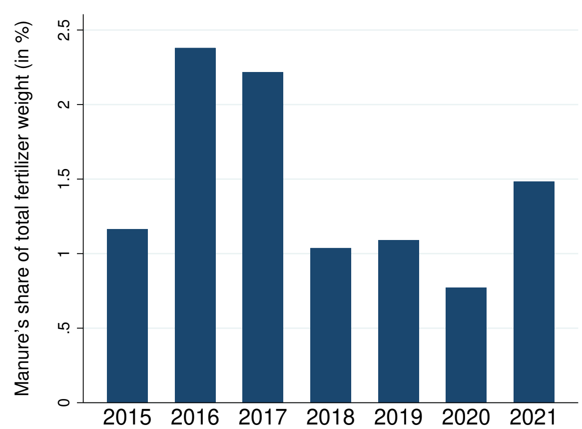

(b) Manure's Share (by weight)

Notes: for each year between 2001 and 2021, Figure A3a shows the total nitrogen content (measured in MM kg) of the manure that is applied to the soil as an organic fertilizer in Sri Lanka, net of the nitrogen that is lost through leaching or volatilization. Appendix Figure A3b shows the same quantity as in Appendix Figure A3a, but as a share of the total weight of all fertilizer used in Sri Lankan agriculture, which includes both manure and (23 types of) chemical fertilizer products. The data is from the FAO.

Table A4: Imports of Products Subject to Import Bans

Year	Import share	Imports (in MM USD)
2017	17.18%	2,544
2018	18.64%	2,873
2019	15.69%	2,088
2020	12.41%	1,436
2021	10.19%	1,444

Notes: For each year, the table displays the total import value of the set of HS8 products whose imports were banned at some point starting in 2020. This variable is shown both in absolute terms (in MM USD) and as a % of total imports. The data comes from Sri Lanka's Extraordinary Gazettes and the S&P Global Market Intelligence's Panjiva data platform.

Appendix Table A5 displays pre-ban and post-ban summary statistics for some of our key variables of interest. Some of the ban's effects are apparent even in these basic statistics: comparing the pre-ban to the post-ban period, we see that the average value of fertilizer imports plummeted from 2 million USD to only 491 thousand USD per importer firm per year. Similarly, the average value of agricultural exports fell from 1.62 million USD to 1.53 million USD per agricultural exporter firm per year.

Panel B of Appendix Table A5 focuses on the evolution of crop yields over time, as estimated by our remote sensing methodology. Again, we see large decreases in crop yields between the pre-ban (2021) and post-ban (2022) periods: the average rice yield across DSs falls from more than 4,000 kg/ha to less than 3,000 kg/ha, while average tea yields fall from 1,644 kg/ha to 1,404 kg/ha.

B Details on Estimation of Stylized Facts

B.1 Stylized Fact 1

In Section 4.1, we estimate the fertilizer import ban's effects on fertilizer imports through state-of-the-art dynamic DID regressions following De Chaisemartin and d'Haultfoeuille (forthcoming). The regressions are estimated on a panel of HS8 products at the monthly level, with the treatment group consisting of fertilizers who had their imports banned from May 2021 onward, and the control group consisting of non-banned products. We exclude from the analysis products whose imports were banned at any point prior to May 2021. The estimating equation is:

$$y_{ct} = \sum_{\tau \neq -1} \beta_\tau \times ban_{ct}^{\tau} + \omega_t + \omega_c + \epsilon_{ct} \tag{39}$$

where c indexes HS8 products and t indexes time periods (months), y_{ct} is the outcome variable, ω_t and ω_c are time and product fixed effects, respectively, and ϵ_{ct} is an error term. Our main variable of interest is ban_{ct}^{τ}, a dummy variable indicating whether an import ban on product c was first imposed

Table A5: Summary Statistics

Panel A: International Trade		
Firm-level variables	Pre-Ban (1)	Post-Ban (2)
Fertilizer Imports (USD)	2,036,434 (7,267,773) [N=106]	491,025.4 (1,378,846) [N=106]
Agricultural Exports (USD)	1,619,553 (6,206,280) [N=1,565]	1,531,465 (6,030,005) [N=1,562]
Panel B: Crop Yields		
Divisional Secretariat-level variables	Pre-Ban (1)	Post-Ban (2)
Rice Yields (kg/ha)	4,236.585 (1,013.066) [N=323]	2,881.95 (680.9113) [N=323]
Tea Yields (kg/ha)	1,644.278 (1,165.624) [N=134]	1,403.798 (971.3812) [N=134]

Notes: The table shows summary statistics for key variables of interest, separately for pre-ban and post-ban periods. In addition to each variable's average, the table displays its standard deviation (in parentheses) and sample size N (in square brackets). The unit of observation is the firm in Panel A and the DS in Panel B. The data comes from the S&P Global Market Intelligence's Panjiva data platform in Panel A and from our remote sensing estimates in Panel B. For each variable and period, the statistics are computed using all available units for which that variable was positive in that period. In Panel A, we define fertilizers as products listed in Chapter 31 of the HS system, and agricultural goods as products listed between Chapters 06 and 24. Imports and exports are measured in US dollars (USD), and crop yields are measured in kg/ha. In the case of tea yields, we use a DS-level confidence index that reflects the precision of the remote sensing estimates as a weight to compute weighted averages and standard deviations. Our definitions of pre-ban and post-ban periods differ across variables: for fertilizer imports, the pre-ban period is May 2020-April 2021 and the post-ban period is May 2021-April 2022; for agricultural exports, the pre-ban period is November 2020-October 2021 and the post-ban period is November 2021-October 2022; in Panel B, the pre-ban period is 2021 and the post-ban period is 2022.

τ periods before period t. Note that ban_{ct}^{τ} is always equal to zero for products whose imports were never banned.

We estimate the evolution of the ban's effects on two alternative outcome variables (y_{ct}): an "extensive-margin" dummy indicating non-zero imports and an "intensive-margin" variable consisting of the log of import values. To define our extensive-margin dummy variable, we expand our original HS8 product-month level import data by adding an entry with zero imports for each month (within the January 2017-October 2022 period) when imports of a product are not reported.

The sample period used in the estimation ranges from January 2019 to September 2022. For presentational purposes, the figures with our import results (Figure 1 and Appendix Figures A5, A6, and A7) only show coefficients for the four months prior and twelve months subsequent to the ban, even though the sample period used in the estimation is longer.

B.2 Stylized Fact 3

In Section 4.3, we estimate the fertilizer import ban's downstream effects on agricultural exports. We follow a firm-level approach, focusing on firms who exported agricultural products in the pre-ban 2017-2019 period. For each firm f, we define the following exposure measure U_f:

$$U_f = \sum_{v=1}^{23} FIC_v \times \left(\frac{X_{fv}^{2017-2019}}{\sum_{v'} X_{fv'}^{2017-2019}} \right), \tag{40}$$

where f indexes firms and v crops, FIC_v is fertilizer intensity (in kg/ha) of crop v, and $X_{fv}^{2017-2019}$ is total exports of crop v by firm f in 2017-2019. The summation in equation (40) only considers the 23 crops for which we have fertilizer intensity data (see Appendix A.5). This exposure measure is the weighted average of fertilizer intensity across all crops that a firm exports (with weights given by each crop's share in the firm's export portfolio).

We then use a firm-quarter panel and the dynamic DID estimator from De Chaisemartin and d'Haultfoeuille (forthcoming) once again to estimate the ban's impact on firm exports. The treatment group consists of firms above the upper quartile of fertilize exposure variable U_f, with the remaining firms in the sample serving as the control group. The estimating equation is as follows:

$$\ln(X_{ft}) = \sum_{\tau \neq -1} \gamma_\tau \times \mathbb{1}\{t = T_0 + \tau\} \times \mathbb{1}\{U_f > U_{p75}\} + \Omega_t + \Omega_f + \mu_{ft} \tag{41}$$

where f indexes firms, t indexes quarters, X_{ft} is agricultural exports, Ω_t and Ω_f are quarter and firm fixed effects, respectively, T_0 is the quarter when the fertilizer import ban was introduced (i.e. Q2/2021), and μ_{ft} is an error term.

The sample period used in the estimation ranges from January 2019 to September 2022. For presentational purposes, the figures with our export results (Figure 4 and Appendix Figure A8) only show coefficients for the four quarters prior and twelve quarters subsequent to the ban, even though the sample period used in the estimation is longer.

C Additional Empirical Findings

C.1 Robustness of Stylized Facts

Figure A4: Dynamic Ban Effects on Imports: Fertilizer Only

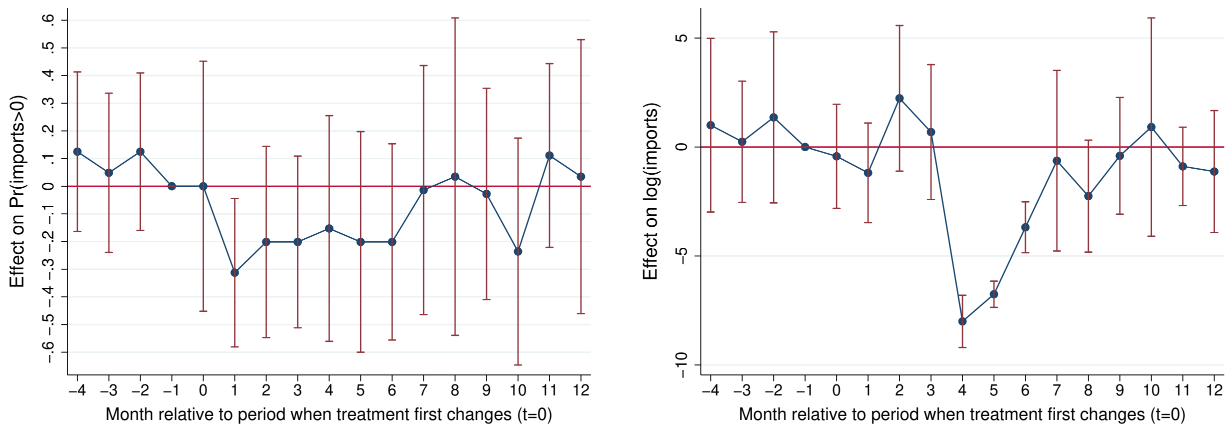

(a) Extensive Margin, fertilizer only

(b) Intensive Margin, fertilizer only

Notes: The figures show De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of import bans on the extensive and intensive margins of imports, relative to the last month before the ban takes effect, using an HS8 product-month level panel. The sample covers the January 2019-September 2022 period (but only coefficients in January 2021-May 2022 are reported) and includes only fertilizer HS8 products. The treatment variable is a set of dummies ban_{ct}^{τ} indicating whether product c was first banned τ months before month t, and the not-yet-treated products serve as the control group. The y-axis shows the effects of fertilizer import bans on the probability of nonzero imports (panel A4a) or on the log value of imports (panel A4b), where negative values correspond to a decrease in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.846 (panel A4a) and 0.925 (panel A4b).

Figure A5: Dynamic Ban Effects on Import Quantities

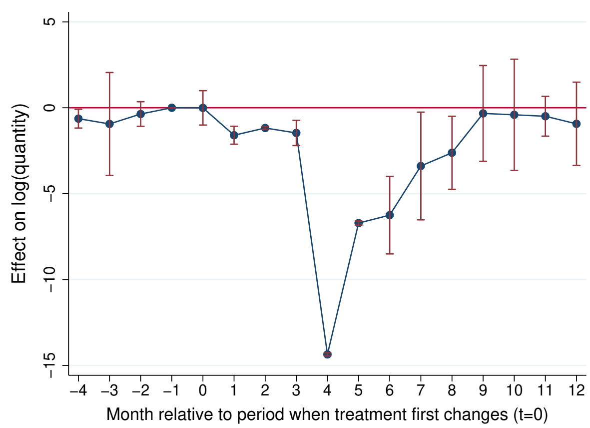

(a) Quantity Effect

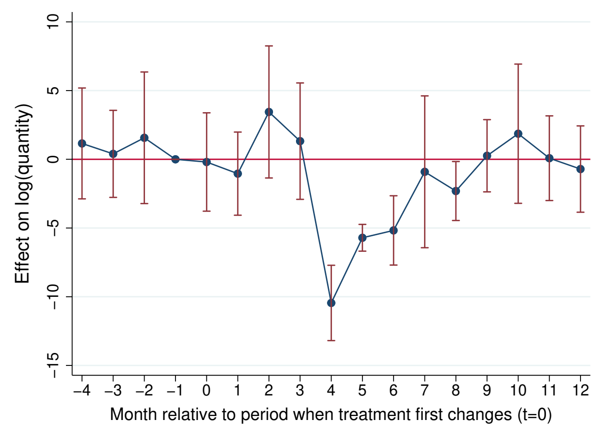

(b) Quantity Effect, fertilizer only

Notes: The charts show De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of import bans on import quantities, relative to the last month before the ban takes effect, using an HS8 product-month level panel. The sample covers the January 2019-September 2022 period (but only coefficients in January 2021-May 2022 are reported) and excludes non-fertilizer HS8 products whose imports were ever banned. In panel A5b, the sample also excludes all non-fertilizer products, even the ones who were never banned. The treatment variable is a set of dummies ban_{ct}^{τ} indicating whether product c was first banned τ months before month t, and the not-yet-treated products serve as the control group. The y-axis shows the effects of import bans on (the log of) the net weight of imports (measured in kg), where negative values correspond to quantity decreases. The p-values of the joint significance tests for the placebo estimators are 0.005 (panel A5a) and 0.93 (panel A5b).

