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

Devaki Ghose† Eduardo Fraga‡ Ana Fernandes§

December 2023

Abstract

This paper quantifies the value of fertilizer for agricultural production and trade in a developing economy where agriculture is centrally important by using an unprecedented natural experiment whereby the government of Sri Lanka imposed an abrupt and unexpected ban on the imports of all chemical fertilizers in May 2021. The analysis combines novel high-frequency firm-level trade data, detailed agricultural ground production data, crop yield estimates from state-of-the-art remote sensing techniques, and dynamic event study designs. The findings show that the fertilizer ban led to dramatic declines in agricultural production, fertilizer imports, and exports of fertilizer-dependent crops. Using a quantitative trade model, the paper finds that the ban's welfare effects were equivalent to a 1.5 percent income reduction on average, with losses disproportionately concentrated on farmers (whose income is tied to agriculture) relative to workers and on regions specialized in the cultivation of relatively fertilizer-intensive crops. The findings quantify the equilibrium value of fertilizer in agriculture, an important estimate for any fertilizer-related policy (such as fertilizer subsidies) and for the public debate on the costs and benefits of environmental regulation more generally.

JEL Codes: D58; F13; F14; O13; Q17.
Keywords: Agriculture, trade, import ban, fertilizer, non-tariff measures.

*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, John Keyser, Gonzalo Varela, Athula Senaratne, Sheu Salau, Manoj Thibbotuwawa, Camille Reverdy, and Alejandro Forero for helpful discussions. We thank Mutlu Ozdogan for his work on the remote sensing estimates. 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.

†World Bank, DECRG. Email: dghose@worldbank.org
‡World Bank, DECRG. Email: epinheirofraga@worldbank.org
§World Bank, DECRG. Email: afernandes@worldbank.org

1 Introduction

Increased utilization of agricultural inputs, specifically fertilizers, has often been associated with enhanced agricultural productivity, a necessary condition for structural transformation and growth (Foster and Rosenzweig, 2008; Gollin et al., 2002; Bustos et al., 2016). Governments worldwide have subsidized fertilizers with the aim of increasing agricultural yields. How do we quantify the value of fertilizer for agricultural production and trade in a developing economy where agriculture is of central importance? The answer to this question is gaining increasing importance due to the Russian Federation-Ukraine war, which is causing fertilizer shortages and rising prices globally (Zereyesus et al., 2022; Hebebrand and Glauber, 2023), as well as heightened scrutiny of chemical fertilizers for their environmental impact.

The main attempt to answer this question has been through a series of field experiments in which fertilizer use has led to large yield increases. However, these trials are typically conducted on closely supervised experimental plots under carefully controlled conditions, so returns differ from those on real-world farms (Beaman et al., 2013; Duflo et al., 2011). While small-scale experiments are essential in assessing local effects, it is difficult to assess the general equilibrium effects once a policy is scaled up (Bergquist et al., 2022; Muralidharan and Niehaus, 2017). For example, the economy-wide value of access to fertilizer (or lack thereof) in general equilibrium can differ depending on regional differences in fertilizer use and crop specialization, the international tradability of crops, and the endogenous responses of economic agents.

In this paper, we quantify the value of fertilizer to an economy by taking advantage of an unprecedented natural experiment whereby the Government of Sri Lanka imposed an abrupt and unexpected ban on the imports of all chemical fertilizers in May 2021. Since most chemical fertilizers in Sri Lanka were imported and the market for organic fertilizer was virtually non-existent, the import ban resulted in an economy-wide fertilizer shortage. Using novel high-frequency granular firm-level customs trade data for Sri Lanka combined with state-of-the-art event study designs, we show that, at its peak, the ban resulted in a 99% decrease in fertilizer imports and a 33% decrease in the exports of relatively fertilizer-dependent agricultural producers. We complement these findings by developing new measures of rice yields and extent of cultivation over 20 years, utilizing satellite imagery and detailed weather variables. These reveal a historically unusual 32% post-ban decline in rice yields, which is particularly noteworthy since the cultivated rice area remained unchanged. We then build and estimate a quantitative model of agriculture and trade to quantify the general equilibrium effects of the ban. We find average welfare losses equivalent to a 1.5% decline in income, with stronger incidence on farmers (whose income is linked to agriculture) relative to workers (who are more sectorally mobile) and on geographic regions specialized in relatively fertilizer-intensive crops.

We begin by documenting three reduced-form stylized facts on the impacts of the fertilizer import ban by combining a rich set of novel data sources. To examine effects on imports, we rely on high-frequency granular customs transaction data on the universe of 2017-2022 firm-level imports for Sri Lanka and on import ban data hand-coded from regulatory gazettes. The import ban was imposed unexpectedly at a critical moment in the cultivation season and was subsequently withdrawn in the face of social upheaval (as discussed in Section 2). Given the ban's sudden nature, we can use our high-frequency data to examine its dynamic impacts over time based on the dynamic difference-in-differences (DiD) approach proposed by De Chaisemartin and d'Haultfoeuille (2022), which corrects the biases in standard dynamic DiD estimates that arise in cases (such as ours) in which the treatment is discontinued during the sample period. Our first stylized fact shows that fertilizer imports fell by more than 99% four months after the ban's introduction and remained low thereafter. This fall was entirely driven by a decline in import quantities, as import prices actually increased over the period.

To examine the ban's impacts on agricultural production, we zoom in on Sri Lanka's primary food crop, rice, whose production is highly reliant on fertilizer. We utilize government data on district-level yields and novel remote sensing yield estimates at the pixel level, which we develop through satellite imagery and pixel-level weather data. Our second stylized fact is a drastic decline of more than 30% between 2021 and 2022 in the rice yields of the major growing season (Maha season), whether those yields are measured from government statistics or remote sensing estimates. This is accompanied by a dramatic increase in rice imports, which were trivial prior to the fertilizer ban. Yield declines due to the ban are observed in all regions of Sri Lanka but are stronger in the North and West where major rice-producing districts are located.

To capture the downstream effects of the fertilizer import ban on agricultural exports, we combine our customs data on imports and exports with government data on fertilizer usage in the production of multiple agricultural products, and follow two research designs. First, we propose a firm-based design using exporters of agricultural products. We construct a firm-level measure of fertilizer usage based on firms' pre-ban export portfolio of agricultural products and their fertilizer intensity and define treated firms as those above the third quartile of fertilizer usage. We then apply our dynamic DiD approach to estimate the ban's effect on firm exports. Second, we propose a product-based design using agricultural exports. Focusing on single-export-product firms, we construct a product-level vector of input requirements based on these firms' history of imports of intermediate inputs (both fertilizer and others). We then construct a measure of product exposure to the ban by interacting fertilizer input requirements with the evolution of fertilizer import bans over time, with treatment status defined as being in the top quartile of exposure. We apply our dynamic DiD approach to estimate the ban's impacts on the exports of affected agricultural products. Our third stylized fact shows a 33% decline in agricultural exports for firms with high fertilizer dependence (relative to less dependent firms) four quarters after the import ban is introduced and a 95% decline in exports for highly exposed agricultural products (relative to less exposed products) in the same time frame. This finding is robust to alternative specifications and to seasonality corrections. Declines in export value are entirely driven by declines in export quantities.

While these three stylized facts demonstrate the fertilizer ban's impact on imports, agricultural production, and exports, it is hard to quantify the general equilibrium effects of the resulting lack of access to fertilizer (and the associated changes in welfare) without a theoretical model. This is because, while in general equilibrium the whole economy is affected, the DiD estimates only allow us to infer the effect of lack of fertilizer on the directly affected products (i.e. the treatment group) relative to the control group, which does not directly rely on fertilizer. Thus, we build a quantitative model of trade and agriculture based on a recent literature that allows for heterogeneity across crops in production technology (including fertilizer intensity) and for rich geographic heterogeneity across sub-national regions in terms of population, productivity, and crop specialization (Sotelo, 2020; Pellegrina, 2022; Farrokhi and Pellegrina, 2022). We follow Eckert and Peters (2023) and augment this quantitative model with non-homothetic preferences, generating a negative relationship between household income and agriculture's share of expenditure, which is an important feature of our data. We further assume that land is unequally distributed within each region, following a log-normal distribution (also a feature of the data). We show that, as a result, regional agricultural expenditure is decreasing in the regional level of land inequality (while holding aggregate income constant), a prediction that is borne out in the data. The model also features costly fertilizer trade, so the fertilizer cost increases caused by the ban can be modeled as an increase in the parameter governing the cost of international fertilizer transportation.

To bring the model to the data, we estimate the model's parameters using a diverse set of empirical strategies and a very rich set of mostly novel data sources. Parameters governing preference non-homotheticity are estimated using the household survey on income and expenditure. Importantly, we estimate the income elasticity of food's expenditure share (i.e., the Engel elasticity) through an instrumental variables (IV) approach utilizing unexpected income shocks (lottery winnings and disaster relief). Additional survey data on crop prices and expenditures by household allows us to estimate the elasticity of substitution *across crops* through an IV approach leveraging the effect on local crop prices of regional crop suitability, which is measured by the Global Agro-Ecological Zones (GAEZ) project from the Food and Agriculture Organization of the United Nations (FAO), and which is determined by exogenous regional geography.

We estimate the elasticity of substitution across different origins of the *same crop* by using a panel IV approach. This method examines variations in agricultural export performance over time across districts, which are influenced by differences in crop specialization, heterogeneity in fertilizer requirements across crops, and changes in fertilizer prices over time. We use the properties of our assumed Cobb-Douglas agricultural production function to estimate technology parameters for each crop by combining information from Sri Lanka's Input-Output table with additional data on fertilizer prices and crop-level output and fertilizer requirements. Finally, we estimate the parameters governing district-level land inequality using household survey data on landholdings.

The model itself is estimated by backing out the set of unobservable exogenous variables (productivities, taste shifters, global fertilizer endowment) and unobservable endogenous variables (incomes, land rents, expenditures, price indices) that solve the equilibrium system of equations given observed exogenous variables (population, labor force, land endowments, trade costs) and observed endogenous variables (wages, fertilizer prices, crop outputs, crop prices) from 2019, a pre-ban year that we use as our baseline.[1] The estimated model is then used for a counterfactual exercise in which we alter a single exogenous variable – the international fertilizer trade cost – and recompute the equilibrium, interpreting any differences between the resulting equilibrium and the original baseline equilibrium as being caused by the change in fertilizer trade cost. We calibrate the magnitude of the trade-cost increase so that the model matches the increase in real fertilizer prices observed in the Sri Lankan data (from \$0.31/kg to \$0.58/kg) between 2019 and 2022.