Figure A6: Dynamic Ban Effects: Deseasonalized Imports

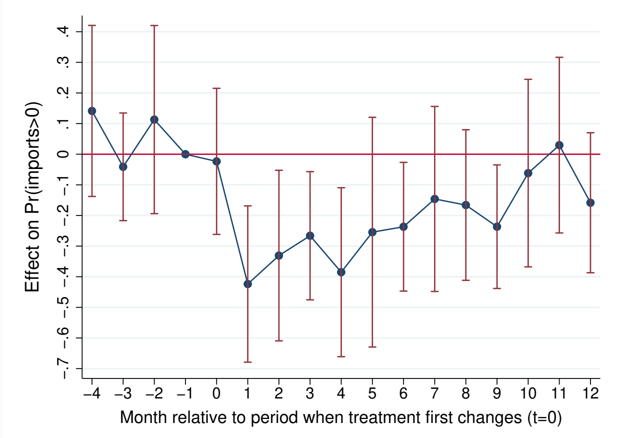

(a) Extensive Margin

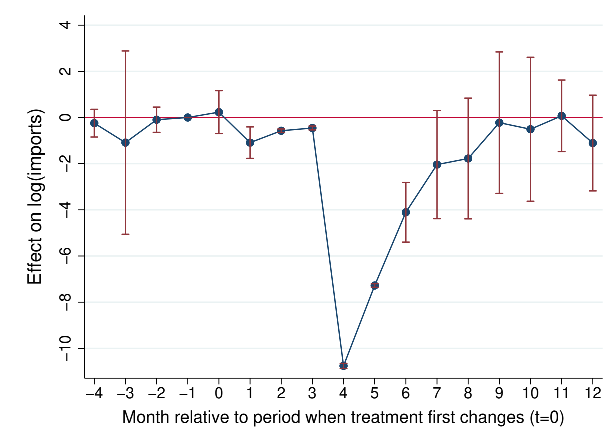

(b) Intensive Margin

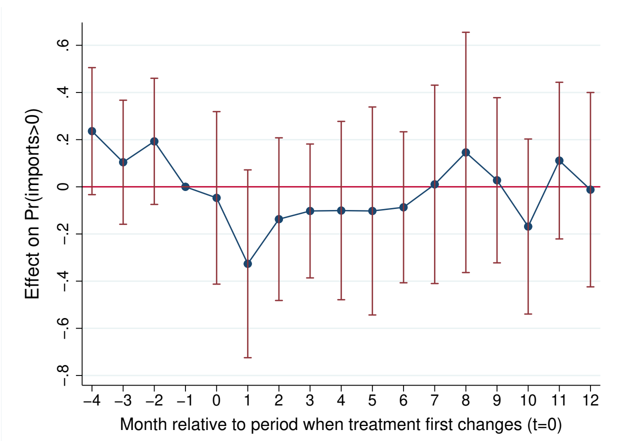

(c) Extensive Margin, fertilizer only

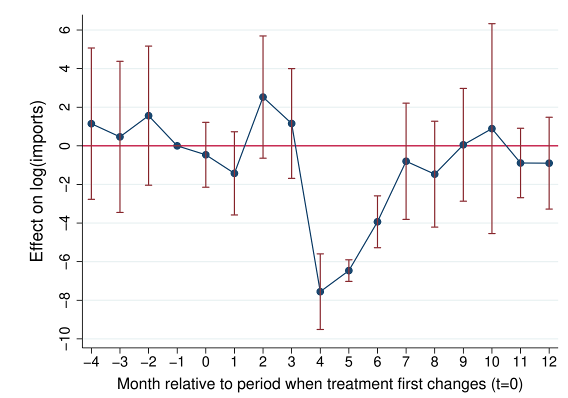

(d) Intensive Margin, fertilizer only

Notes: The charts show De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of import bans on both the extensive and intensive margins of (deseasonalized) imports, relative to the last month before the ban takes effect, using an HS8 product-month level panel. The sample covers the January 2019-September 2022 period (but only coefficients in January 2021-May 2022 are reported) and excludes non-fertilizer HS8 products whose imports were ever banned. In panels A6c and A6d, the sample also excludes all non-fertilizer products, even the ones who were never banned. For each month of the year, we compute the (log of the) product's historical average import value and the historical average probability of nonzero imports for that month in 2017-2019; then, for every product-month we subtract the relevant historical average from log import value and from the indicator of nonzero imports, thus obtaining deseasonalized versions of these two outcome variables. The treatment variable is a set of dummies ban_{ct}^{τ} indicating whether fertilizer product c was first banned τ months before month t, and the not-yet-treated products serve as the control group. The y-axis shows the effects of import bans on the (deseasonalized) probability of nonzero imports (panels A6a and A6c) or on the (deseasonalized) log value of imports (panels A6b and A6d), where negative values correspond to decreases in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.71 (panel A6a), 0.68 (panel A6b), 0.35 (panel A6c), and 0.85 (panel A6d).

Figure A7: Event Study: Imports and the Arrival of Covid-19

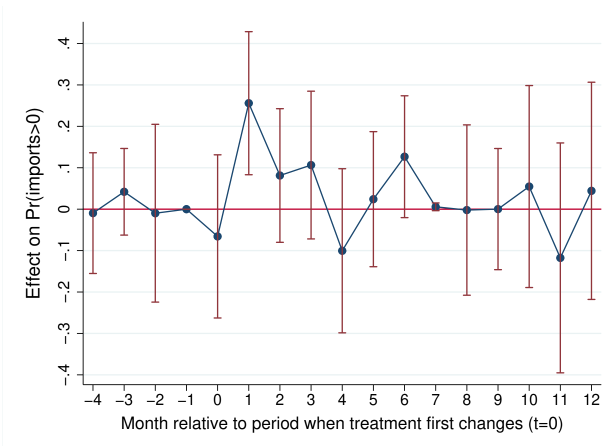

(a) Extensive Margin

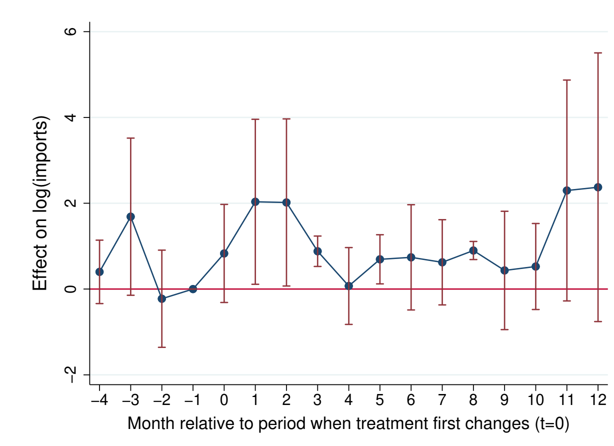

(b) Intensive Margin

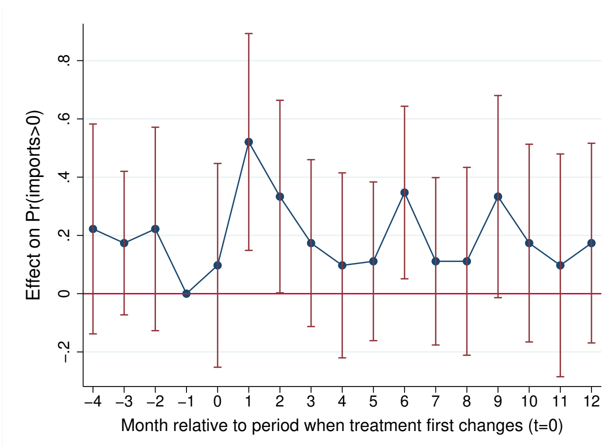

(c) Extensive Margin, fertilizer only

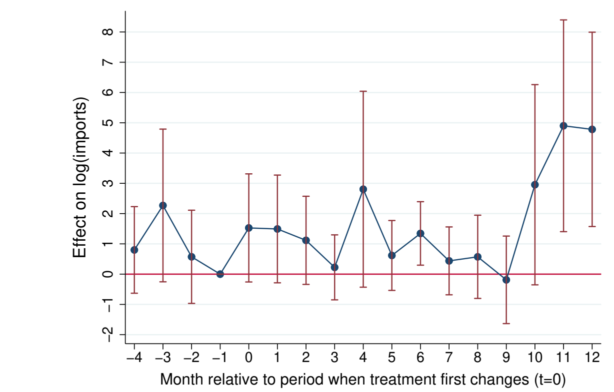

(d) Intensive Margin, fertilizer only

Notes: The charts show De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the differential effects of the Covid-19 pandemic on both the extensive and intensive margins of fertilizer imports, relative to the last month before the start of Covid-19, using an HS8 product-month level panel. The sample covers the January 2019-April 2021 period (but only coefficients in November 2019-March 2021 are reported) and excludes non-fertilizer HS8 products whose imports were ever banned. In panels A7c and A7d, the sample also excludes all non-fertilizer products, even the ones who were never banned. The treatment variable is a dummy $covid_{ct}$ which equals one if period t is later than February 2020 **and** product c was banned in some period starting in May 2021. Never-banned products serve as the control group. The y-axis shows the differential effects of Covid-19 on the probability of nonzero imports (panels A7a and A7c) or on the log value of imports (panels A7b and A7d), where negative values correspond to decreases in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.63 (panel A7a), 0.27 (panel A7b), 0.47 (panel A7c), and 0.36 (panel A7d).

Figure A8: Dynamic Ban Effects on Firms: Deseasonalized Agricultural Exports (High vs Low Fertilizer Intensity)

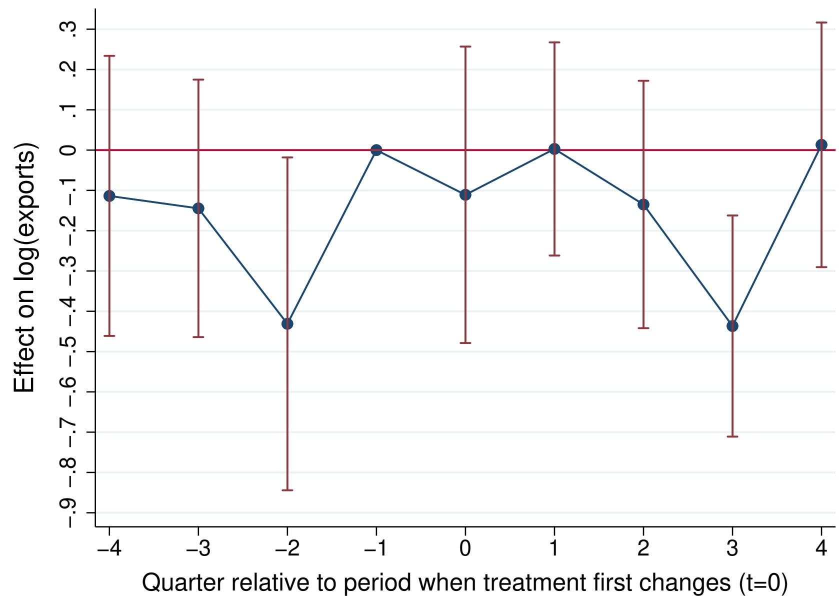

Notes: The chart shows De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effects of fertilizer import bans on the (deseasonalized) value of firms' agricultural exports, relative to the last quarter before the ban takes effect, using a firm-level panel. The sample covers the Q1/2019-Q3/2022 period (but only coefficients in Q2/2020-Q2/2022 are reported) and includes all firms who exported at least one of 23 crops for which we have fertilizer intensity data in 2017-2019. We compute each firm's fertilizer exposure as a weighted average of crops' fertilizer intensities with weights given by each crop's share in the firm's 2017-2019 exports. The treatment group consists of firms whose fertilizer exposure is above the sample's 75th percentile. Firms below the 75th percentile serve as the control group. The treatment period is defined as the set of all quarters when any fertilizer import bans were in place. For each quarter of the year, we compute the (log of the) firm's historical average export value for that quarter in 2017-2019; then, for every firm-quarter, we subtract that historical average from the (log) export value, thus obtaining a deseasonalized version of the outcome variable. The y-axis shows the effects of fertilizer import bans on the log value of agricultural exports, where negative values correspond to a decrease in exports. The p-value of the joint significance test for the placebo estimators is 0.196.