Results from the counterfactual exercise indicate that the import ban decreased fertilizer imports by \$11.5 million, halving the quantity of fertilizer used in Sri Lankan agriculture and thus decreasing crop yields by anything between 1.3% (groundnuts) to 14.3% (potatoes), depending on the crop's fertilizer intensity. Lower production led to a decrease of \$137.7 million in agricultural exports, suggesting the ban failed even when judged solely according to its presumed objective of saving foreign exchange through a reduction in the trade deficit. The incidence of the resulting welfare losses was highly heterogeneous. The representative farmer, whose source of income is tied to the agricultural sector, suffered losses equivalent to 4%-8% of her baseline income, depending on her geographic location. Workers, in contrast, are more sectorally mobile, so their losses were limited to the equivalent of 3.2% of their baseline income even in the worst affected, most agricultural districts, where the employment "buffer" provided by manufacturing is smallest. Moreover, farmers' losses varied across space and were strongly associated with district-level specialization in crops that are relatively fertilizer-intensive. Therefore, the fertilizer ban's effects were not only unequal across economic occupations but also across geographies.

Our quantification of the value of fertilizer in general equilibrium adds to the extensive literature concerned with evaluating the effects of fertilizer (or, more broadly, modern agricultural inputs) on yields and agricultural productivity, which mainly uses randomized control experiments (Carter et al., 2021; Beaman et al., 2013; Duflo et al., 2008, 2011). In contrast, our setting of a nationwide, sudden, and unprecedented lack of fertilizer allows us to estimate the value of fertil-

[1]This process of backing out unobservable from observable variables using the model's structure is often referred to as model "inversion".

izer for a whole country in an open economy model with international trade in agricultural as well as non-agricultural goods, and realistic features such as geography-driven comparative advantage in crop production, worker switching between occupations, non-homothetic preferences, and land inequality. Our paper is complementary to recent work by Bergquist et al. (2022) who develop a quantitative model of agricultural trade and propose a new approach for quantifying the effects of large-scale agricultural policy counterfactuals, including the provision of agricultural inputs. Our analysis is based on a unique trade policy experiment involving the removal of available fertilizer, in contrast to the existing literature, which typically focuses on subsidized input provision, where fertilizer take-up ultimately depends on varying farmer propensities for input usage (Duflo et al., 2011; Suri, 2011). The welfare losses we compute, therefore, stem solely from lack of fertilizer, and help quantify the overall value of fertilizer in an economy, informing the debate on fertilizer-related policy (such as fertilizer subsidies) through an approach that is distinct from but complementary to the existing literature. Our paper also complements Artuc et al. (2023), whose analysis of the consequences of conflict for global agricultural trade includes a simulation exercise in which Ukraine and Russia limit exports of fertilizer (and of multiple agricultural products). Such trade policy affects more countries than ours, but each country suffers less because it can substitute its imports toward unaffected suppliers to some extent.

Our theoretical model borrows tools from the burgeoning literature on quantitative spatial models with a focus on agriculture (Costinot et al., 2016; Aggarwal et al., 2022; Carleton et al., 2023; Farrokhi and Pellegrina, 2022; Nath, 2023; Pellegrina, 2022; Sotelo, 2020). We combine geographically differentiated agricultural and manufacturing sectors, which are connected through their competition for labor inputs, thus allowing shock transmission across sectors while still maintaining rich and tractable model mechanics. To such a quantitative trade model, we add a class of non-homothetic preferences (Price-Independent Generalized Linear, or PIGL) borrowed from the literature on structural transformation (Eckert and Peters, 2023; Boppart, 2014). Following Eckert and Peters (2023), we assume a specific family of income distributions and derive closed-form expressions for aggregate demand and indirect utility, which we then use to compute the equilibrium of a spatial model with multiple heterogeneous regions. However, our model has two types of agents (workers and farmers) while theirs only features workers and, significantly, we assume that the distribution of farmer income is log-normal, thus showing that the convenient aggregation properties of PIGL preferences are not restricted to Frechet distributions.

The natural experiment of Sri Lanka's fertilizer import ban relates our work to the literature on the impacts of non-tariff measures (NTMs).[2] Most studies have focused on a subset of NTMs – regulatory measures or anti-dumping measures – often bundling several measures and generally

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

estimating impacts on trade flows of products directly subject to the NTMs (see e.g., Ederington and Ruta (2016) for a survey). Import prohibitions have rarely been studied, except in a few case studies focusing on bans of imports of certain agricultural products from certain origin countries to prevent animal disease outbreaks.[3] Two recent studies that are more related to our work, despite focusing on import licenses rather than import bans, study the impact of discretionary import approvals in Argentina that varied with macroeconomic conditions to avoid worsening the current account balance on firm imports and import prices (Atkin et al., 2022) and on downstream responses by exporting firms (Bernini and García-Lembergman, 2022). Our paper's contribution to this literature is two-fold. On the one hand, we are the first to be able to rigorously estimate the causal impacts of an import ban – of a crucial intermediate input – on product- and firm-level trade flows for the input itself but also for the downstream exports dependent on the input by exploiting a natural experiment. On the other hand, we add to the only two papers to our knowledge – Atkin et al. (2022) and Dhingra et al. (2023) – that incorporate a non-tariff measure in a quantitative general equilibrium spatial and trade model, which allows us to measure its impacts on domestic outcomes such as agricultural production, productivity, and welfare. Our findings are entirely novel to the NTM literature: we show an NTM-driven decline in imports of fertilizer that ultimately results in a decline in agricultural exports of much larger value, thus jeopardizing the hidden motivation of the NTM of saving foreign exchange.

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

2 Background on the Fertilizer Import Ban in Sri Lanka

On May 6, 2021, the Sri Lanka government abruptly banned the imports of chemical fertilizers with the aim of transitioning the country's agricultural sector into organic farming. The use of agro-chemicals was linked to widespread chronic health problems and adverse environmental impacts across the country. Although the future president Gotabaya Rajapaksa had mentioned in a 2019 election campaign manifesto "Vistas of prosperity and splendor," the aim to provide Sri Lankans with food without harmful chemicals, the imposition of the agro-chemicals import ban

[3] See e.g., Felt et al. (2011) on a Japanese ban on imports of Taiwanese pork; 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.

came overnight unexpectedly.[4] This drastic non-tariff measure was proposed as necessary for health and environmental reasons but the most pressing reason for its introduction was the shortage of foreign currency (to save the USD 400 million the country was spending every year on fertilizer imports). Sri Lanka faced a large current account deficit and balance of payment crisis exacerbated by the adverse shock that Covid-19 imposed on the country's main source of foreign exchange earnings, tourism. To address these external challenges the government of Sri Lanka had already been imposing bans on the imports of non-fertilizer products since March 2020 (USAID, 2021).[5]

Importantly, the president attempted to push this shift into organic farming in a single agricultural season rather than over several years and without providing farmers with the adequate support and training for the shift. Proponents of organic farming in Sri Lanka such as Thilak Kariyawasam, president of the Lanka Organic Agriculture Movement, were appalled at the president's decision arguing "The president's committee of advisers to implement the new agricultural policy had no knowledge of organic farming" that 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." [6] Since domestic production of organic fertilizer was insufficient and imports of organic fertilizer (or of nutrients needed to produce them) were difficult, organic fertilizers were in short supply and very expensive. Some farmers ended up using low-quality fertilizers obtained on the black market.

More generally, agricultural specialists argue that in terms of providing the nutrients needed for crop growth, organic fertilizer can only be a supplement and not a substitute for chemical fertilizer (even if over the long-run organic fertilizer can help reduce the demand for chemical fertilizer).

Quantifying the economic consequences of the ban on agro-chemical imports is the focus of our paper. Preliminary evidence in FAO and WFP (2022) suggests the ban on chemical fertilizer imports had dramatically adverse effects on agricultural production (unfavorable prospects for the 2021/22 Maha season crop due to nutrient deficiency). This resulted in strong food price inflation and reduced food security, leading to heavy protests by farmers and a general upheaval in the country that eventually induced the government to lift the ban on chemical fertilizer imports later in 2021. After the end of the import ban, the international prices of chemical fertilizers were at all-time highs due in part to the onset of the Russia-Ukraine war (initiated in February 2022). Given Sri Lanka's limited capability to import due to low foreign reserves and a strong depreciation of the Sri Lanka rupee during 2022, this led to a continued shortage of agro-chemicals, with adverse

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

[5]The stated policy goal of the import bans was in fact ambiguous, with policy statements outlining objectives ranging from minimizing foreign exchange leakages to import substitution (Wijesinghe et al., 2023).

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

consequences for agricultural production and food security in 2022.[7] Even though the fertilizer import ban was short-lived it was imposed at a very unfortunate time, right before the Maha season harvesting.

As for the organic versus inorganic fertilizers, the general idea of fertilizer additions is to supply the nutrient needs (mostly N, P, K) of plants that are not sufficient (naturally occurring) in the soil. This is regardless of the source (organic vs inorganic) as they both supply this need. However, there are some important differences between them. Organic fertilizers tend to be slow-releasing and may require certain amounts of heat and moisture to extract the nutrients (N, P, K) during the time of plant need because some breakdown needs to occur for these nutrients to be released. Also, the plant need is met with almost the right timing and amounts so leaching (removal of excess nutrients from the plant/soil system) is minimized. In contrast, inorganic fertilizers provide rapid input of nutrients (N, P, K) because they are in plant uptake-ready form so the use is rapid but this could come at the expense of leaching, especially if not timed right.[8] Both types of fertilizers are capable of meeting the plant's needs but the amounts and timing matter a lot.

3 Data

Our reduced-form and model-driven analyses rely on several data sources. In this section, we list and describe the main data sets we use in this paper.[9]

3.1 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 products (defined at the HS 8-digit level - HS8 in what follows) subject to an import ban as well as the start and end date of such bans.[10] Specifically, we digitized the gazettes and then carefully inspected manually all the entries so as to create a database with HS8 products and the start and end dates of import bans. The ban on fertilizer imports started in May 2021, and most bans were eliminated by July 2021 while the rest was eliminated in November 2021.

Appendix Table A3 shows the list of HS8 product codes corresponding to fertilizers that were subject to the import ban, and Appendix Figure A2 shows the timing of the import ban. Combining

[7] Sri Lanka's Central Bank devalued the rupee by up to 15% in March 2022. Over the year 2022 the rupee devalued by nearly 45% against the USD.

[8] Leaching occurs when water removes water-soluble nutrients out of the soil, by runoff or drainage. Leaching is an environmental concern to the extent that chemical fertilizers find their way into water bodies.

[9] This section covers only data sources pertaining to Sri Lanka. For ancillary data on the rest of the world, see Appendix A.3.

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

this information with the high-frequency trade data described next in Section 3.2, Appendix Table A2 shows that, across fertilizer and non-fertilizer categories, the products whose imports were banned accounted for 16%-19% of total imports in 2017-2019, before the bans were introduced. This share declined to 13% in 2020 and 11% in 2021, suggesting the bans may have been at least partially effective.

3.2 International Trade

We use novel high-frequency granular trade data for Sri Lanka obtained from the S&P Global Market Intelligence's Panjiva data platform, covering the universe of export and import transactions of Sri Lanka between January 2017 and October 2022. The data is from bills of lading (BoL) for exports and imports. Crucially, the data includes information on the shippers (exporters) and the consignees (importers) for each shipment: their names and addresses (in most cases). It also includes HS8 codes for the products. We use the data at the level of the Sri Lankan exporting or importing firm-HS8 product-partner country-month-year, with information on value, weight, and quantity (along with unit of measurement for the quantity). The partner country variable is cleaned and mapped to the corresponding ISO 3-digit code. In some analyses, we aggregate the data to the HS8 product-month-year level by summing across firms and partner countries.

Exporting and importing firms in Sri Lanka are identified by their names, as the BoL data does not require the inclusion of tax identification numbers.[11] The firm names are noisy, so we use machine learning and text analysis techniques to clean them and assign a unique firm identifier to each name and thus construct a panel of firms in Sri Lanka engaged in exports or imports.[12] Export and import values are measured in United States (US) dollars. The data is subject to a series of cleaning procedures as detailed in Fernandes et al. (2016).[13] Our analysis focuses on some types of

[11]While Panjiva adds a firm identifier to the Bill of Landing data that Flaen et al. (2019) use for US data, we found such an identifier to be extremely unreliable for Sri Lanka's 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 IDs.

[12]Details on the fuzzy matching algorithms used are provided in Appendix A.1.

[13]The cleaning procedures consist of dropping the very few observations for which trade value is 0 or missing, product codes are missing or not part of the HS 2017 revision list, partner country is missing or is Sri Lanka (this may indicate transactions made with Special Economic Zones but this cannot be systematically assessed), so the information cannot be used. The quality of the data 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 the total yearly exports and imports for Sri Lanka obtained from COMTRADE/WITS (World Integrated Trade Solution) and from the Central Bank of Sri Lanka. The ratios to COMTRADE/WITS in 2017-2021 range from 97%-100% for exports and 83%-94% for imports. To obtain our final datasets, we eliminate exports and imports of oil (HS chapter 27), which is generally not well captured in transaction-level customs datasets. Moreover, to focus on true commercial export and import transactions by firms, we also drop from the datasets transactions of currency paper notes (HS 490700) that are conducted only by Sri Lanka's Central Bank and transactions of large machinery that is entering and exiting Sri Lanka temporarily for construction projects such as dredgers (vessels) (HS 890510, HS 890190) and transactions of arms and ammunition

imports, namely fertilizer products, whose list of HS8 codes is shown in Appendix Table A3, and some types of exports, namely agricultural products, whose HS8 codes belong to the HS 2-digit chapters 06 up to 24.

Among other uses, this trade data allows us to compute the average price of the fertilizers being imported by Sri Lanka over time. Figure 1 displays this average price for each year between 2017 and 2022, showing a clear and marked increase in the post-fertilizer ban period (2021 and 2022).[14]

Figure 1: Average fertilizer prices in Sri Lanka

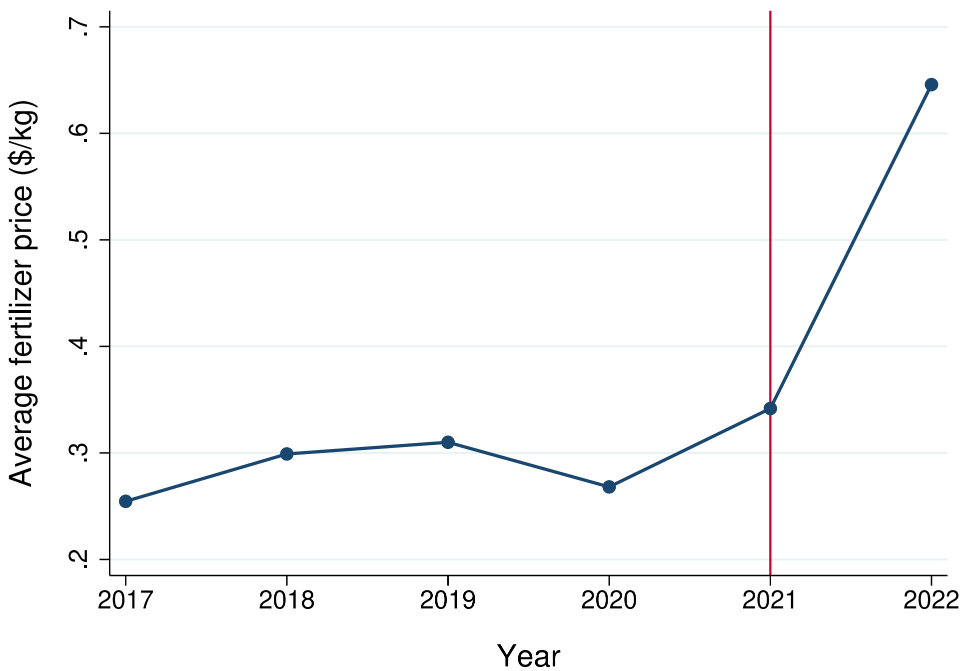

Notes: The figure shows the average fertilizer price in Sri Lanka for each year between 2017 and 2022. The average price in each year is defined as the total value of fertilizer imports (measured in US dollars) divided by the total quantity of fertilizer imports (measured in kilograms). All fertilizers (whether or not banned) are included in the calculation. Fertilizers are defined as products whose HS8 code belongs to HS 2-digit chapter 31. The red line in 2021 marks the year fertilizer import bans were introduced.

3.3 Fertilizer Use

We use data provided by Sri Lanka's National Fertilizer Secretariat to construct chemical fertilizer usage in the production of various agricultural crops. The crop-level fertilizer requirements are

(HS chapter 93).

[14]While Figure 1 displays average fertilizer prices in Sri Lanka by year, we also compute average prices by quarter, which we use to help estimate the elasticity of substitution across origins in Section 6.1.1.

reported as kilograms of fertilizer required per hectare of cultivated land for 23 crops: paddy, tea, rubber, coconut, cocoa, cloves, nutmeg, cinnamon, pepper, maize, kurakkan, undu, green gram, cowpea, horse gram, sesame, groundnuts (peanuts), soybeans, onion, chili pepper, potatoes, vegetables, and fruits (see Appendix A.2 for details). We validate and complement this information with data from the International Fertilizer Development Center and the International Fertilizer Association. Sri Lanka's National Fertilizer Secretariat data also includes estimates of fertilizer prices and of each crop's total cultivated area, which we combine with fertilizer requirements to compute the total country-level expenditure on fertilizer for each crop.[15]

3.4 Remote Sensing Area and Yield

We use remote sensing estimates of rice cultivated area and rice yield at a very granular level. The methodology relies on remote sensing data, namely an expert-based image classification algorithm on satellite observations (from Landstat and Sentinel-2) that have been enhanced to isolate the rice signal based on the premise that paddy rice has specific growing conditions in Sri Lanka.[16]

To estimate rice yields, a statistical model is used that correlates district-level rice yields from government statistics (described in Section 3.6) 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) and the pixel-level paddy rice yield estimates across two decades (2000-2022). Further details on the methodology are provided in Ozdogan et al. (2023).

These remote-sensing rice area and yield estimates, aggregated to the level of the divisional secretariat, are shown to be highly consistent with the area and yield measures reported in government statistics (for area from survey-based crop cutting experiments and for yield from production statistics). The remote-sensing rice yield estimates have the advantage of providing a spatially highly granular view of rice cultivation that allows a detailed assessment of the heterogeneous effects of the fertilizer import ban within Sri Lanka.

3.5 Household Survey

We use the 2016 and 2019 editions of Sri Lanka's Household Income and Expenditure Survey (HIES) for data on wages, employment, land ownership, land distribution, income, and consump-

[15]See Section 6.1.2 for details on how we use each crop's country-level fertilizer expenditure to estimate technological parameters.

[16]The presence of rice-growing areas is determined by factors related to climate, water availability, 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.

tion patterns. Conducted by the Department of Census and Statistics, 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.

We extract wage information from HIES 2019, which lists household members who "were paid employees during last four weeks / last calendar month". To compute each district's average wage, we take 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.[17] We multiply this average wage by 12 to express wages in annual terms.

We use information from HIES 2016 to count the number of workers and farmers in each district.[18] We define a "worker" as a person with a main occupation that paid her above-zero employment compensation in the last month. To compute the number of farmers, we first count the number of households in each district with positive landholdings.[19] We then estimate the number of farmers in each district by multiplying the number of farming households by the average number of economically active agents per household.[20] Finally, we increase the resulting numbers of workers and farmers by the rate of population growth in Sri Lanka from 2016 to 2019 (2.8%) to estimate the corresponding numbers for year 2019, our baseline period for model estimation (see Section 6).

We also use data from HIES 2019 on household size (all members regardless of age), income, and expenditure patterns, in particular the fraction of income spent on food, for our parameter estimation in Section 6.1.1. 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. Information on "income by chance or ad hoc gains" is also for instrumental variables estimation in Section 6.1.1.

3.6 Agricultural Production and Cultivated Area

We obtain data on the production and cultivated area of rice in each district from the "Paddy Statistics" of Sri Lanka's Department of Census and Statistics for all growing seasons ranging from the 2012-13 Maha season to the 2022 Yala season.[21] Crucially, this time frame includes the 2019-

[17] Household survey files provide household-level statistical weights. All averages and standard deviations we compute using household survey data are weighted using these statistical weights.

[18] While HIES 2016 does not cover the universe of the population, it is possible to use its representative statistical weights to estimate total counts.

[19] We only count the land that is owned by the household (as opposed to leased, rented, etc.) and that does not have housing built on top of it.

[20] 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 and farmer households are similar in terms of household size, structure, and member participation in the economic sustenance of the family.