C.2 Food's Expenditure Share

The fraction of the consumer's income that is spent on agricultural goods is an important variable in our analysis. To illustrate how much this variable can vary, we use data from the 2019 edition of the HIES to display the empirical distribution of food expenditure shares across Sri Lankan households in Appendix Figure A9, which shows a wide dispersion.

Moreover, as detailed in Section 5.3.2, we model consumer preferences as non-homothetic. The main motivation for this assumption is the marked negative relationship observed in the data between income and the food expenditure share. This can be seen in Appendix Figure A10, which plots household income against the food share of household expenditures, showing a strong and clear negative correlation between the two variables. The two variables plotted in Appendix Figure A10 are "residualized" versions, i.e. deviations from the (district-level) average.

Figure A9: Distribution of Food Expenditure Shares

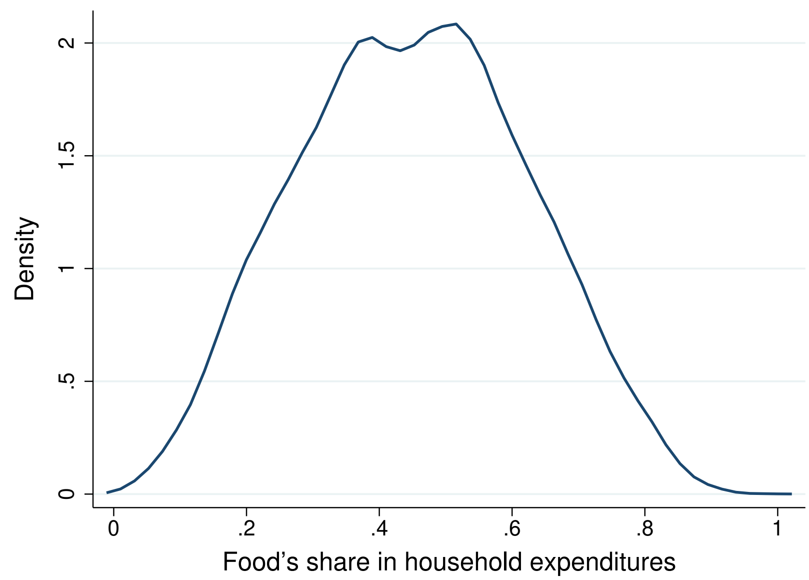

Notes: The figure shows kernel density estimates for the share of total household expenditure that is spent on food. A unit of observation is a household. The data comes from the 2019 edition of the HIES. The density is estimated using an Epanechnikov kernel with optimal bandwidth (0.0213) and household-specific statistical weights provided by the HIES.

Figure A10: Association Between Food Share and Income

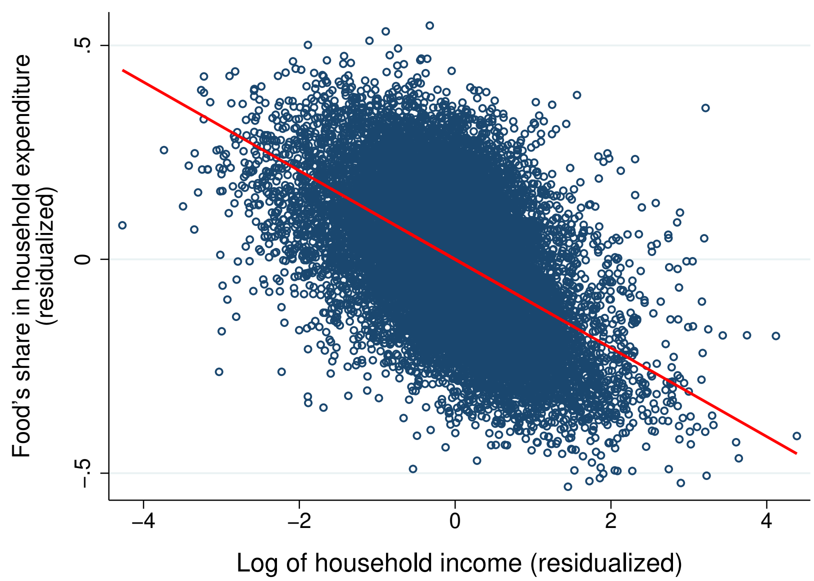

Notes: The figure plots a residualized version of the log of household income against a residualized version of the share of total household expenditure that is spent on food. A unit of observation is a household. The data comes from the 2019 edition of the HIES. The residualized version of each variable is obtained by extracting the residual from a regression of that variable on a set of district dummy variables, with the regression estimated by Weighted Least Squares (WLS) using the statistical weights provided by the HIES. A red regression line is overlaid on the graph representing the regression (also estimated by WLS) of residualized food share on residualized log household income.

Figure A11: Real Crop Prices over Time

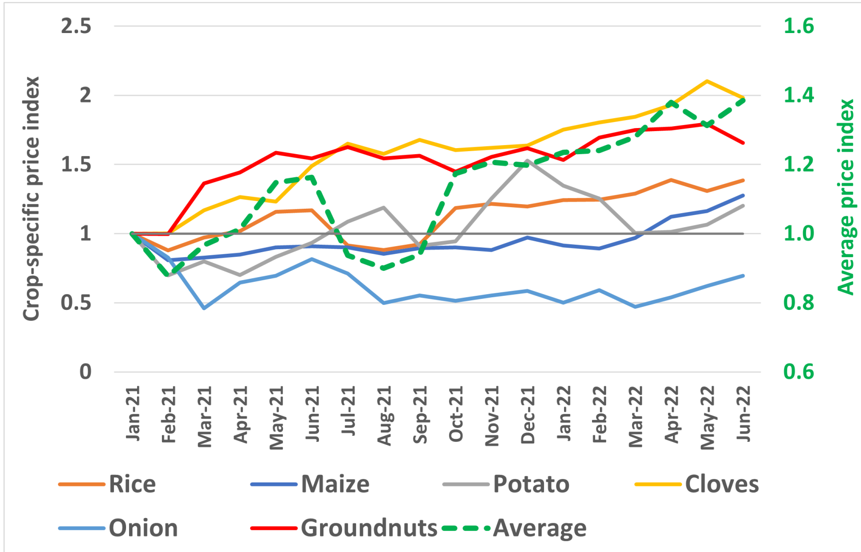

Notes: for each month between January 2021 and June 2022, the figure shows the real price index of six crops in Sri Lanka, as well as a cross-crop (weighted) average price index, with weights given by each crop's national revenue in 2019, which is computed from data on crop production and producer prices (see Appendices A.3 and A.4). Cumulative nominal price increases are deflated by cumulative inflation (computed using the National Consumer Price Index) to obtain real price indices, which are then renormalized to assume a value of one in January 2021. All price data is from Bulletins of Selected Retail and Producer Prices (see Appendix A.4).

C.3 Land Distribution

Land ownership in Sri Lanka is unequally distributed, with landholding sizes varying considerably across landowning households within any of the country's districts, as seen in Appendix Figure A12a, which shows the distribution of (the log of) landholding size. Note that this variation is not solely driven by regional differences in average landholding size, as dispersion remains high after controlling for district in Appendix Figure A12b.

The importance of land inequality in the data motivates us to introduce inequality explicitly as a feature of our model in Section 5.3.1 by parametrizing the land distribution in each district n as a log-normal distribution with parameters μ_n and σ^2_{Ln}. It then becomes necessary to estimate these parameters as part of the process of bringing the model to the data (see Section 6.1.4), which we do using 2016 HIES data.

Specifically, for each district n in the data, we start by computing the average of the log of landholding sizes across households in that district, denoted $\bar{lL}_n$:

$$\bar{lL}_n = \left(\frac{1}{\sum_{h' \in n} w_{h'}}\right) \sum_{h \in n} w_h \ln(L^h)$$

where h indexes households in the sample, w_h is the household's statistical weight, and L^h is the area (in ha) of the land it owns. Only households with positive landholdings are included in the

Figure A12: Distribution of Landholding Sizes

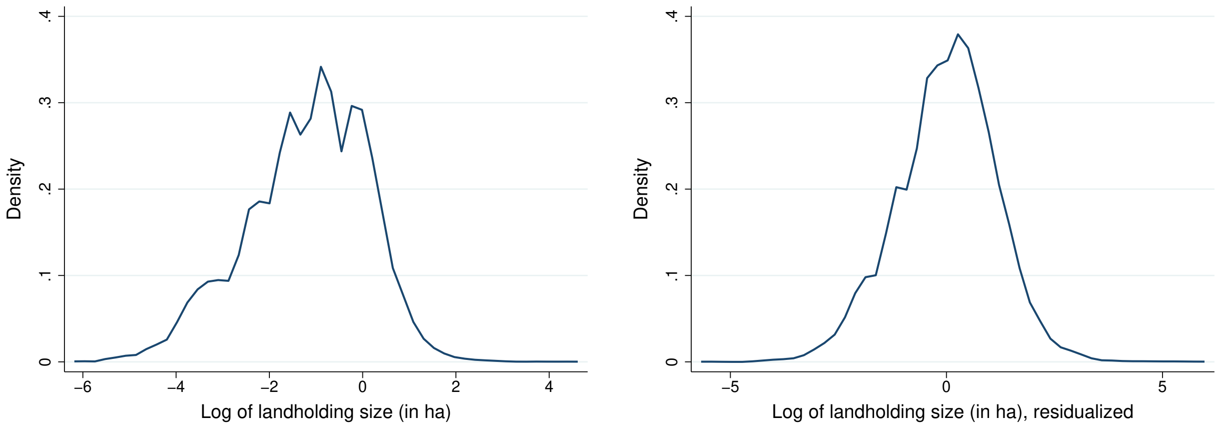

(a) Unadjusted (b) Adjusted for Region

Notes: The figures show kernel density estimates for (the log of) the size of households' landholdings (measured in ha). A unit of observation is a household. The data comes from the 2016 edition of the HIES. Panel A12a shows unadjusted estimates. Panel A12b shows estimates for a residualized version of the log landholding size variable which is obtained by extracting the residual from a regression of that variable on a set of district dummy variables, with the regression estimated by Weighted Least Squares (WLS) using the statistical weights provided by the HIES. Densities are estimated using Epanechnikov kernels with optimal bandwidth (0.2021 in Panel A12a and 0.1647 in Panel A12b) and statistical weights.

sample. We then estimate the dispersion parameter σ^2_{Ln} for district n as:

$$\hat{\sigma}^2_{Ln} = \frac{|\mathbb{S}_n|}{|\mathbb{S}_n| - 1}\left(\frac{1}{\sum_{h'\in n} w_{h'}}\right)\sum_{h\in n} w_h\big(\ln(L^h) - \bar{lL_n}\big)^2$$

where $|\mathbb{S}_n| \equiv \sum_{h'\in n} 1$ is the number of surveyed households in district n that own land.

Given estimates $\hat{\sigma}^2_{Ln}$ of the dispersion parameters, we can then estimate the scale parameters μ_n so as to match the average landholding size in each district. Specifically, the properties of the log-normal distribution imply that the average landholding size in district n is $\exp(\mu_n + \sigma^2_{Ln}/2)$, so we set the empirical analogue of this expression equal to the average landholding size observed in the data and rearrange terms to obtain our estimator of μ_n:

$$\hat{\mu}_n = \ln\left(\frac{L_n}{N^F_n}\right) - \frac{\hat{\sigma}^2_{Ln}}{2}$$

where L_n is total arable land (in ha) and N^F_n is the number of farmers in district n.

C.4 Evolution of Rice and Tea Agriculture over Time

We combine cutting-edge remote sensing data (described in Appendix A.3.2) with detailed weather data (described in Appendix A.10) to show the evolution of yields of rice and tea (and, for rice, of cultivated area) over time.

For rice, we estimate the following regression for each of the two growing seasons (Maha and Yala):

$$A^s_{dy} = \theta^s_y + \omega^s_d + \mathbf{B^s_W W_{dy}} + \epsilon^s_{dy} \quad (42)$$

where s indexes seasons, d indexes DSs, y indexes years, A^s_{dy} is the cultivated rice area in district d during season s of year y, θ^s_y (resp. ω^s_d) is a set of year (resp. DS) fixed effects, $\mathbf{W^s_{dy}}$ is a vector of weather variables, and ϵ^s_{dy} is an error term. We estimate equation (42) by Ordinary Least Squares with standard errors clustered at the DS level. The vector of weather variables includes: rainfall, number of rainy days, dry spells, solar radiation, VPD, average minimum and maximum temperatures. We then repeat the process substituting cultivated area for yields as the dependent variable. The results are shown in Appendix Figure A13.