[21] The "Paddy Statistics" website is: http://www.statistics.gov.lk/Agriculture/StaticalInformation/Paddy$_{S}tatistics$.

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 crops (maize, groundnuts, potatoes, onions, cinnamon, cloves, and others), we obtain data on production and cultivated area at the Divisional Secretariat (DS) level from the Department of Census and Statistics. It covers all seasons from Maha 2019-20 to Yala 2022. We aggregate the data from the DS level to the district level using the mapping between DS and district. Production is measured in kilograms and cultivated area is measured in hectares. For crops that have harvests in both Yala and Maha seasons (maize, rice, and groundnuts), we compute total production and total cultivated area in each year for each district by summing across the two seasons.

In addition to district-level production data, we need national production data for each crop to estimate the fertilizer coefficients in our model's Cobb-Douglas agricultural production functions (see Section 6.1.2). For each crop, we use total 2022 production data (in metric tons) from the Economic Statistics of Sri Lanka 2023 report, produced by the Department of Census and Statistics.[22]

3.7 Producer Prices

To estimate our model, we need data on producer prices for each crop and district. We use data from the Bulletin of Selected Retail and Producer Prices 2016-2019 (published in September 2020), from Sri Lanka's Department of Census and Statistics.[23] We extract average producer prices (measured in LKR per kilogram) in year 2019 for the following crops: coconut, potatoes, cinnamon, cloves, onions, soybeans, groundnuts, and maize.[24] Whenever the producer price of a crop is missing for a given district, we impute that district's price using the crop's average price across the other Sri Lankan districts.[25]

3.8 Input-Output Matrix

To help estimate labor and land coefficients in our model's Cobb-Douglas agricultural production functions, we use data from Sri Lanka's input-output tables for year 2000, produced by the Institute of Policy Studies (Amarasinghe and Bandara, 2005). Specifically, we extract information on

[22]We extract data for the following crops: cinnamon, cloves, coconut, groundnuts, maize, onions, potato, rice, and soy. The report is available at: http://www.statistics.gov.lk/Publication/Economic-Statistics-2023.

[23]The Bulletins of Prices can be found at: http://www.statistics.gov.lk/InflationAndPrices/StaticalInformation/Bulletinsgsc.tab=0. We use the 2016-2019 edition.

[24]Onions and rice have prices available for multiple varieties. We use the prices of red onions, and red rice (raw). Coconut prices refer to medium-sized coconuts and are not given in kilograms but in 100 units.

[25]Producer prices are missing for 42.5% of the region-crop pairs that have non-zero agricultural production. Nevertheless, the imputation procedure is unlikely to create meaningful biases because producer prices are similar across space: the coefficient of variation of (non-missing) producer prices across districts ranges from 0.009 to 0.058, depending on the crop, with the average being 0.032.

employee compensation and value added (at factor cost prices) for a variety of agricultural sectors. See Section 6.1.2 for details on how this data is used in estimation.

3.9 Potential Crop Yields

To estimate the elasticity of substitution across origins in our model's demand system, we use an instrumental variables approach that leverages variation in agricultural productivity across space due to exogenous geographic factors (such as climate and geology).[26] To do so, we use data on district-level potential attainable yields for each of seven crops, which we obtain from the Global Agro-Ecological Zones (GAEZ) project of the Food and Agriculture Organization of the United Nations (FAO), a data source that has often been used in the literature on agriculture and trade (Costinot et al., 2016; Bustos et al., 2016; Sotelo, 2020; Farrokhi and Pellegrina, 2022).[27]

We extract potential yield data (in kilograms per hectare) at the grid cell level from FAO-GAEZ and then aggregate up to the district level. GAEZ offers different yield estimates depending on the irrigation system (rainfed or irrigated) and on the climate conditions (e.g. historical conditions, various climate change scenarios, etc.). 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.

4 Stylized Facts on Fertilizer Imports, Agricultural Production and Exports

In this section, we present three stylized facts related to the imports of fertilizer, the exports of agricultural goods that utilize fertilizers with varying degrees of intensity, and the variation in fertilizer intensity across districts in Sri Lanka depending on the crops they produce.

4.1 Fact 1: Fertilizer imports declined after the ban

In this section, we show that the ban on the imports of chemical fertilizers led to a decrease in fertilizer imports relative to non-banned fertilizers and other non-banned categories of products.

[26] See Section 6.1.1 for details on how FAO-GAEZ data is used in estimation.

[27] The seven crops we use are: coconut, groundnut, maize, onion, rice, tea, and white potato. We did not use cinnamon because it was not available on GAEZ.

4.1.1 Empirical Specification

To understand the effects of the fertilizer ban on imports, we estimate dynamic difference-in-differences regressions using import data at the HS8 product-month level, with the treated group consisting of fertilizer products whose imports were banned after April 2021 and the control group consisting of products whose imports were not banned, which includes both non-banned fertilizers and all other types of non-banned products. We compare the imports of fertilizers in the treatment group over time with their control group counterparts, using the state-of-the-art difference-in-differences methods proposed by De Chaisemartin and d'Haultfoeuille (2022). The advantage of using this research design over other difference-in-differences designs is the fact that this research design allows units to switch in and out of treatment.[28] This is especially important for our study as the import ban was periodically lifted and later re-imposed on certain products. Our first estimating equation for imports is given below:

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

where c designates an HS8 product, t designates a time period (a month-year), ω_t and ω_c are time and product fixed effects, respectively, 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 τ periods before period t.[29] The outcome variable y_{ct} can be either a dummy indicating non-zero imports of product c in period t (i.e., the extensive margin) or the log of the U.S. dollar value of product c imports in period t (i.e., the intensive margin).[30]

Our dynamic difference-in-differences estimator (De Chaisemartin and d'Haultfoeuille, 2022) recovers the evolving effects of the import bans over time. The estimating sample period ranges from January 2019 to September 2022 but for presentational purposes we will show only the coefficients for the four months prior and then twelve months subsequent to the import ban on fertilizers. We exclude from the analysis any products whose imports were banned at any point prior to May 2021.

4.1.2 Results

Figure 2 plots the results of estimating equation 1 and shows a clear decline in imports right after the fertilizer ban was imposed, with the average probability of nonzero imports for banned products

[28] See Callaway (2022) for a complete review of difference-in-differences methods.

[29] Note that ban_{ct}^{τ} is always equal to zero for products whose imports were never banned.

[30] To define the extensive margin outcome we expand our original data on import flows at HS8 product-month level by adding zero import flows for every month between January 2017 and October 2022 when imports of a given HS8 product are not reported.

Figure 2: Dynamic Ban Effects on Imports

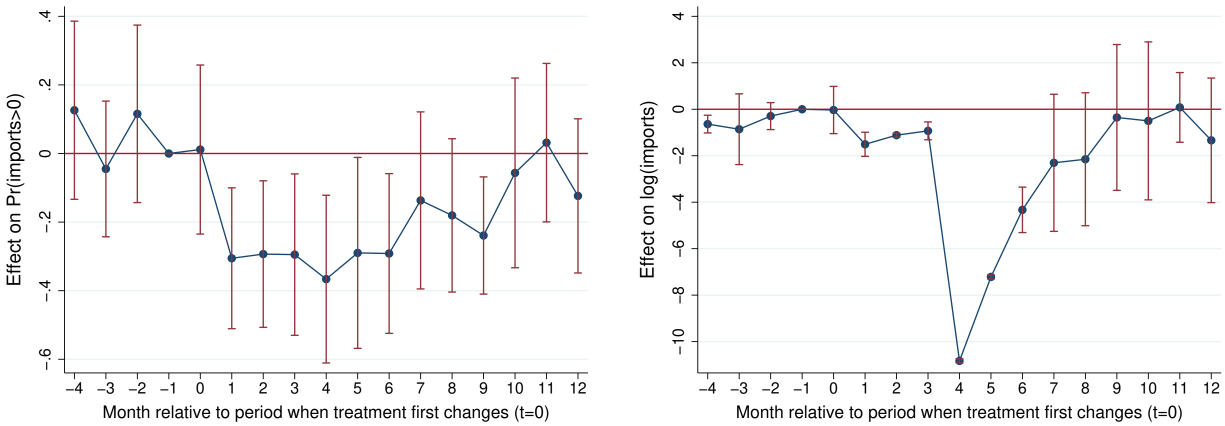

(a) Extensive Margin

(b) Intensive Margin

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of import bans on the both the extensive and intensive margins of imports, relative to the last month before the ban takes effect, using an HS8 product-level panel and methods described in section 4.1.1 to isolate the ban effects. The sample covers the period between January 2019 and September 2022 (but coefficients are only reported from January 2021 to May 2022) 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 non-banned fertilizers as well as other products. The y-axis shows the effects of fertilizer import bans on the probability of nonzero imports (panel 2a) or on the log value of imports (panel 2b), where negative values correspond to a decrease in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.56 (panel 2a) and 0.0001 (panel 2b).

dropping by 31 percentage points in the second month of the ban (panel 2a). Along the intensive margin, negative effects peaked four months after the ban was introduced, with imports falling by 99.998% (panel 2b), and remaining low for the next several months.[31]

We also estimate equation 1 for a sample including only fertilizer products, in which we compare how the imports of banned versus non-banned fertilizers evolve over time as the ban is introduced. Figure 3 presents results that are generally similar to those of Figure 2, albeit less precise. These results indicate that, relative to non-banned fertilizers, the average probability of nonzero imports for banned fertilizers declined by 31 percentage points one month after the fertilizer ban was first imposed, although the difference in probability is not statistically significant in subsequent months. Chemical and other types of fertilizer are often used as complements in agricultural

[31] The coefficient for month four in Figure 2b equals -10.8. Given the log specification, we can recover the corresponding percentage effect on imports through the expression: $(\exp(-10.8) - 1) \times 100\% \simeq -99.998\%$.

Figure 3: Dynamic Ban Effects on Imports (fertilizer only)

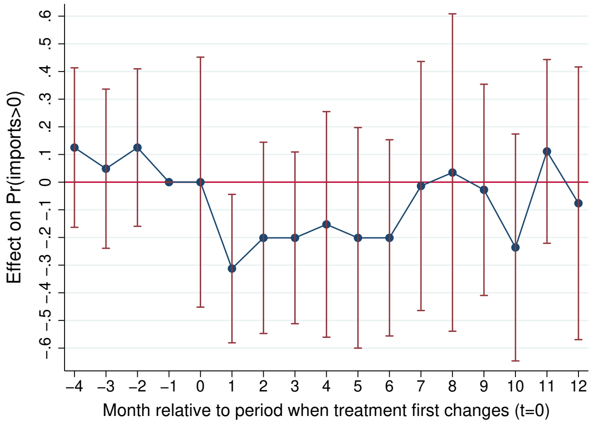

(a) Extensive Margin

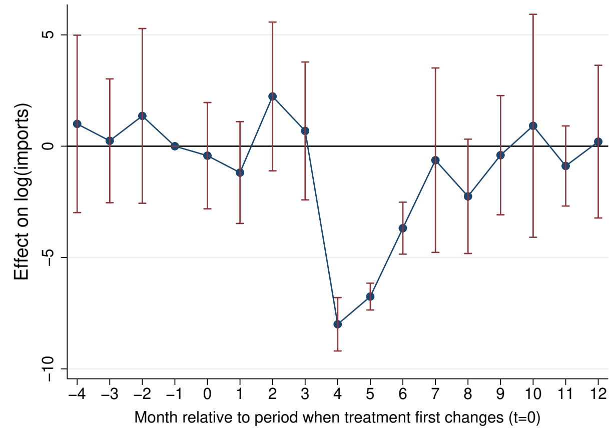

(b) Intensive Margin

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of import bans on the both the extensive and intensive margins of imports, relative to the last month before the ban takes effect, using an HS8 product-level panel and methods described in section 4.1.1. The sample covers the period between January 2019 and September 2022 (but coefficients are only reported from January 2021 to May 2022) and includes only HS8 products corresponding to fertilizer. 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 non-banned fertilizers. The y-axis shows the effects of fertilizer import bans on the probability of nonzero imports (panel 3a) or on the log value of imports (panel 3b), 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 3a) and 0.925 (panel 3b).

production. Therefore, banning the imports of some chemical fertilizers may also have negatively affected other fertilizers used in production. 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.

Next, we investigate potential mechanisms through which the strong negative effects of the ban on fertilizer imports operate. We decompose the ban's effect on import values (presented in Figures 2 and 3) into the contributions of price effects and quantity effects. To do that, we estimate equation (1) using as outcome variable y_{ct} the log of the imported quantity of product c in period t (panels A3a and A3c) or with the log of the average import price of product c in period t (panels A3b and A3d). The results presented in Appendix Figure A3 show that import declines are entirely driven by decreases in imported quantities. In fact, prices effects, when statistically significant, tend to be *positive*, which is expected in a context of a policy that radically reduced fertilizer supply.

Finally, we conduct two empirical exercises to establish the robustness of our estimated effects of the ban on fertilizer import values. First, since fertilizer bans happened in a specific part of the year, one concern with our results is that they could be confounded by seasonality in fertilizer imports. To address this concern we estimate equation (1) using deseasonalized versions of our outcome variables instead of the original variables (y_{ct}).[32] The results presented in Appendix Figure A6 strongly resemble the original results from Figures 2 and 3, thus allaying potential seasonality concerns. Second, since the import bans happened when the Covid-19 pandemic was still an ongoing concern for the global economy, another concern is that our estimates could be partially driven by the effects of the pandemic, which we would then be misattributing to the ban. However, this would only introduce bias in our estimates if the pandemic affected imports of banned and non-banned fertilizers to different degrees, which we do not have a strong reason to expect. In any case, to allay this concern, we implement a falsification test in which we estimate equation (1) keeping the same assignment of products to treatment/control groups but redefining the treatment period to start in March 2020 (when Covid-19 arrived) instead of May 2021. Results in Appendix Figure A8 show that, if anything, fertilizers whose imports would eventually be banned in mid-2021 actually saw an *increase* in imports relative to other products in the early months of the pandemic. This allows us to argue that our main results are not driven by the effects of the pandemic.

4.2 Fact 2: Rice production declined significantly after the ban on fertilizer imports, with regional heterogeneity

We show that, right after the fertilizer ban, rice production in Sri Lanka declined significantly and rice imports, in turn, increased. We focus on rice because it is Sri Lanka's primary food crop, and about 40% of the total arable land in Sri Lanka is under rice production.[33] Additionally, we are able to use very granular pixel-level estimates of rice yield using remote sensing, as crop type mapping and yield can be better predicted for rice than for other crops using remote sensing tools because the water and vegetation index is more visible for rice (Dong et al., 2016).

Panel a) in Figure 4 plots the yearly average yield of Maha paddy in kilograms (kg) per hectare for the past 10 years. We focus on the Maha season because this is the main season for paddy cultivation in Sri Lanka, running from September to March of the following year. The second season, Yala, runs from May to August.[34] Panel a) in Figure 4 shows that Maha rice yields declined precipitously in 2022 from a high of 4,500 kg/hectare in 2020 to below 3,000 kg/hectare. This represents a more than 30% decline in yield in 2022 relative to the average of the previous nine years. Panel b) in Figure 4 shows Sri Lanka's monthly rice imports from January 2020 to October

[32] Both outcome variables (extensive and intensive margins) are deseasonalized by subtracting their HS8 product-month-specific averages in the 2017-2019 period, before the fertilizer ban.

[33] https://rb.gy/jy8gq

[34] https://rb.gy/4tvo9

2022. Right until the import ban, rice imports were negligible as Sri Lanka had achieved self-sufficiency in rice production. From late 2021 onwards rice imports surged up to January 2022, remaining elevated throughout the rest of the year.

In Figure 5, we show rice yields (in kg per hectare) at the divisional secretariat level, where our yields are obtained using state-of-the-art remote sensing methods that make use of the water and vegetation index of rice, as described in Ozdogan et al. (2023). While obtained independently using remote sensing tools, these estimates have been verified against district-level statistics from the Sri Lankan government and the spatial variance in yield has been matched against plot-level data from Sri Lanka. The left panel of Figure 5 shows the divisional secretariat-level rice yield for the time period between September 2020 and March 2021 (Maha season), while the right panel map shows the same but for September 2021 to March 2022, which follows the introduction of the fertilizer import ban in May 2021. A comparison across the left and right panels shows that, with a handful of exceptions, paddy production fell in almost all regions of the country, with spatial heterogeneity in the degree of yield reduction. These patterns seem more pronounced in the North and Western parts of the country, including major rice-producing districts such as Ampara, Anuradhapura, Batticaloa, and Polonnaruwa, among others. Ozdogan et al. (2023) 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.

Figure 4: 10-year time-series of yield per hectare of Maha rice in Sri Lanka

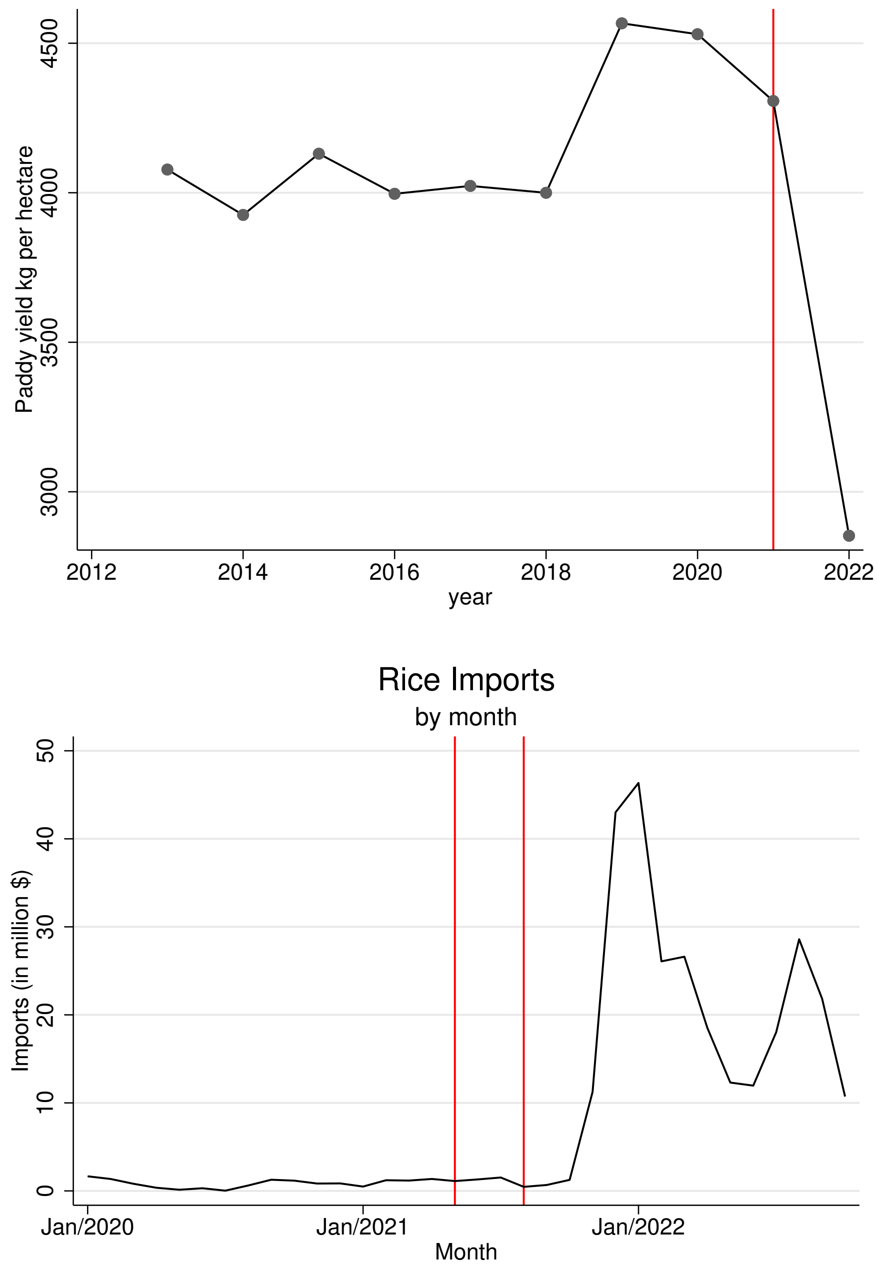

Notes: The top panel shows the Maha rice yield (in kg per hectare) for the 10-year period 2012-2022. The Maha season for rice is from September to March of the following year. The red line in 2021 marks the year fertilizer import bans were introduced. The bottom panel shows the monthly imports of rice. The first red line in May 2021 marks the beginning of the ban, and the second red line in August 2021 marks the end of the ban. The source of this data is the Department of Census and Statistics, Sri Lanka. The graph looks very similar if we instead use our remote-sensing estimates.