Figure A13: Rice Yields and Cultivated Area, by year

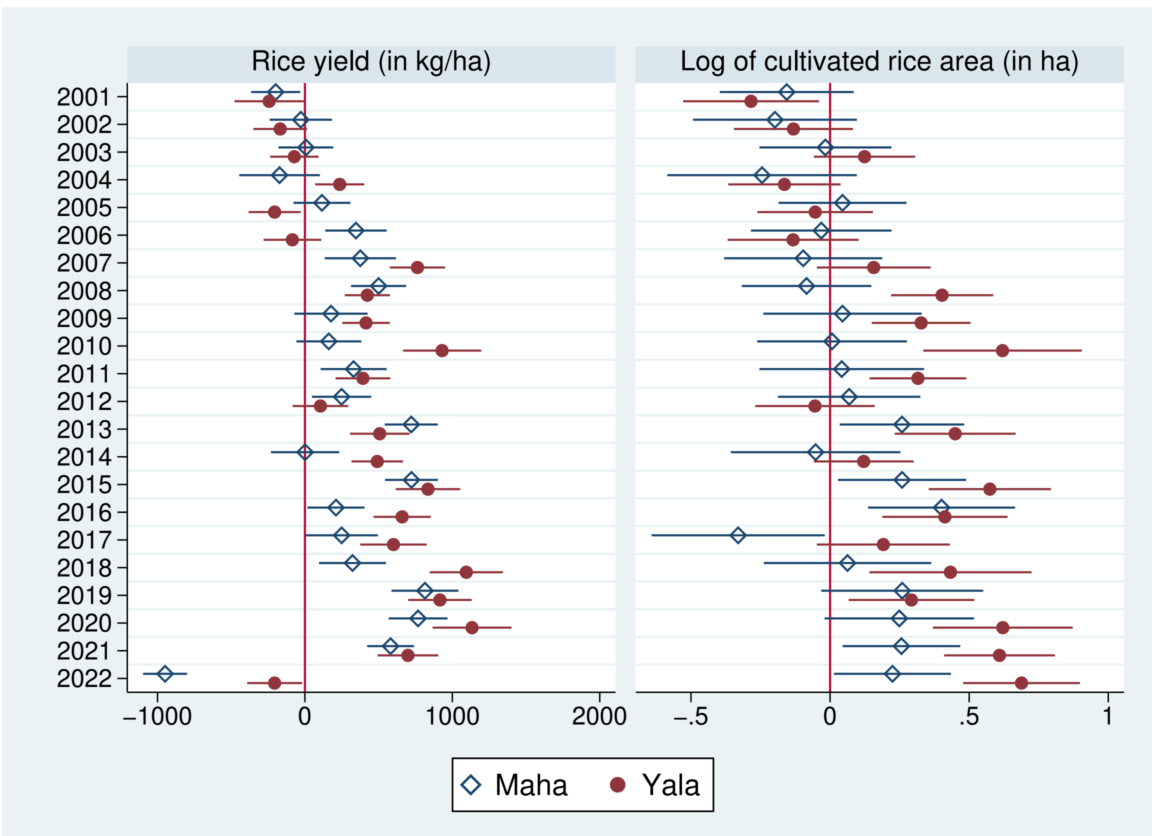

Notes: The figure plots coefficients and 95% confidence intervals from four panel regressions at the DS-year level in which we regress rice yields (left panel) and the log of cultivated rice area (right panel) on a set of year dummies and DS fixed effects. Regressions are run separately for the Maha season (diamonds) and the Yala season (circles). Yields are measured in kg/ha, and area is measured in ha. Standard errors are clustered at the DS level. All regressions control for the following season- and DS-specific weather variables: rainfall (in mm/season), number of rainy days, dry spells, solar radiation (in W/m2), VPD (in hectopascals), and average minimum and maximum temperatures (in degrees Celsius).

Since we control for weather and DS fixed effects, the year coefficients displayed in Appendix

Figure A13 can be interpreted as the portion of the year's deviation from the historical average that cannot be explained by abnormal weather. Rice yields in the 2022 Maha season (the first growing season after the fertilizer import ban) were 948 kg/ha below the baseline year (2000), with the null hypothesis of zero deviation from baseline being strongly rejected at conventional significance levels. A 948 kg/ha decrease in rice yields corresponds to a 26.9% decrease from the 2000 baseline, or to a 24.4% decrease from the the average yield of the 2000-2021 period. In short, the first rice growing season after the ban was, by far, the least productive in at least 20 years.

In contrast, in the cultivated area regression, the 2022 Maha coefficient is positive. Therefore, we have no evidence of a reduction in cultivated rice area after the ban. This is consistent with our model's assumption that land cannot be reallocated across crops, which is expected given the short time horizon under study. Moreover, the combination of this finding with the dramatic yield declines in the same season supports our hypothesis that the fertilizer ban had a major negative impact on Sri Lankan rice agriculture.

For tea, we estimate a similar regression as for rice, but pooling Maha and Yala yields together into a single observation per district-year:

$$Y_{dy} = \theta_y + \omega_d + \mathbf{B_W^{yala} W_{dy}^{yala}} + \mathbf{B_W^{maha} W_{dy}^{maha}} + \epsilon_{dy} \tag{43}$$

where d indexes DSs, y indexes years, Y_{dy} is the tea yield in district d in year y, θ_y (resp. ω_d) is a set of year (resp. DS) fixed effects, $\mathbf{W_{dy}^{yala}}$ (resp. $\mathbf{W_{dy}^{maha}}$) is a vector of weather variables in DS d during the Yala (resp. Maha) season of year y, and ϵ_{dy} is an error term. Equation (43) is estimated by Weighted Least Squares (WLS) with standard errors clustered at the DS level. Since the precision of the remote sensing tea yield metric varies across regions, the regression is weighted by a DS-specific confidence index that reflects the degree of precision of the metric in each DS. Regression results are displayed in Appendix Figure A14.

As in the rice regression, the year coefficients in Appendix Figure A14 can be interpreted as the yield deviations from the historical baseline (i.e. year 2014) that cannot be explained by abnormal weather. Post-ban tea yields in 2022 were 428 kg/ha (or 26.8%) below the baseline year (2014), a difference that is strongly statistically significant and that corresponds to a 28.7% drop relative to the average yield of the 2014-2021 period. Therefore, controlling for weather fluctuations reveals that the first post-ban agricultural season had remarkably low yields not only for rice but also for tea.

C.5 Relationship Between Land Inequality and Agricultural Expenditures

Equation (14) in our model predicts a negative relationship at the regional level between aggregate agricultural expenditure (X_n^A) and land inequality (σ_{Ln}^2) when holding aggregate income (E_n) and

Figure A14: Tea Yields, by year

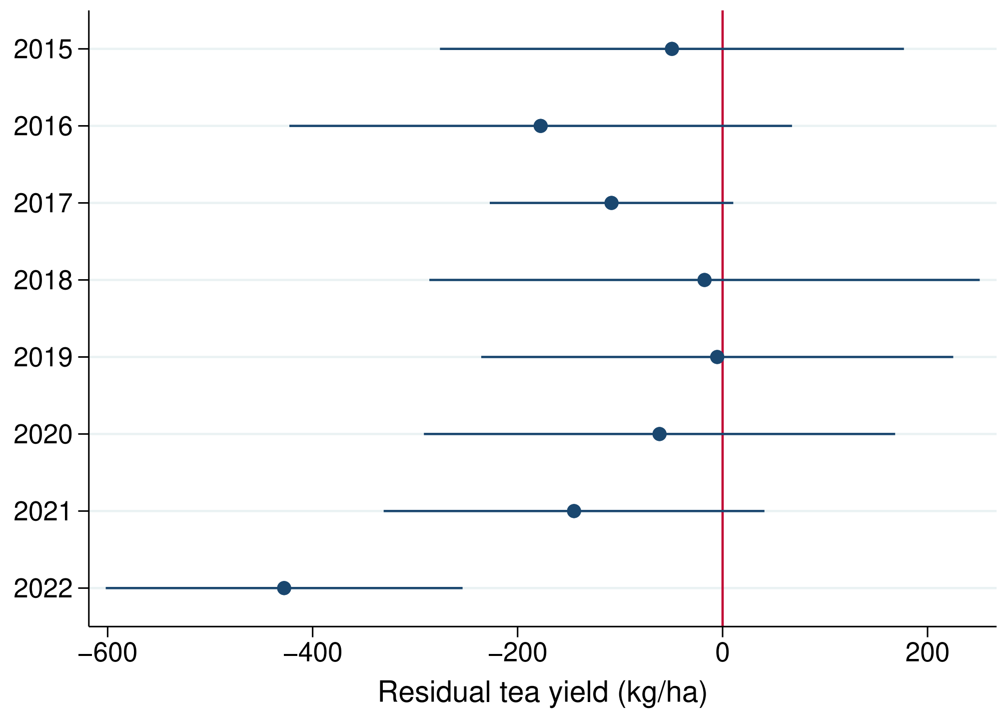

Notes: The figure plots coefficients and 95% confidence intervals from a panel regression at the DS level in which we regress tea yields (measured in kg/ha) in each DS-year on a set of year dummies and DS fixed effects. The regression is estimated by Weighted Least Squares (WLS), with weights given by DS-specific confidence weights reflecting the degree of precision of the remote sensing tea yield metric (see Appendix A.3.2 for details), and standard errors clustered at the DS level. All regressions control for the following weather variables: rainfall (in mm/season), number of rainy days, dry spells, solar radiation (in W/m2), VPD (in hectopascals), average minimum and maximum temperatures (in degrees Celsius). These season- and DS-specific weather variables are included separately for the Yala and Maha seasons. See Appendix A.10 for definitions of weather variables and further details.

other variables constant. That is, for a given level of total income, more unequal regions tend to spend a lower share of that income on agricultural goods. The intuition is that redistributing income away from poor farmers into the the hands of richer farmers moves money into the hands of consumers who spend a relatively lower share of this additional income on agriculture, due to non-homothetic preferences. This intuition is reinforced by the fact that the effect of inequality on agricultural expenditure disappears in the homothetic case ($\eta = 1$).

In this section, we test this prediction using 2019 HIES data. Specifically, we use data on landholdings and expenditures on various goods to construct district-level measures of aggregate food expenditures, aggregate total expenditures, and the variance of (the log of) landholding size as proxies for X_n^A, E_n, and σ_{Ln}^2, respectively. We then regress the log of food expenditures on land inequality and the log of total expenditures. Results are displayed in Appendix Table A6. The coefficient on log total expenditures is 0.893 in Column (1). A 10% increase in total expenditure is associated with an increase of only around 9% in our agricultural expenditure proxy, which is itself indicative of non-homothetic preferences as it implies that agriculture's expenditure *share* declines with income. More importantly for us, the coefficient on the land inequality proxy is -0.0035, pointing to a negative association between inequality and agricultural expenditure, as predicted by equation (14). Column (2) repeats the estimation strategy of Column (1) but uses an adjusted version of the dependent variable that maps more closely to the theoretical constructs in equation (14), with very similar results.[100]

Table A6: District-Level Food Expenditure and Land Inequality

	Dependent variable: Log of food expenditures	
	(1)	(2)
Land inequality ($\hat{\sigma}_{Ln}^2$)	-0.0035	-0.0036
	(0.0040)	(0.0041)
Log of total expenditures ($\ln(E_n)$)	0.893***	0.890***
	(0.0259)	(0.0265)
Constant	1.700***	1.745**
	(0.6305)	(0.6469)
Adjusted $\ln(X_n^A)$?	NO	YES
Observations	25	25

Notes: *** denotes significance at the 1% level, ** at the 5% level. The table shows results of district-level regressions relating aggregate food expenditures to land inequality and aggregate total expenditures. Regressions are estimated by Ordinary Least Squares (OLS) using data from the 2019 edition of the HIES. Land inequality in each district is defined as the variance of the log of landholding size across its landholders. Standard Errors are shown in parentheses. Column (2) uses an adjusted version of the dependent variable that maps more closely to the theoretical constructs in equation (14), namely $\ln(X_n^A - \hat{\phi}E_n)$, where X_n^A is aggregate food expenditures and E_n is aggregate total expenditures in district n, and $\hat{\phi} = 0.0105$.

[100] Specifically, in Column (2) the dependent variable is not $\ln(X_n^A)$ but $\ln(X_n^A - \hat{\phi}E_n)$, with $\hat{\phi} = 0.0105$ being our estimate of ϕ as displayed on Table 2 and explained in Section 6.1.1.