Figure 5: Yield per hectare of Maha rice in Sri Lanka

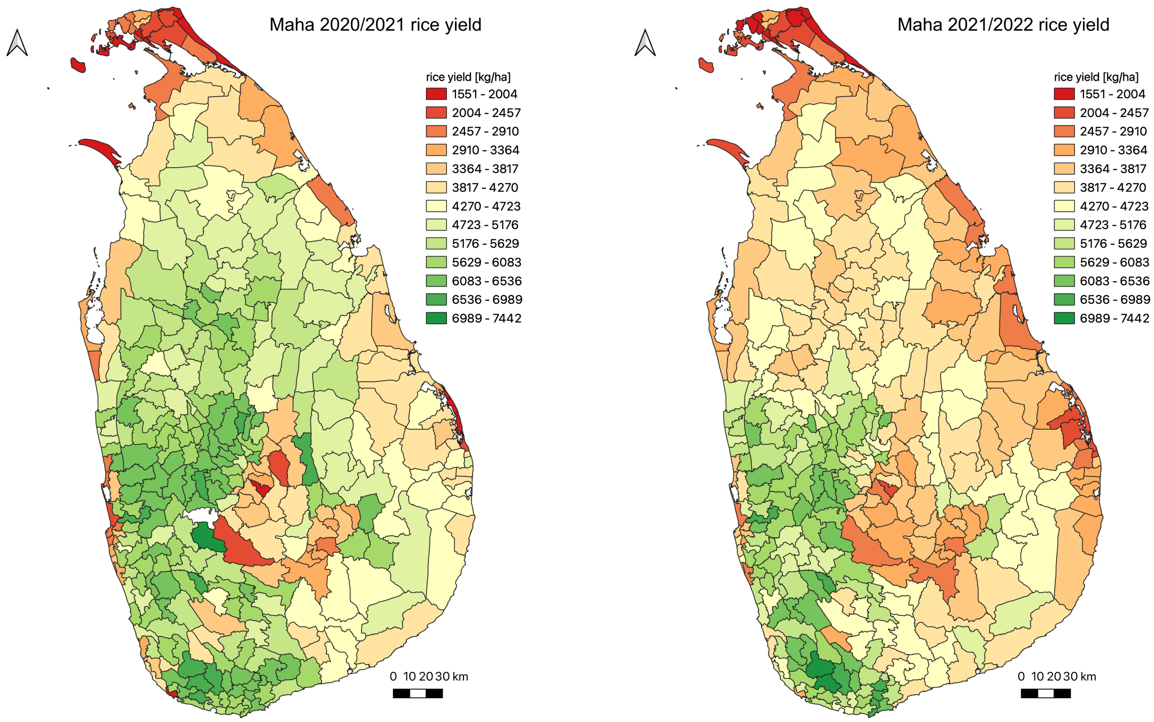

Notes: The left panel shows the Maha rice yield (in kilogram per per hectare) in the 2020-2021 seasons, and the right panel shows the Maha rice yield (in kilogram per hectare) in the 2021-2022 season. The Maha season for rice runs from September of a year to March of the following year.

4.3 Fact 3: Agricultural exports declined after the ban on fertilizer imports

Facts 1 and 2 show that fertilizer imports declined, rice imports increased, and production declined after the fertilizer import ban was implemented. In this section we show that agricultural exports declined following the ban.

4.3.1 Empirical Specification

To assess the fertilizer ban's effects on the export performance of agricultural firms, we consider the sample of firms who, during the 2017-2019 period (before the 2020 wave of import bans was imposed), exported at least one of the 23 crops for which we have fertilizer intensity data. To understand which firms required more fertilizer for production before the ban and were thus more adversely affected by the ban, we define an index of firm-level fertilizer usage as:

$$U_f = \sum_{v=1}^{23} FIC_v \times \left(\frac{X_{fv}^{2017-2019}}{\sum_{v'} X_{fv'}^{2017-2019}} \right), \tag{2}$$

where v indexes crops, f indexes firms, FIC_v is fertilizer intensity (in kilogram per hectare) of crop v, and $X_{fv}^{2017-2019}$ is total exports of crop v by firm f between 2017 and 2019. Ranking all firms by fertilizer usage U_f, we compute the third quartile of fertilizer usage U_{p75}. Then, the estimating equation for the impact of import bans on exports can be written as:

$$\ln(X_{ft}) = \sum_{\tau \neq -1} \gamma_\tau \times \mathbb{1}\{t = T_0 + \tau\} \times \mathbb{1}\{U_f > U_{p75}\} + \Omega_t + \Omega_f + \mu_{ft} \tag{3}$$

where f indexes firms, t indexes quarters, X_{ft} is agricultural exports of firm f in quarter t, Ω_t and Ω_f are quarter and firm fixed effects, respectively, T_0 is the first quarter when there was a fertilizer import ban (i.e. T_0 is quarter 2 of 2021), and μ_{ft} is an error term. Unlike the import regressions, which rely on a monthly time dimension, here we rely on a quarterly time dimension. The reason for this is that firms may have an inventory of fertilizers which are not necessarily applied to agricultural production in the month they are imported. A quarterly analysis allows the effects on exports to materialize after some time, not necessarily in the same month the ban was implemented. As in Fact 1, we use the De Chaisemartin and d'Haultfoeuille (2022) estimator to allow for the fact that bans on some fertilizers may have switched on and off.

The advantage of this research design is that we directly use information on a firm's agricultural exports and agricultural products' fertilizer requirements to compute the firm's potential exposure to fertilizers. The disadvantage is that the control group may be imperfect because all agricultural exporters might have been affected to some degree. Therefore, we also inspect the effects on agricultural exports using a second, product-centric (as opposed to firm-centric) research design.

In this second design, we divide export products into treatment and control groups based on the degree to which they require fertilizer as an input in their own production process. To estimate these requirements, we use import-export data from single-export-product firms in the 2017-2019 period. Formally, for each pair of HS8 export product c and HS8 import product d, we estimate the requirements of input d in the production process of output c as:

$$a_{cd} = \frac{\sum_y \sum_{f \in S_c^y} M_{f,d}^y}{\sum_y \sum_{f \in S_c^y} X_f^y} \tag{4}$$

where $y \in \{2017, 2018, 2019\}$ indexes years, f indexes firms, S_c^y is the set of firms who exported product c (and no other products) in year y, $M_{f,d}^y$ is imports of product d by firm f in year y, and X_f^y is exports of firm f in year y. In this process, we obtain the input-output (IO) matrix corresponding to 931 agricultural products, each using on average 37.3 different products as inputs.

Using the requirements, we define a variable sev_{ct} to capture how severely an export product c was affected by fertilizer import bans in quarter t:

$$sev_{ct} = \sum_{d \in \mathbb{F}} a_{cd} \times fctban_{dt} \tag{5}$$

where $fctban_{dt}$ is the fraction of months within quarter t during which imports of product d were banned, and $\mathbb{F}$ is the set of all fertilizer products.[35] We then define a treatment dummy T_{ct} that indicates whether a product c was facing particularly severe import bans on its inputs in a given quarter t:

$$T_{ct} \equiv \mathbb{1}[sev_{ct} > sev_{p75}] \tag{6}$$

where sev_{p75} is defined as the third quartile of the severity variable sev_{ct} in the sample *conditional on* $sev_{ct} > 0$. Finally, to estimate the dynamic effects of fertilizer import bans on product exports, we use again the De Chaisemartin and d'Haultfoeuille (2022) estimator:

$$\ln(X_{ct}) = \sum_{\tau \neq -1} \gamma_\tau \times PT^{\tau}_{ct} + \Omega_t + \Omega_c + \mu_{ct} \tag{7}$$

where c indexes products, t indexes quarters, X_{ct} is exports of product c in quarter t, Ω_t and Ω_c are quarter and product fixed effects, respectively, and μ_{ct} is an error term. The variable of interest is PT^{τ}_{ct}, a dummy variable indicating whether product c was treated for the first time τ periods before period t.

4.3.2 Results

Figure 6 presents the results from estimating equation (3) to obtain the ban effects on firms' agricultural exports. We find that, three quarters after the fertilizer ban was implemented, agricultural exports declined for firms facing higher fertilizer exposure before the ban, compared to firms facing lower fertilizer exposure. The magnitude of the estimated effects corresponds to a reduction of 33% in firm agricultural export value.

[35] We define an HS8 product d as a "fertilizer" if it is listed in chapter 31 of the HS system. The variable $fctban_{dt}$ can assume values 0, 1/3, 2/3, or 1 depending on whether imports of product d were banned for zero, one, two, or three months within a quarter.

Figure 6: Dynamic Ban Effects on Firms' Agricultural Exports (High vs Low Fertilizer Intensity)

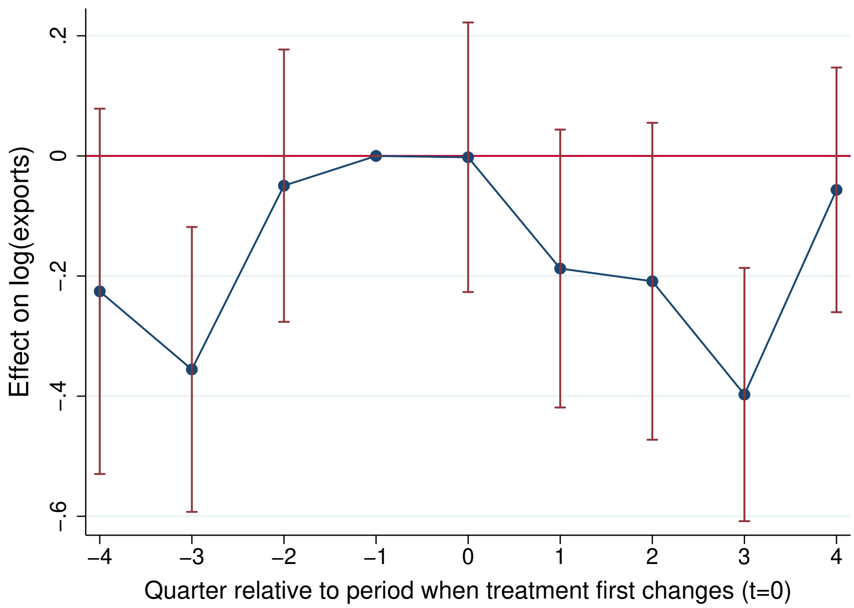

Notes: The figure shows De Chaisemartin and d'Haultfoeuille (2022) estimatesof the effects of fertilizer import bans on the value of firms' agricultural exports, relative to the last quarter before the ban takes effect, using a firm-level panel and methods described in section 4.3.1 to isolate the ban effects. The sample covers the period between quarter 1 of 2019 and quarter 3 of 2022 (but coefficients are only reported from quarter 2 of 2020 to quarter 2 of 2022) and includes all firms who exported at least one of 23 crops for which we have fertilizer intensity data during the 2017-2019 period. We compute each firm's fertilizer usage U_f as a weighted average of crops' fertilizer intensities with weights given by each crop's share in firm f's 2017-2019 export record. The treatment group consists of firms whose fertilizer intensity U_f 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 there were any fertilizer import bans in place. 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 0.015.