Figure A15: Dynamic Ban Effects on Districts' Night Lights (High versus Low Fertilizer Importance)

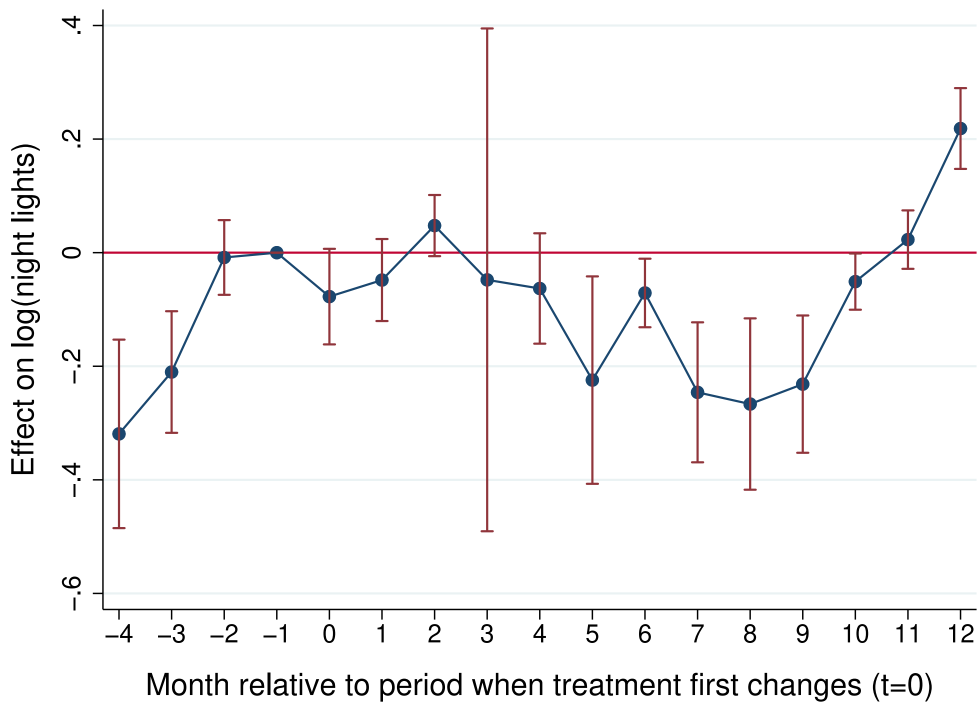

Notes: The figure shows De Chaisemartin and d'Haultfoeuille (forthcoming) estimates of the effect of the fertilizer import ban on night light intensity, relative to the last month before the ban takes effect, using a district-level panel. The outcome variable is the log of the average night light intensity across the district's cells in a given month. We compute each district's baseline fertilizer importance index as the ratio of the value of fertilizer inputs (gross of subsidies) used in its agricultural sector to the district's aggregate income in our estimated baseline model. The treatment variable is a dummy which takes the value of one from May 2021 onwards for districts above the median of the index. Below-median districts serve as the control group. The y-axis shows the effects of the fertilizer import ban on the log of average night light intensity, where negative values correspond to decreases in night lights.

D Estimation Details

D.1 Preference Parameters

Engel Elasticity. To estimate the Engel elasticity (η), we start from model equation (13). Rearranging, taking logs, and adding household (h) subscripts, we obtain the following equation:

$$\ln(\xi^A_{nh} - \phi) = \ln(\nu) - \eta \ln(y_{nh}) + \eta\phi \ln(P^A_n) + (1-\eta)\phi \ln(P^M_n)$$

We then collect relevant terms into a region fixed effect and add a household size control variable and an error term, obtaining equation (32).[101] As explained in Section 6.1.1, equation (32) is estimated using Two-Stage Least Squares (2SLS) with two IVs: income obtained from lottery (and other *ad hoc* gains) and income obtained from disasters and other relief payments. We use Weighted 2SLS estimation, weighing observations with the sampling weights available in the HIES data, and with standard errors clustered at the Primary Sampling Unit level. Complete estimation results can be seen in Column (1) of Appendix Table A7, including the estimated coefficients, standard errors, and first-stage results.

Cross-crop Elasticity of Substitution. To estimate σ_A, we first take logs of equation (17) and add household (h) subscripts, obtaining:

$$\ln(\beta^A_{nhk}) = -(\sigma_A - 1)\ln(P_{nhk}) + \ln(b_k) + (\sigma_A - 1)\ln(P^A_{nh})$$

We then collect relevant terms into crop- and household-specific fixed effects.[102] Finally, we add an error term to obtain equation (33), which is then estimated using Weighted 2SLS with standard errors clustered at the household level, with crop price $\ln(P_{nhk})$ instrumented by a dummy variable D^{GAEZ}_{nk} that indicates whether region n has any growing potential for crop k (according to the FAO-GAEZ data). The crops included in the estimation sample are: rice, maize, groundnut, onion, potato, coconut, and tea. See Column (2) of Appendix Table A7 for the complete estimation results.

Cross-origin Elasticity of Substitution. As explained in the text, σ_K is set so that the average increase in agricultural producer prices generated by the model-driven counterfactual exercise of Section 7.1 matches the corresponding increase that is observed in the data. Formally, for each candidate value s_K for parameter σ_K, we compute the corresponding counterfactual average price increase $\hat{\bar{p}}^{cf}(s_K)$ as follows:

$$\hat{\bar{p}}^{cf}(s_K) \equiv \sum_{k\in\mathcal{K}} \left(\frac{\sum_{i=1}^{I} p_{ik}Q_{ik}}{\sum_{l\in\mathcal{K}}\sum_{i=1}^{I} p_{il}Q_{il}} \right) \times \left(\frac{1}{\mathcal{N}_k} \sum_{i=1}^{I} \left(\frac{p^{cf}_{ik}(s_K)}{p_{ik}} \right) \mathbb{1}\{Q_{ik} > 0\} \right) \tag{44}$$

[101] Region-specific fixed effects are defined as: $\omega_n \equiv \ln(\nu) + \eta\phi\ln(P^A_n) + (1-\eta)\phi\ln(P^M_n)$.

[102] Crop- and household-specific fixed effects are defined respectively as: $\omega_k \equiv \ln(b_k)$ and $\omega_h \equiv (\sigma_A - 1)\ln(P^A_{nh})$.

Table A7: Estimates of Preference Parameters (2SLS)

	Panel A: Second Stage	
	Dependent variable:	
	Log of food's expenditure share (1)	Log of crop's expenditure share (2)
Log of household income [$\hat{\eta}$]	-0.656*** (0.129)	
Log of crop price [$1-\hat{\sigma}_A$]		-0.714*** (0.153)
Household size	0.135*** (0.029)	
	Panel B: First Stage	
	Dependent variable:	
	Log of household income (1)	Log of crop price (2)
Lottery/*ad hoc* income	8.81e-07*** (1.37e-07)	
Disaster/other relief payments	8.16e-07*** (1.42e-07)	
Regional potential yield > 0		-0.062*** (0.0044)
F-statistic	37.2***	199.4***
p-value	0.0000	0.0000
Fixed Effects?	PSU	Crop, HH
Clustered SEs?	by PSU	by HH
Weighted?	Yes	Yes
Observations	19,910	91,416

Notes: *** denotes significance at the 1% level. Columns (1) and (2) display results for the estimation of equations (32) and (33), respectively, by Weighted Two-Stage Least Squares (2SLS), with standard errors in parentheses. PSU stands for "Primary Sampling Unit", and HH stands for "household". Weighted regressions use the statistical weights provided by the 2019 edition of the HIES. In Column (1), the dependent variable in the second stage is the log of the *adjusted* food expenditure share, which is computed as $\ln(\xi^A_{nh} - \hat{\phi})$ for household h of district n, where ξ^A is food expenditure share and $\hat{\phi} = 0.0105$ is our estimate of parameter ϕ (see Table 2). See Section 6.1.1 for details on data sources, identification strategies, and variables' definitions.

where p_{ik} and Q_{ik} are baseline producer prices and quantities of crop k in district i, $\mathcal{N}_k$ is the number of Sri Lankan districts with nonzero production of crop k, $\mathcal{K}$ is the set of crops for which data on price increases is available, and $p_{ik}^{cf}(s_K)$ is the counterfactual producer price of crop k in district i assuming $\sigma_K = s_K$. The crops for which price increase data is available are: rice, maize, potatoes, cloves, onions and groundnuts.

We test several possible candidate values s_K, ultimately choosing the value $s_K^* = 2.377$ that generates a counterfactual average price increase of $\hat{\bar{p}}^{cf}(s_K^*) = 1.2265$, which matches the 22.65% increase observed in the data between January 2021 and January 2022 (shown in Appendix Figure A11). Specifically, our average crop price increase calculation is done in three steps: (i) computing the nominal price increase of each crop in $\mathcal{K}$ between January 2021 and January 2022; (ii) deflating these nominal increases using the period's cumulative inflation (17.61%) to obtain real price increases; and (iii) computing the average crop price increase as a weighted average of crop-level real price increases, with each crop's weight given by its revenue share at baseline (as in equation 44). We set $\hat{\sigma}_K = 2.377$ as our estimate for the elasticity of substitution across origins of the same crop.

D.2 Trade Costs

In Section 6.1.2, we estimate the parameters (α, $\{\beta_k\}_k$) that govern trade costs according to equation (35). To estimate semi-elasticity α, we use two key pieces of information: first, the tea transportation cost $p_{T,100}$ is 1,496 LKR/MT for a distance of 100 km (according to the Ceylon Petroleum Corporation); second, the price of tea p_{tea} is 455,473 LKR/MT (according to the DCS). The tea transportation cost is calculated based on the average distance from farms to tea brokers' warehouses in the city of Wattala, assuming that the cargo is transported by a lorry with a capacity of 7,500 kg charging typical prices.[103] To compute domestic "factory-gate" tea prices (p_{tea}), the price of green leaf needed to produce one MT of made tea is added to processing costs (labour, fuel, electricity, machinery upkeep, packing materials, and factory sundries). The corresponding iceberg trade cost can be written as $\tau_{ni,tea}^{A} = 1 + \frac{p_{T,100}}{p_{tea}}$, which we substitute in equation (35) to obtain the following relationship:

$$1 + \frac{p_{T,100}}{p_{tea}} = \exp(\hat{\alpha} \times 100) \tag{45}$$

where $\hat{\alpha}$ is our estimate for the semi-elasticity α. We solve this equation to obtain: $\hat{\alpha} = 3.3 \times 10^{-5}$.

To estimate crop-specific international trade costs ($\{\beta_k\}_k$), we use FAO data on each crop's producer price in the RoW ($p_{RoW,k}$), in USD/kg, and data on each crop's "benchmark price" at the

[103] The data is available at: https://ceypetco.gov.lk/historical-prices/.

port of Colombo ($p_{Colombo,k}$), in USD/MT (which we then convert to USD/kg).[104] Since data on port prices ($p_{Colombo,k}$) is only available for five crops, we use equation (36) to compute $\hat{B}_k$ for each of these five crops, then impute the $\hat{B}_k$ of the other three crops (cinnamon, cloves, and groundnuts) with the average $\hat{B}_k$ among the five crops for which data is available.

D.3 Production Function Parameters

As explained in Section 6.1.3, we estimate Cobb-Douglas coefficients γ_k^f by matching fertilizer cost shares (see equation 37). For that, we use data on national crop prices p_k^{LKA}, national annual crop production Q_k^{LKA}, and crop-level national annual fertilizer expenditure $M_k^{f,LKA}$.

As shown in equation (38), we then apportion residual input shares to land and labor using data from IO tables on value added and labor compensation. To do that, we map our eight crops to the IO table's agricultural sectors, which include "minor export crops", "highland crops", "tea growing" (separately for high, medium, and low elevations), "paddy", and "potatoes" as sectors. We classify cinnamon and cloves as "minor export crops", groundnuts, maize, and onions as "highland crops", and rice as "paddy". For tea, we aggregate high, medium, and low elevations into a single sector.

E Inversion Details

In Section 6.2 we describe our general approach to backing out unobserved exogenous variables (productivities, preference shifters, fertilizer endowments, and number of estate workers) by combining observed variables with the equilibrium structure of the model. In this section, we explain the details of this "inversion" procedure.

E.1 Inversion Procedure

It is possible to rewrite the market clearing equations (22, 25, 26, 27) of the model without QRs to obtain the following system of equations:[105]

$$p^f_{RoW} F^G = \sum_k \gamma_k^f p_{RoW,k} Q_{RoW,k} + \sum_{i=1}^{I} \sum_k \frac{\gamma_k^f p_{ik} Q_{ik}}{(1-s_k)}$$

$$Q_{ik} = \frac{1}{p_{ik}} \sum_n X_n^A b_k \left(\frac{P_{nk}}{P_n^A}\right)^{1-\sigma_A} b_{ik} \left(\frac{\tau_{ni,k}^A p_{ik}}{P_{nk}}\right)^{1-\sigma_K}$$

[104] This data is from the Central Bank of Sri Lanka's Economic and Social Statistics 2023.