However, as discussed before, all agricultural firms were likely affected by the fertilizer import ban. To address this issue, we estimate the product-level equation (7). The estimates in Figure 7a show a 95% decline in the exports of agricultural products with high fertilizer input requirements (compared to their counterparts with relatively lower requirements) three quarters after the ban is implemented. Figure 7b shows the results change very little if we exclude from the sample any products that, during the sample period, faced import bans on some of their non-fertilizer inputs (but no ban on their fertilizer inputs).

Next, we examine the mechanisms for the adverse impact of the fertilizer ban on export values. We decompose the ban's effect on export values (which was presented in Figures 6 and 7) into the

Figure 7: Dynamic Ban Effects on Agricultural Exports (IO linkage)

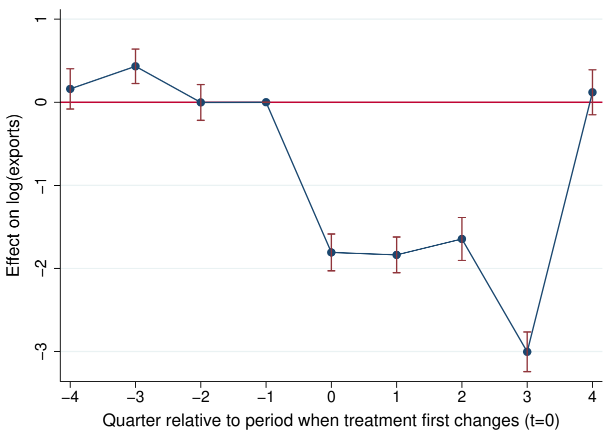

(a) All agricultural products

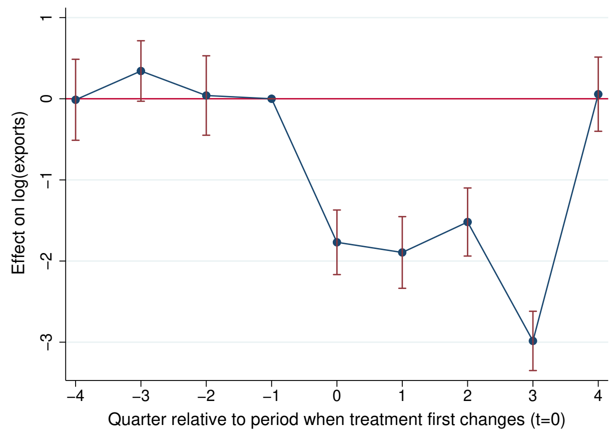

(b) Excluding products with banned non-fertilizer inputs

Notes: The figures show the De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of fertilizer import bans on agricultural product exports, relative to the last quarter before the ban takes effect, using an HS8 product-level panel and methods described in section 4.3.1 to isolate the ban effects. The sample covers the period between quarter 1 of 2019 and quarter 3 of 2022 (but coefficients are only reported from quarter 2 of 2020 to quarter 2 of 2022). In panel 7b, the sample excludes products for which some non-fertilizer inputs (but no fertilizer inputs) were banned during the period. The treatment variable is a dummy variable T_{ct} indicating whether the severity of the import bans on product c's fertilizer inputs in quarter t (sev_{ct}) is above the sample's 75th percentile. The not-yet-treated products serve as the control group. The severity of the fertilizer import bans faced by a product c in quarter t (sev_{ct}) is defined as the (adjusted) count of HS8 fertilizer products that were banned in quarter t, with adjustment terms given by each fertilizer's unit requirements in the production of product c. The y-axis shows the effects of fertilizer import bans on the log value of exports, where negative values correspond to a decrease in exports. The p-values of the joint significance tests for the placebo estimators are 0.00002 (panel 7a) and 0.147 (panel 7b).

contributions of price effects and quantity effects. Results are presented in Appendix Figures A4 and A5 for firm-level and product-level specifications, respectively. The product-level decline in exports three quarters after the start of the ban is entirely driven by quantity effects, since the price of treated products actually *increased* (relative to control products) 3-4 quarters into the ban.

Finally, we conduct a robustness check to the estimated effects of the fertilizer ban on exports. We address seasonality concerns by estimating firm-level equation (3) using a deseasonalized version of the outcome variable. The results, presented in Figure A7, are similar to the original results from Figure 6. Therefore, it does not appear that results are driven by seasonality confounders.

While these three stylized facts demonstrate the ban's impact on imports, agricultural production, and exports, it is hard to quantify the general equilibrium effect of the loss from the lack of access to fertilizer and the effects of the ban on households without a spatial model. This is because, while in general equilibrium, the whole economy is affected, the difference-in-difference estimates allow us to only infer the effect of lack of fertilizer on the directly affected products (treatment group that relies directly on fertilizer) relative to the control group that does not rely directly on fertilizer.

5 A Quantitative Spatial Model of Trade and Agriculture

This section presents a quantitative spatial trade model that will be used to make sense of the import ban effects shown above and to perform counterfactual exercises that allow us to estimate the ban's causal effect on relative prices, income, and welfare across different regions within Sri Lanka.

5.1 General Environment

In the model, the world is composed of R geographic regions indexed by i or n. There are $I = R-1$ regions within Sri Lanka, plus one additional region representing the rest of the world (RoW).[36] Each region i has a fixed population of size N_i which cannot move across regions.[37] The population is composed of two types of agents: N_i^W workers and N_i^F landowning farmers, with $N_i^W + N_i^F = N_i$.

There are two main economic sectors, agriculture (**A**) and manufacturing (**M**). Agriculture encompasses a finite set of K crops indexed by k. 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

[36] When bringing the model to the data, the regions within Sri Lanka are districts (the 25 second-level administrative divisions of Sri Lanka).

[37] 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.

an intermediate input in agriculture. All markets are perfectly competitive.

5.2 Production

This subsection presents the model's assumptions concerning production and trade costs in each of the three sectors of the economy: agriculture, manufacturing, and fertilizer.

5.2.1 Agricultural Production

Each region i has a mass of agricultural land area of total size L_{ik} that can only be used to grow crop k.[38] This land mass is 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:

$$q_{ik}(\omega) = T^A_{ik}(n_{ik}(\omega))^{\gamma^n_k}(f_{ik}(\omega))^{\gamma^f_k}(l_{ik}(\omega))^{\gamma^l_k} \tag{8}$$

where T^A_{ik} is a land productivity parameter and (γ^n_k,γ^f_k,γ^l_k) are Cobb-Douglas coefficients.

The assumptions above allow fertilizer intensity to vary across crops, since the parameter γ^f_k depends on k. This reflects the fact that some crops are more fertilizer-dependent than others, as shown earlier. Furthermore, since productivity parameter T^A_{ik} varies across regions, different regions may have different absolute and comparative advantages within agriculture.

Individual production decisions can be aggregated up 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 $c_{ik}(\omega)$ is the solution to the cost minimization problem:

$$c_{ik}(\omega) \equiv \min_{(n,l,f)\geq 0} w_i n + r_{ik}(\omega)l + p^f_i f, \text{ such that: } q_{ik}(\omega) \geq 1 \tag{9}$$

where production $q_{ik}(\omega)$ is given by equation (8), w_i is the wage in region i, p^f_i is the price of fertilizer in region i, and $r_{ik}(\omega)$ is the rental rate of plot ω. Solving this problem allows us to write the unit cost $c_{ik}(\omega)$ as a function of input and output prices:

$$c_{ik}(\omega) = \kappa_k[w_i^{\gamma^n_k}(r_{ik}(\omega))^{\gamma^l_k}(p^f_i)^{\gamma^f_k}]^{\frac{1}{S_k}}(T^A_{ik})^{-1} \tag{10}$$

[38] This assumption implies that land cannot be reallocated across crops, which is consistent with the short-term horizon of our empirical implementations of the model. Importantly, this assumption finds strong support in the Ozdogan et al. (2023) remote-sensing rice area estimates 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). Also see Appendix Section C.3, which shows there is no evidence of a post-ban reduction in cultivated rice area.

where $S_k \equiv \gamma_k^n + \gamma_k^l + \gamma_k^f$, and κ_k is a composite parameter.[39] Perfect competition then implies that the unit cost $c_{ik}(\omega)$ must equal the local crop price (denoted p_{ik}) in equilibrium.[40] This equality then allows us to solve for land rent $r_{ik}(\omega)$ as:

$$r_{ik}(\omega) = (T_{ik}^A)^{\frac{S_k}{\gamma_k^l}} \underbrace{(p_{ik})^{\frac{S_k}{\gamma_k^l}} (w_i)^{\frac{-\gamma_k^n}{\gamma_k^l}} (p_i^f)^{\frac{-\gamma_k^f}{\gamma_k^l}} (\kappa_k)^{\frac{-S_k}{\gamma_k^l}}}_{\equiv h_{ik}} \tag{11}$$

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

$$R_{ik} = (T_{ik}^A)^{\frac{S_k}{\gamma_k^l}} h_{ik} L_{ik} \tag{12}$$

However, due to the properties of the Cobb-Douglas production function, we also know that total payments to land must equal a share $\frac{\gamma_k^l}{S_k}$ of the agricultural revenue from crop k in region i:

$$R_{ik} = \frac{\gamma_k^l}{S_k} p_{ik} Q_{ik} \tag{13}$$

where Q_{ik} is the total physical production of crop k in region i. Combining equation (12) and (13), we can solve for production Q_{ik} as:

$$Q_{ik} = (\frac{S_k}{\gamma_k^l})(T_{ik}^A)^{\frac{S_k}{\gamma_k^l}} \frac{h_{ik}}{p_{ik}} L_{ik} \tag{14}$$

Intuitively, we conclude from equation (14) that equilibrium regional crop production depends positively on land productivity, land endowment, and crop prices, and depends negatively on the price of inputs (labor and fertilizer).