[105] More specifically, to obtain the system of equations, substitute equation (20) and $p_i^f = p_{RoW}^f$ into fertilizer market clearing condition (22), substitute equations (7), (17), and (18) into crop market clearing condition (25), and substitute equations (16) and (19) into labor market clearing condition (26).

$$w_i^m N_i^m = \Big(\sum_{k\neq\text{tea}} \gamma_k^n p_{ik} Q_{ik}\Big) + \sum_n ((1-t)E_n - X_n^A)\Big(\frac{\tau_{ni}^M w_i}{T_i^M P_n^M}\Big)^{1-\sigma_M}$$

$$w_i^e N_i^e = \gamma_{tea}^n p_{i,\text{tea}} Q_{i,\text{tea}}$$

In these four equations, black variables are either directly observed or previously estimated in a prior stage, red variables are the key unobserved variables that the system solves for, and blue variables are "intermediary" variables in the sense that they are functions of black and red variables only.[106] We solve this system of equations numerically to recover all unobserved variables ($\boldsymbol{T}^M$, $\boldsymbol{b}$, F^G, $\boldsymbol{N}^e$, t, $\boldsymbol{p}^-$, $\boldsymbol{P}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$), and then recover agricultural productivities ($\boldsymbol{T}^A$) through equation (6), which links observed crop production to underlying productivity, crop prices, and input prices.

Having recovered the global fertilizer endowment (F^G) we apportion it across regions in two stages. First, the total Sri Lankan fertilizer endowment ($\sum_{i=1}^I F_i$) is chosen such that the ratio of domestic fertilizer endowment to domestic fertilizer use equals 3.5%, with the remainder of the global fertilizer endowment being given to the RoW.[107] Second, the Sri Lankan fertilizer endowment ($\sum_{i=1}^I F_i$) is distributed across Sri Lankan districts ($\{F_i\}_{i=1}^I$) in proportion to each district's number of mobile workers, N_i^m.

We should note that the equation system is homogeneous of degree zero in the preference shifters ($\boldsymbol{b}$). Therefore, normalizations are necessary. We set $b_{RoW,k} = 2$ for each crop k, as well as $b_{rice} = 3 \times 2^{\frac{\hat{\sigma}_A - 1}{1-\hat{\sigma}_K}}$. Naturally, these particular normalization choices do not affect the results in any meaningful way.

E.2 Data Requirements for Inversion

To solve the system of equations in Appendix E.1, we must plug in the "black" observed variables, which include both previously estimated parameters and data from various sources. In that process, it is first necessary to make decisions concerning the basic geographic units of analysis, the set of crops, and the time period for which to run the inversion. We use Sri Lanka's 25 districts (i.e. $I = 25$), which are the country's second-level administrative divisions, because that is the most granular level of aggregation for which we have data on important variables such as wages and producer prices. We use the eight crops (i.e. $K = 8$) for which we have data on both fertilizer intensity and producer prices, namely: cinnamon, cloves, groundnut, maize, onions, potatoes, rice, and tea. We choose to use 2019 for inversion because it is prior to the 2021 fertilizer import ban

[106]Blue variables can be recovered from black and red variables by using equations (5), (8), (14), (15), (17), (18), (19), and (21).

[107]That is, the fertilizer endowment of the RoW is given by: $F_{RoW} = F^G - \sum_{i=1}^I F_i$. Sri Lanka's 3.5% fertilizer endowment-to-use ratio is based on FAO data on fertilizer production and use by country and year.

(and to other potentially influential events such as the Covid-19 pandemic and Russia's invasion of Ukraine) while also offering rich household survey data.

Preference parameters (σ_A, σ_K, σ_M) are taken from Table 2 and production parameters ($\{\gamma_k^f$, γ_k^n, $\gamma_k^l\}_k$) from Table 3. Additionally, we use data on crop-specific fertilizer subsidy rates ($\{s_k\}_k$) and crop-district-specific prices ($\{p_{ik}\}_{i,k}$), production ($\{Q_{ik}\}_{i,k}$), and cultivated area ($\{L_{ik}\}_{i,k}$). And we use district-level annual wages ($\{w_i^m, w_i^e\}_i$), and numbers of mobile workers and farmers ($\{N_i^m$, $N_i^F\}_i$). Bilateral transportation costs ($\{\tau_{ni,k}^A\}_{n,i,k}$, $\{\tau_{ni}^M\}_{n,i}$) are computed using the coefficients estimated in Section 6.1.2.

For the RoW, we use data on 2019 crop prices ($\{p_{RoW,k}\}_k$), production ($\{Q_{RoW,k}\}_k$), and cultivated area ($\{L_{RoW,k}\}_k$), as well as on the 2019 wage (w_{RoW}). We assume that the RoW has no estate workers ($N_{RoW}^e = 0$) and that its tea sector uses mobile workers instead, and we rely on the measures of the number of workers (N_{RoW}^m) and farmers (N_{RoW}^F) described in Appendix A.11. All monetary variables expressed in USD are converted into LKR using the 2019 exchange rate. The international price of fertilizer, p_{RoW}^f =0.31 USD/kg, from our trade data (see Appendix A.2), is converted to LKR using the 2019 exchange rate.[108]

F Equivalent Variation

Since we are interested in the welfare effects of the fertilizer import ban, we use our estimated QSM of trade and agriculture to estimate such effects in Section 7. However, to do that we must first address some difficulties that emerge when using PIGL preferences. Careful inspection of equation (12) shows us that indirect utility can be negative for some combinations of prices and income.[109] As a result, it is not straightforward to express utility as "real income" (i.e. nominal income divided by a price index), as it is often done for other classes of preferences such as CES. Therefore, we propose an alternative welfare metric to measure changes in welfare between any two equilibria, namely, the EV.

Suppose we are comparing a given initial equilibrium to a given final equilibrium. Starting from the initial equilibrium, the EV for a given agent is defined as the change in income that would leave that agent with the same level of utility as she has in the final equilibrium if prices were held constant at their initial levels. Formally, the EV for that hypothetical agent living in region n is given by:

$$EV = V^{-1}(v_1; P_{n,0}^A, P_{n,0}^M) - y_0 \tag{46}$$

where v_1 is the agent's utility in the final equilibrium, y_0 is her income in the initial equilibrium, ($P_{n,0}^A$, $P_{n,0}^M$) are the prices indices in the initial equilibrium, and V^{-1} is the inverse of the indirect

[108] See Appendix A for more details on the data that is mentioned in this section.

[109] For example, if $\nu > 0$ and $P^A > P^M$, then utility will be negative for a sufficiently small level of income y.

utility function (12) with respect to the income argument (y).[110] This formula shows that the EV depends not only on regional prices, but also on the agent's initial income level y_0 and her final utility level v_1, both of which differ between workers and farmers of the same district and across farmers with different landholding sizes.

Thus, to assess the welfare changes faced by different agents, we define multiple types of agents in each district n for which we separately compute welfare effects in Section 7. For each type and district, the EV between the initial equilibrium (indexed by a 0 subscript) and the final equilibrium (indexed by a 1 subscript) is obtained from equation (46) by choosing the appropriate final utility (v_1) and initial income (y_0):[111]

1. **Worker**: all mobile workers in district n are identical and earn the same (net-of-tax) wage $(1-t)w_n^m$. Their EV is obtained by substituting $v_1 = V_{n,1}^m$ and $y_0 = (1-t)w_{n,0}^m$ in equation (46):

$$EV_n^m = V^{-1}(V_{n,1}^m; P_{n,0}^A, P_{n,0}^M) - (1-t_0)w_{n,0}^m$$

The same considerations apply to tea estate workers, whose EV can be written as:

$$EV_n^e = V^{-1}(V_{n,1}^e; P_{n,0}^A, P_{n,0}^M) - (1-t_0)w_{n,0}^e$$

2. **Median farmer**: defined as the farmer with a landholding size equal to the median landholding size in district n, which is $\exp(\mu_n)$ due to the log-normality of the land distribution (see Section 5.3.1). Her initial disposable income level is thus $y_0 = \exp(\mu_n)(1-t_0)R_{n,0}/L_i$, and her final utility v_1 is obtained by evaluating indirect utility function (12) at final disposable income $y_1 = \exp(\mu_n)(1-t_1)R_{n,1}/L_i$ and final prices ($P_{n,1}^A$, $P_{n,1}^M$). Her EV can thus be expressed as:

$$EV_n^{medF} = V^{-1}(V_{n,1}^{medF}; P_{n,0}^A, P_{n,0}^M) - \exp(\mu_n)(1-t_0)R_{n,0}/L_i,$$

$$\text{with: } V_{n,1}^{medF} = V\Big(\exp(\mu_n)(1-t_1)R_{n,1}/L_i, P_{n,1}^A, P_{n,1}^M\Big)$$

3. **Average farmer**: defined as the farmer with a landholding size equal to the average landholding size in district n, which is $\exp(\mu_n + \sigma_{Ln}^2/2)$ due to the log-normality of the land distribution (see Section 5.3.1). Similarly to the median farmer, it can be shown the average

[110] Simple algebraic manipulation shows that the inverse of the indirect utility function (with respect to the income argument y) is: $V^{-1}(v; P^A, P^M) = \eta^{\frac{1}{\eta}}(v + \nu \ln(P^A/P^M))^{\frac{1}{\eta}}(P^A)^{\phi}(P^M)^{1-\phi}$. For definitions of the agricultural and manufacturing price indices ($P_{n,0}^A$, $P_{n,0}^M$), see Sections 5.3.3 and 5.3.4.

[111] Some EV expressions also use welfare variables (V_n^m, V_n^e, V_n^F, V_n^{avg}) as inputs. For the definitions of these variables, see Section 5.6.

farmer's EV is:

$$EV_n^{avgF} = V^{-1}(V_{n,1}^{avgF}; P_{n,0}^A, P_{n,0}^M) - \exp(\mu_n + \sigma_{Ln}^2/2)(1-t_0)R_{n,0}/L_i,$$

$$\text{with: } V_{n,1}^{avgF} = V\Big(\exp(\mu_n + \sigma_{Ln}^2/2)(1-t_1)R_{n,1}/L_i, P_{n,1}^A, P_{n,1}^M\Big)$$

4. **Representative farmer:** defined as the farmer whose utility level is equal to V_n^F, the district's average farmer utility. Her EV can be obtained by substituting $v_1 = V_n^F$ and $y_0 = V^{-1}(V_{n,0}^F; P_{n,0}^A, P_{n,0}^M)$ in equation (46):

$$EV_n^F = V^{-1}(V_{n,1}^F; P_{n,0}^A, P_{n,0}^M) - V^{-1}(V_{n,0}^F; P_{n,0}^A, P_{n,0}^M)$$

5. **Representative agent**: defined as a hypothetical agent whose income level is just enough to obtain the district's average utility level V_n^{avg}. Her EV is obtained by substituting $v_1 = V_{n,1}^{avg}$ and $y_0 = V^{-1}(V_{n,0}^{avg}; P_{n,0}^A, P_{n,0}^M)$ in equation (46):

$$EV_n^{avg} = V^{-1}(V_{n,1}^{avg}; P_{n,0}^A, P_{n,0}^M) - V^{-1}(V_{n,0}^{avg}; P_{n,0}^A, P_{n,0}^M)$$

When reporting welfare effects in Section 7, we use the relative version (in %) of the EV to make relative welfare losses more comparable across regions and agent types. This relative EV can be computed by dividing the EV by the corresponding initial income level: $\frac{EV}{y_0}$.

G Proofs

G.1 Aggregate Agricultural Expenditure

In Section 5.3.2, we claim that agricultural expenditure share in region n is given by equation (14):

$$X_n^A = \phi(1-t)E_n + \nu(1-t)^{1-\eta}\left(\frac{N_n^m(w_n^m)^{1-\eta} + N_n^e(w_n^e)^{1-\eta} + N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}}}{(P_n^A)^{-\eta\phi}(P_n^M)^{-\eta(1-\phi)}}\right)$$

where P_n^A and P_n^M are sectoral price indices, $r_n \equiv \frac{R_n}{N_n^F}$ is the average land rent earned by the region's farmers, t is the income tax rate, and E_n is aggregate income.

To derive this equation, we first separately derive aggregate agricultural expenditure by middlemen ($X_n^{A,C}$), by workers ($X_n^{A,W}$), and by farmers ($X_n^{A,F}$), then add them up. Aggregate expenditure by middlemen is simply given by multiplying their total disposable income by ϕ, the asymptotic agricultural share of expenditure:

$$X_n^{A,C} = \phi(1-t)(p_n^f F_n + QR_i^{rent}) \tag{47}$$

Aggregate expenditure by workers ($X_n^{A,W}$) is given by multiplying the numbers of mobile and estate workers (N_n^m and N_n^e), earnings per worker (i.e. wages w_n^m and w_n^e), the net-of-tax rate $(1-t)$, and the fraction of income spent on agricultural goods (given by equation 13 with income y evaluated at the relevant net-of-tax wages):

$$\begin{aligned} X_n^{A,W} &= N_n^m(1-t)w_n^m \xi^A((1-t)w_n^m, P_n^A, P_n^M) + N_n^e(1-t)w_n^e \xi^A((1-t)w_n^e, P_n^A, P_n^M) = \\ &= \phi(1-t)(N_n^m w_n^m + N_n^e w_n^e) + \nu(1-t)^{1-\eta}((P_n^A)^\phi (P_n^M)^{1-\phi})^\eta (N_n^m (w_n^m)^{1-\eta} + N_n^e (w_n^e)^{1-\eta}) \end{aligned} \tag{48}$$

For an individual farmer h who lives in region n and earns land rent R^h, expenditure on agriculture is $(1-t)R^h\xi^A((1-t)R^h, P_n^A, P_n^M)$, the product of her net-of-tax income R^h and her agricultural expenditure share (see equation 13). Substituting equation (13) and taking the expectation operator across the region's farmers, the average agricultural expenditure per farmer ($\bar{X}_n^{A,F}$) can be written as:

$$\bar{X}_n^{A,F} = \phi(1-t)r_n + \nu(1-t)^{1-\eta}((P_n^A)^\phi (P_n^M)^{1-\phi})^\eta \mathbb{E}[(R^h)^{1-\eta}] \tag{49}$$

From equation (11), land rent R^h has a log-normal distribution. Therefore, the expectation $\mathbb{E}[(R^h)^{1-\eta}]$ can be obtained by using the distribution's properties. Specifically, a power of a log-normal variable is also log-normal:

$$(R_n^h)^a \sim \text{log-N}(a[\mu_n + \ln(R_n) - \ln(L_n)], a^2\sigma_{Ln}^2) \tag{50}$$

for any constant $a \neq 0$. Evaluating equation (50) with $a = 1-\eta$ and using the relationship between the parameters of a log-normal distribution and its average, we conclude that:[112]

$$\mathbb{E}[(R^h)^{1-\eta}] = \exp((1-\eta)[\mu_n + \ln(R_n) - \ln(L_n)] + (1-\eta)^2\sigma_{Ln}^2/2)$$

By using the properties of the exponential and rearranging terms, this expression can be rewritten as follows:[113]

[112] Namely, we use that fact that the average of any random variable that is distributed log-N(m, n^2) is equal to $\exp(m + n^2/2)$.

[113] Namely, we use the properties that $\exp(m+n) = \exp(m \times n)$ and $(\exp(m))^n = \exp(mn)$, for any real numbers m and n.

$$\mathbb{E}[(R^h)^{1-\eta}] = \Bigg(\exp\Big(\mu_n + \frac{\sigma_{Ln}^2}{2}\Big) \frac{N_n^F}{L_n} \times \frac{R_n}{N_n^F} \times \exp\Big(- \eta \frac{\sigma_{Ln}^2}{2}\Big)\Bigg)^{1-\eta}$$

It turns out that $\exp(\mu_n + \sigma_{Ln}^2/2)$ is the average landholding size in region n, which can also be written as $\frac{L_n}{N_n^F}$.[114] Using this fact, cancelling terms, and applying the definition of r_n, we can rewrite the previous equation as follows:

$$\mathbb{E}[(R^h)^{1-\eta}] = r_n^{1-\eta} \exp\big(- \eta(1-\eta)\frac{\sigma_{Ln}^2}{2}\big) \tag{51}$$

Substituting (51) into equation (49) and multiplying $\bar{X}_n^{A,F}$ by the number of farmers (N_n^F), we obtain the total agricultural expenditure by farmers of region n:

$$X_n^{A,F} = \bar{X}_n^{A,F} \times N_n^F = \phi(1-t)r_n N_n^F + \nu(1-t)^{1-\eta}((P^A)^{\phi}(P^M)^{1-\phi})^{\eta} N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}} \tag{52}$$

Total agricultural expenditure X_n^A in region n is obtained by summing the agro expenditures of middlemen, farmers, and workers:

$$X_n^A = X_n^{A,C} + X_n^{A,F} + X_n^{A,W}$$

Using equations (47), (52), and (48) to substitute for $X_n^{A,C}$, $X_n^{A,F}$, and $X_n^{A,W}$, rearranging terms, and noting that $R_n = r_n N_n^F$ and $E_n = N_n^m w_n^m + N_n^e w_n^e + R_n + p_n^f F_n + QR_i^{rent}$ (see equation 15), we obtain equation (14).

G.2 Average Farmer Welfare

In Section 5.6, we claim that average farmer welfare in region i is given by equation (30):

$$V_i^F = \frac{1}{\eta}\Big(\frac{1}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} e^{\eta[\mu_i + \ln(\frac{(1-t)R_i}{L_i})] + \eta^2\frac{\sigma_L^2}{2}} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big)$$

To prove this equation, consider that the indirect utility of an individual farmer h who lives in region i, earns land rent R^h, and pays a tax rate t is given by $V((1-t)R^h, P_i^A, P_i^M)$. Using equation (12), taking the expectation operation across farmers, and noting that P_i^A and P_i^M are the same for all the farmers in the region, average farmer indirect utility can be written as:

[114] This comes from the fact that landholding size follows a log-normal distribution with parameters μ_n and σ_{Ln}^2, as described in Section 5.3.1.

$$V_i^F \equiv \mathbb{E}[V((1-t)R^h, P_i^A, P_i^M)] = \frac{1}{\eta}\Big(\frac{1-t}{(P^A)^{\phi}(P^M)^{1-\phi}}\Big)^{\eta}\mathbb{E}[(R^h)^{\eta}] - \nu \ln\Big(\frac{P^A}{P^M}\Big) \tag{53}$$

So our remaining task is to find an expression for $\mathbb{E}[(R^h)^\eta]$. From equation (50) with $a = \eta$, we know that random variable $(R^h)^\eta$ has a log-normal distribution with parameters $\eta(\mu_n + \ln(R_n) - \ln(L_n))$ and $\eta^2\sigma_{Ln}^2$. Using the relationship between the log-normal's parameters and its average, we then conclude that:

$$\mathbb{E}[(R^h)^\eta] = \exp(\eta(\mu_n + \ln(\frac{R_n}{L_n})) + \eta^2\sigma_{Ln}^2/2)$$

Substituting this expression into equation (53), substituting $\exp(\eta\ln(1-t))$ for $(1-t)^\eta$, and rearranging, we obtain equation (30).

Table A8: Regional Ban Effects

	Worker		**Farmer**			**Representative**
District	Mobile	Estate	Median	Average	Representative	**Agent**
Ampara	-2.89%	.	-22.38%	-22.38%	-22.38%	-8.25%
Anuradhapura	-2.93%	.	-23.08%	-23.08%	-23.08%	-9.97%
Baddulla	+0.68%	-16.58%	-22.48%	-22.47%	-22.47%	-5.11%
Batticaloa	-0.37%	.	-22.67%	-22.67%	-22.67%	-4.26%
Colombo	+2.55%	-16.58%	-20.41%	-20.32%	-20.35%	+2.40%
Galle	+2.19%	-16.58%	-9.38%	-9.37%	-9.37%	+0.46%
Gampaha	+2.51%	.	-21.96%	-21.94%	-21.94%	+2.08%
Hambantota	-0.85%	.	-19.76%	-19.76%	-19.76%	-5.93%
Jaffna	+2.18%	.	-19.51%	-19.50%	-19.50%	+0.48%
Kaluthara	+2.37%	-16.58%	-16.69%	-16.67%	-16.68%	+1.40%
Kandy	+2.30%	-16.58%	-15.72%	-15.70%	-15.71%	+0.70%
Kegalle	+2.30%	-16.58%	-16.91%	-16.89%	-16.90%	+0.77%
Killinochchi	-6.09%	.	-20.74%	-20.74%	-20.74%	-10.96%
Kurunagala	+0.53%	.	-22.78%	-22.78%	-22.78%	-4.32%
Mannar	-4.51%	.	-21.02%	-21.02%	-21.02%	-9.11%
Matale	+0.93%	-16.58%	-21.22%	-21.21%	-21.21%	-3.19%
Matara	+1.94%	-16.58%	-10.29%	-10.28%	-10.28%	-0.31%
Monaragala	-2.89%	.	-23.34%	-23.33%	-23.33%	-10.28%
Mullaittivu	-11.62%	.	-17.74%	-17.74%	-17.74%	-14.34%
Nuwareliya	+1.91%	-16.58%	-20.09%	-20.08%	-20.08%	-3.65%
Polannaruwa	-8.51%	.	-20.02%	-20.02%	-20.02%	-13.30%
Puttalam	+1.71%	.	-20.87%	-20.86%	-20.86%	-0.74%
Rattnapura	+2.07%	-16.58%	-14.46%	-14.45%	-14.45%	-0.80%
Trincomalee	-1.88%	.	-22.33%	-22.33%	-22.33%	-6.87%
Vavuniya	-0.32%	.	-22.15%	-22.15%	-22.15%	-5.71%
Average	-0.67%	-16.58%	-19.52%	-19.51%	-19.51%	-4.35%

Notes: for each district, the table shows the % welfare changes of six types of agents between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs. The median (resp. average) farmer is defined as a farmer whose landholding size is equal to the median (resp. average) of her district. The representative farmer of district i is defined as a farmer whose utility is equal to the district's average farmer utility (V_i^F). Similarly, the representative agent of district i is defined as a hypothetical agent whose (net-of-taxes) income is just enough to provide her with the district's average utility (V_i^{avg}). Welfare changes are defined as EVs: starting from the agent's baseline (net-of-taxes) income, the EV is the % change in (net-of-taxes) income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels.

H Frictions to Worker Intersectoral Mobility

Our main model assumes that mobile workers can costlessly move across different sectors of employment. There may be some concern that our results are being driven by this potentially unrealistic assumption. In particular, estimated welfare effects could be more negative for mobile workers if our model recognized that some of them are stuck in agriculture due to mobility frictions, and are thus unable to move to manufacturing once the ban hits.

In this section, we show that the assumption of frictionless worker intersectoral mobility is not an important driver of our key results. To do that, we extend our main model to account for mobility frictions and then show that the model-predicted welfare losses to the representative agent varies very little with the severity of these frictions.

H.1 Model

To account for frictions to worker intersectoral mobility, we propose a dynamic version of our model.[115] The main change is that, in each period, each mobile worker finds herself employed in a certain sector and must then choose whether she will stay in her current sector or move to the other sector in the next period.

We assume mobile workers are infinitely lived and discount future utility with a discount factor $\delta < 1$. In each period, a worker obtains "flow" utility from consumption and from work enjoyment (represented by a idiosyncratic sector-specific preference shock). She also chooses whether to stay in her current sector or change sectors, in the latter case suffering a "mobility cost" paid in disutility. Using Bellman's principle of optimality, the expected utility in period T of a worker who lives in region i and works in sector $s \in \{A, M\}$ is given by:

$$U_{isT}(\epsilon_T) = V(P^A_{iT}, P^M_{iT}, (1 - t_T)w^s_{iT}) + \delta \times \max_{r\in\{A,M\}} \{\epsilon_{rT} - \tilde{\mu}_{sr} + \mathbb{E}_t[U_{ir,T+1}(\epsilon_{T+1})]\} \qquad (54)$$

where $\epsilon_T \equiv [\epsilon_{AT}, \epsilon_{MT}]$ is a vector containing two sector-specific idiosyncratic preference shocks, V is the indirect utility function given by equation (12), P^A_{iT} (resp. P^M_{iT}) is the agricultural (resp. manufacturing) price index, w^s_{iT} is the wage in sector s, t_T is the income tax rate, and $\tilde{\mu}_{sr}$ is the mobility cost of moving from sector s to sector r.[116] The preference shocks ϵ_{sT} are independently distributed across workers, sectors, and periods, and follow a Gumbel distribution with location parameter equal to zero and scale parameter equal to $\frac{1}{\theta} > 0$:

[115] This dynamic formulation is influenced by Caliendo et al. (2019) and Artuc et al. (2010).

[116] We assume that a worker pays no mobility cost when she chooses to stay in her current sector, i.e. $\tilde{\mu}_{AA} = \tilde{\mu}_{MM} = 0$. We further assume that mobility costs are symmetric, i.e. $\tilde{\mu}_{AM} = \tilde{\mu}_{MA} = \tilde{\mu}$.

$$\Pr(\epsilon_{sT} \leq x) = e^{-e^{-\theta x}}, \forall x \in \mathbb{R} \tag{55}$$

The max operator in equation (54) captures the decision each worker makes in each period of which sector she will work in in the next period.[117] Using the properties of the Gumbel distribution in equation (55), it can be shown that the average utility across workers in a given region-sector-period can be written as the following closed-form expression:

$$V^m_{isT} \equiv \mathbb{E}_{T-1}[U_{isT}(\epsilon_T)] = V(P^A_{iT}, P^M_{iT}, (1-t_T)w^s_{iT}) + \frac{\delta\gamma}{\theta} + \frac{\delta}{\theta} \ln\left(\sum_{r\in\{A,M\}} [\exp(V^m_{ir,T+1} - \tilde{\mu}_{sr})]^\theta \right)$$

where V^m_{isT} is the average utility across workers of sector s and region i in period T, and $\gamma \approx 0.577$ is the Euler-Mascheroni constant. Furthermore, the properties of the Gumbel distribution also imply that the probability that a worker who is in sector s and region i in period T will choose to work in sector r in period $(T+1)$ is given by: $\frac{[\exp(V^m_{ir,T+1} - \tilde{\mu}_{sr})]^\theta}{\sum_{k\in\{A,M\}} [\exp(V^m_{ik,T+1} - \tilde{\mu}_{sk})]^\theta}$. This closed-form expression for the transition probabilities allow us to write a law of motion for employment levels:

$$N^s_{i,T+1} = \sum_{k\in\{A,M\}} \left(\frac{\bar{V}^m_{is,T+1}}{\mu_{ks}\Pi_{ik,T+1}} \right)^\theta N^k_{i,T}$$

where $N^s_{i,T}$ is employment in sector s of region i in period T, $\bar{V}^m_{isT} \equiv \exp(V^m_{isT})$, $\mu_{ks} \equiv \exp(\tilde{\mu}_{ks})$, and $\Pi_{isT} \equiv \sum_{r\in\{A,M\}} \mu^{-\theta}_{sr} \bar{V}^\theta_{irT}$ can be interpreted as the option value of working in sector s in terms of its "access" to other sectors (as measured by $\{\mu^{-\theta}_{sr}\}_r$).

In the model's steady state, endogenous variables do not change, so we can eliminate the time subscripts and rewrite the key equilibrium relationships as follows:

$$\ln(\bar{V}^m_{is}) = V(P^A_i, P^M_i, (1-t)w^s_i) + \frac{\delta\gamma}{\theta} + \delta \ln(\Pi_{is}) \tag{56}$$

$$\Pi^\theta_{is} \equiv \sum_{r\in\{A,M\}} \mu^{-\theta}_{sr} (\bar{V}^m_{ir})^\theta \tag{57}$$

$$N^s_i = \sum_{k\in\{A,M\}} \left(\frac{\bar{V}^m_{is}}{\mu_{ks}\Pi_{ik}} \right)^\theta N^k_i \tag{58}$$

The new mobility module represented by equations (56), (57), and (58) can be plugged into the main model with minor modifications. First, because intersectoral mobility is imperfect, the equilibrium agricultural wage and manufacturing wage may now differ within the same region. To

[117] Note that, due to idiosyncratic preference shocks, in any given period there exist otherwise identical workers in the same sector and region who will choose to work in different sectors in the following period.

reflect this possibility, we must thus adjust the equilibrium relationships governing relevant agricultural prices (equation 3), aggregate expenditure on agriculture (equation 14), aggregate income (equation 15), and manufacturing expenditure shares (equation 19), respectively, as follows:

$$h_{ik} = \begin{cases} (p_{ik})^{\frac{1}{\gamma_k^l}} (w_i^e)^{\frac{-\gamma_k^n}{\gamma_k^l}} (p_i^{f,k})^{\frac{-\gamma_k^f}{\gamma_k^l}} (\kappa_k)^{\frac{-1}{\gamma_k^l}} & \text{, if } k = \text{tea and } i \neq \text{RoW} \\ (p_{ik})^{\frac{1}{\gamma_k^l}} (w_i^A)^{\frac{-\gamma_k^n}{\gamma_k^l}} (p_i^{f,k})^{\frac{-\gamma_k^f}{\gamma_k^l}} (\kappa_k)^{\frac{-1}{\gamma_k^l}} & \text{, otherwise} \end{cases} \tag{59}$$

$$X_n^A = \phi(1-t)E_n + \frac{\nu(1-t)^{1-\eta}}{((P_n^A)^{\phi}(P_n^M)^{1-\phi})^{-\eta}} (N_n^A (w_n^A)^{1-\eta} + N_n^M (w_n^M)^{1-\eta} + N_n^e (w_n^e)^{1-\eta} + N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}}) \tag{60}$$

$$E_n = w_n^A N_n^A + w_n^M N_n^M + w_e N_n^e + R_n + p_n^f F_n + QR_i^{rent} \tag{61}$$

$$\beta_{ni}^M = \left(\frac{\tau_{ni}^M w_i^M}{T_i^M P_n^M}\right)^{1-\sigma_M}, \quad P_n^M = \left(\sum_i (\tau_{ni}^M w_i^M / T_i^M)^{1-\sigma_M}\right)^{\frac{1}{1-\sigma_M}} \tag{62}$$

Naturally, the other equilibrium relationship that must be adjusted is labor market clearing. Instead of equation (26), we now need two equations for each region, one for agricultural labor and one for manufacturing labor:

$$N_i^A = \frac{1}{w_i^A} \sum_{k \neq \text{tea}} \gamma_k^n p_{ik} Q_{ik} \tag{63}$$

$$N_i^M = \frac{1}{w_i^M} \sum_n X_n^M \beta_{ni}^M \tag{64}$$

With these adjustments, we are now ready to state the equilibrium definition for the case with frictions:

Definition 3 (Equilibrium with Frictions) *Given the model's parameters (R, K, η, ν, ϕ, $\boldsymbol{\sigma}$, $\boldsymbol{\gamma}$, θ, δ) and exogenous variables ($\boldsymbol{F}$, $\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{T}$, $\boldsymbol{\tau}$, $\boldsymbol{b}$, $\bar{f}$, $\boldsymbol{s}$, $\boldsymbol{\mu}$), a **steady-state equilibrium with frictions to worker intersectoral mobility** is a set of endogenous variables (t, $\boldsymbol{p^A}$, $\boldsymbol{p^f}$, $\boldsymbol{p^-}$, $\boldsymbol{P}$, $\boldsymbol{w^A}$, $\boldsymbol{w^M}$, $\boldsymbol{w^e}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$, $\boldsymbol{Q}$, $\boldsymbol{QR}^{rent}$, $\boldsymbol{N^A}$, $\boldsymbol{N^M}$, $\boldsymbol{\Pi}$, $\bar{\boldsymbol{V}}$) satisfying equations 5, 6, 7, 8, 16, 17, 18, 20, 21, 25, 27, 56, 57, 58, 59, 60, 61, 62, 63, 64. Additionally, if QRs are **not** in place, equation 22 must hold, with $p_i^f = p_{RoW}^f$ for all regions $i \in \{1, ..., R\}$. Alternatively, if QRs **are** in place, equations 23 and 24 must hold, with $p_i^f = p_{LKA}^f$ for all domestic regions $i \in \{1, ..., I\}$.* [118]

[118] Variables in bold are collections of region- and/or sector-level terms, as follows: $\boldsymbol{\mu} = \{\{\mu_{rs}\}_{r\in\{A,M\}}\}_{s\in\{A,M\}}$, $\bar{\boldsymbol{V}} = \{\{\bar{V}_{is}\}_{s\in\{A,M\}}\}_{i=1}^R$, $\boldsymbol{\Pi} = \{\{\Pi_{is}\}_{s\in\{A,M\}}\}_{i=1}^R$, $\boldsymbol{w^A} = \{w_i^A\}_{i=1}^R$, $\boldsymbol{w^M} = \{w_i^M\}_{i=1}^R$, $\boldsymbol{w^e} = \{w_i^e\}_{i=1}^R$, $\boldsymbol{N^A} = \{N_i^A\}_{i=1}^R$, $\boldsymbol{N^M} = \{N_i^M\}_{i=1}^R$. The other variables in bold are exactly the same as in the equilibrium of the original

Given the equilibrium variables, we can compute average welfare for various agent types. For the average mobile worker in sector $s \in \{A, M\}$ of region i, average welfare is simply given by $V_{is}^m \equiv \ln(\bar{V}_{is}^m)$. For the sake of comparability, we slightly modify the definition of welfare for the other agent types, assuming that they are also infinitely lived with discount rate δ. The average welfare in region i can then be computed as follows:

$$V_i^{avg} = \frac{N_i^A}{N_i} V_{iA}^m + \frac{N_i^M}{N_i} V_{iM}^m + \frac{N_i^e}{N_i} \frac{V_i^e}{(1-\delta)} + \frac{N_i^F}{N_i} \frac{V_i^F}{(1-\delta)}$$

where average per-period utility for estate workers (V_i^e) and farmers (V_i^F) are given by equations (29) and (30), respectively. To compare agent welfare across different equilibria, we can still use the concept of EV, as explained in Appendix F.

H.2 Counterfactual Results

To assess the effect of the ban in a model with frictions to worker intersectoral mobility, we proceed in two steps. First, starting from the frictionless baseline equilibrium estimated in Section 6, we choose a specific value for the parameter μ that governs the severity of frictions (i.e. $\mu_{AM} = \mu_{MA} = \mu$) and compute a "baseline equilibrium with frictions" by applying Definition 3 with no QRs.[119] Second, we use Definition 3 again to recompute the equilibrium with frictions, but now imposing QRs ($\bar{f} = 0$), as in Section 7.

By repeating this process for several different values of μ, we can evaluate the sensitivity of the ban's welfare impacts on each agent type to the severity of intersectoral mobility frictions. Appendix Table A9 displays the results of the exercise. Note that we explore a very wide range of friction intensities, going from the frictionless case ($\mu = 1$ or $\tilde{\mu} = 0$) to extremely severe frictions ($\mu = e^{21}$). The table shows us that welfare losses vary remarkably little with friction severity. For example, when there are no frictions the welfare losses for the representative agent are equivalent to a 2.08% income loss, but in the last row of the table, when frictions are most severe, her loss increases only to 3.06%.

Therefore, the takeaway message from the exercise seems to be that while the reallocation of mobile workers across sectors does affect the ban's welfare consequences, this influence is relatively less important compared to other, more significant factors.

model with perfect mobility (see Definition 1).

[119]To compute equilibria with frictions, it is also necessary to take a stand on the values of parameters θ and δ. Following Caliendo et al. (2019), we set $\delta = 0.99^4$, and $\theta = \frac{1}{5.34}$.

Table A9: Ban Effects on Welfare, by agent type (With Frictions to Worker Mobility)

Friction	**Mobile Worker**			**Estate**	**Farmer**			**Representative**
μ	Agro	Manufacturing	All	**Worker**	Median	Average	Representative	**Agent**
1	-0.95%	-0.81%	-0.85%	-15.49%	-14.55%	-14.55%	-14.55%	-2.08%
1.3	-0.96%	-0.82%	-0.86%	-15.52%	-14.62%	-14.62%	-14.62%	-2.11%
e^1	-1.00%	-0.84%	-0.88%	-15.63%	-14.83%	-14.82%	-14.82%	-2.21%
e^2	-1.04%	-0.86%	-0.90%	-15.76%	-15.11%	-15.10%	-15.10%	-2.33%
e^4	-1.08%	-0.85%	-0.90%	-15.98%	-15.70%	-15.69%	-15.70%	-2.53%
$e^{6.566}$	-1.08%	-0.79%	-0.84%	-16.19%	-16.47%	-16.46%	-16.46%	-2.73%
e^{10}	-1.02%	-0.64%	-0.70%	-16.37%	-17.40%	-17.39%	-17.40%	-2.88%
e^{15}	-0.92%	-0.46%	-0.54%	-16.50%	-18.38%	-18.37%	-18.37%	-3.00%
e^{21}	-0.86%	-0.35%	-0.44%	-16.56%	-18.97%	-18.96%	-18.97%	-3.06%

Notes: for each level of frictions to mobile workers' intersectoral mobility (μ) and for each agent type, the table shows the average % welfare change between the baseline equilibrium with no QRs and the counterfactual equilibrium with QRs, where the (unweighted) average is computed across Sri Lankan districts. Equilibria are computed assuming that the utility cost faced by a mobile worker when moving to a different sector of employment ($\mu_{AM} = \mu_{MA}$) is given by the value μ in the "Friction" Column (see Appendix H.1 for details on equilibria with frictions). Welfare changes are defined as EVs: starting from the agent's baseline (net-of-tax) income, the EV is the % change in (net-of-tax) income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels (see Appendix F for details on EVs and agent types).