A crop k produced in region i can be sold either locally in i itself 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:

$$p_{ni,k} = \tau_{ni,k}^A p_{ik} \tag{15}$$

Finally, it will be convenient to also define an aggregate land rent variable R_n combining the land rent from all crops in region n:

$$R_i = \sum_k R_{ik} \tag{16}$$

[39] $\kappa_k \equiv S_k[(\gamma_k^n)^{\gamma_k^n}(\gamma_k^l)^{\gamma_k^l}(\gamma_k^f)^{\gamma_k^f}]^{-S_k^{-1}}$

[40] 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.

5.2.2 Manufacturing Production

We assume that manufacturing production uses labor as its sole input. Given a total number of n_i workers employed in manufacturing production in region i, output q_i^M is given by the following production function:

$$q_i^M = T_i^M n_i^M \tag{17}$$

where T_i^M is a productivity parameter.

As in agriculture, the manufacturing good produced in a region i can be either consumed locally or shipped to another region n, but the latter incurs 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 produced in i and offered for sale in n must be:

$$p_{ni}^M = \frac{\tau_{ni}^M w_i}{T_i^M} \tag{18}$$

5.2.3 Fertilizer Production

As in Farrokhi and Pellegrina (2022), we assume that fertilizer is simply given as a natural endowment (rather than being produced). Therefore, we can write the quantity of fertilizer Q_i^f in region i as follows:

$$Q_i^f = F_i \tag{19}$$

where F_i is the fertilizer endowment of region i.

We further assume that Sri Lanka is not endowed with any fertilizer.[41] Therefore, only the rest of the world (RoW) is capable of producing fertilizer (i.e. $F_i = 0$ if $i \neq$ RoW), implying that Sri Lankan agriculture depends on imports to fulfill its demand for fertilizer inputs. As in the other two sectors, shipping fertilizer from region i to region n incurs an "iceberg" trade cost $\tau_{ni}^f \geq 1$. In particular, changes in the fertilizer trade costs between RoW and Sri Lankan regions ($\tau_{i,RoW}^f$) can be used to represent a fertilizer import ban, and will thus play a critical role in our counterfactual exercises that assess the effects of the ban.

Since fertilizer is a tradable homogeneous good, its prices must follow the Law of One Price. Combined with the assumptions above, this implies that fertilizer prices p_i^f in region i are determined by the international fertilizer price (p_{RoW}^f) as follows:

$$p_i^f = \tau_{i,RoW}^f p_{RoW}^f \tag{20}$$

[41] This assumption is consistent with the preponderance of imported chemical fertilizer in Sri Lankan agriculture.

5.3 Consumers

Each region n has a fixed population of N_n agents who earn income and consume. Workers (numbering N_n^W) inelastically supply one unit of labor to local firms and are perfectly mobile across sectors, thus earning the local wage w_n independently of the sector in which they work. Farmers (numbering N_n^F), who own land, rent out this land for agricultural production, earning land rents, and consume.

5.3.1 Land Ownership

We allow for inequality in land ownership. Specifically, 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 μ_n and σ_{Ln}^2.[42] Conditional on its total land size L_n^h, we also assume that each farmer receives a random sample of plots.

These assumptions imply that the land rent income R_n^h of a farmer h in region n is proportional to its land size:

$$R_n^h = L_n^h \frac{R_n}{L_n}$$

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:

$$R_n^h \sim \text{log-N}(\mu_n + \ln(R_n) - \ln(L_n), \sigma_{Ln}^2) \tag{21}$$

5.3.2 Consumption across Sectors

All consumers have non-homothetic preferences with regard to their taste for agricultural versus manufacturing goods.[43] We represent this non-homotheticity using the price-independent generalized linear (PIGL) class of preferences (Eckert and Peters, 2023), in which an agent's indirect utility function is given by:

$$V(y, P^A, P^M) = \frac{1}{\eta}\Big(\frac{y}{(P^A)^{\phi}(P^M)^{1-\phi}}\Big)^{\eta} - \nu \ln\left(\frac{P^A}{P^M}\right) \tag{22}$$

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.[44] This formulation implies that agriculture's

[42] This assumption is consistent with the empirical distribution of landholdings observed in Sri Lanka's household survey data. See Appendix Section C.2 for details.

[43] This assumption is strongly supported by Sri Lanka's household survey data, in which agricultural goods' expenditure share decreases with household income, as shown in Figure A10.

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

share of expenditure (ξ^A) can be written as:

$$\xi^A(y, P^A, P^M) = \phi + \nu\Big(\frac{y}{(P^A)^{\phi}(P^M)^{1-\phi}}\Big)^{-\eta} \tag{23}$$

Because $\eta > 0$, equation (23) implies that agriculture's expenditure share decreases as the agent's income y increases, with an asymptote ϕ. As a consequence, the aggregate agricultural expenditure of a region n (denoted X_n^A) will depend on the distribution of income across agents within the region. As we prove in Appendix F.1, it can be shown that X_n^A can be written as:[45]

$$X_n^A = \phi E_n + \nu((P_n^A)^{\phi}(P_n^M)^{1-\phi})^{\eta}(N_n^W w_n^{1-\eta} + N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}}) \tag{24}$$

where P_n^A and P_n^M are price indices of the agricultural and manufacturing goods, respectively, in region n that are defined in Sections 5.3.3 and 5.3.4 below, $r_n \equiv \frac{R_n}{N_n^F}$ is the average land rent earned by farmers in region n, and E_n is the aggregate income in region n which can be written as:

$$E_n = \underbrace{w_n N_n^W}_{\text{wage income}} + \underbrace{R_n}_{\text{land rents}} + \underbrace{p_n^f F_n}_{\text{income from fertilizer sales}} \tag{25}$$

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

$$X_n^M = E_n - X_n^A \tag{26}$$

5.3.3 Within-Agriculture Consumption

We assume consumers allocate their agricultural expenditure across multiple crops according to Constant Elasticity of Substitution (CES) preferences with elasticity σ_A. Defining β_{nk}^A as the expenditure share of crop k *within total agricultural expenditure*, we can write β_{nk}^A as:

$$\beta_{nk}^A = \frac{b_k P_{nk}^{1-\sigma_A}}{(P_n^A)^{1-\sigma_A}}, \text{ for } k \in \{1, ..., K\} \tag{27}$$

where b_k is an exogenous preference shifter, P_{nk} is the price index of crop k in region n (defined below), and P_n^A is the price index for agricultural goods given by:

$$P_n^A = (\sum_k b_k P_{nk}^{1-\sigma_A})^{\frac{1}{1-\sigma_A}} \tag{28}$$

Furthermore, for each crop k, agents buy and consume geographically differentiated varieties

[45] In addition to the theoretical proof, we also provide a loose empirical test of equation (24) in Appendix C.4.

from all producing regions following CES preferences with elasticity σ_K. Defining $\beta^A_{ni,k}$ as the share of origin region i within the expenditures of region n on crop k, we can write $\beta^A_{ni,k}$ as:

$$\beta^A_{ni,k} = \frac{b_{i,k}(p_{ik}\tau^A_{ni,k})^{1-\sigma_K}}{P^{1-\sigma_K}_{nk}} \tag{29}$$

where $b_{i,k}$ is an exogenous preference shifter, p_{ik} is the local price of the variety of crop k produced in origin region i, and P_{nk} is the price index of crop k in region n given by:

$$P_{nk} = (\sum_i b_{i,k}(p_{ik}\tau^A_{ni,k})^{1-\sigma_K})^{\frac{1}{1-\sigma_K}} \tag{30}$$

5.3.4 Manufacturing Consumption

As in agriculture, each agent also decides on how to allocate her total manufacturing expenditure across geographically differentiated manufacturing varieties according to CES preferences with elasticity σ_M. Defining β^M_{ni} as the share of origin region i *within the total manufacturing expenditure* of destination region n, we can write β^M_{ni} as:

$$\beta^M_{ni} = \frac{(p^M_{ni})^{1-\sigma_M}}{(P^M_n)^{1-\sigma_M}}$$

where p^M_{ni} is the price of the manufacturing variety from origin region i in region n, and P^M_n is the price index for manufacturing goods in destination region n:

$$P^M_n = (\sum_i b^M_{ni}(p^M_{ni})^{1-\sigma_M})^{\frac{1}{1-\sigma_M}}$$

Substituting pricing equation (18) into the two expressions above, we can rewrite manufacturing expenditure shares and price indices as follows:

$$\beta^M_{ni} = \frac{(\tau^M_{ni}w_i)^{1-\sigma_M}}{(T^M_i P^M_n)^{1-\sigma_M}}, \quad P^M_n = \left(\sum_i (\tau^M_{ni}w_i/T^M_i)^{1-\sigma_M}\right)^{\frac{1}{1-\sigma_M}} \tag{31}$$

5.4 Equilibrium

We close the model by stating market clearing conditions in the markets for fertilizer, crops, and labor. First, the global fertilizer market clearing condition is:

$$\underbrace{F_{RoW}}_{\text{fertilizer supply}} = \sum_{i=1}^{R} \frac{\tau^f_{i,RoW}}{p^f_i} \underbrace{\sum_k \frac{\gamma^f_k}{S_k} p_{ik}Q_{ik}}_{\text{fertilizer expenditure in } i} \tag{32}$$

with fertilizer price p_i^f given by equation (20) and production Q_{ik} given by equation (14). Note that global fertilizer supply equals F_{RoW} because we assume Sri Lankan regions do not produce any fertilizer.

Second, the market clearing condition for each agricultural crop k produced by each region i is given by:

$$\underbrace{Q_{ik}}_{\text{local production of } k} = \sum_n \frac{\tau_{ni,k}^A}{p_{ni,k}} \underbrace{X_n^A \beta_{nk}^A \beta_{ni,k}^A}_{\text{exports of } k \text{ to region } n}, \forall (i,k) \tag{33}$$

where bilateral prices ($p_{ni,k}$) are given by equation (15), aggregate agricultural expenditures X_n^A are given by equation (24), and crop and origin expenditure shares β_{nk}^A and $\beta_{ni,k}^A$ are given by equations (27) and (29), respectively.

Third, the labor market clearing condition in each region i is:

$$\underbrace{N_i^W}_{\text{labor force}} = \frac{1}{w_i}\Big(\sum_k \underbrace{\frac{\gamma_k^n}{S_k} p_{ik} Q_{ik}}_{\text{wage bill of crop-}k\text{ farms}} + \underbrace{\sum_n X_n^M \beta_{ni}^M}_{\text{manufacturing wage bill}}\Big) \tag{34}$$

with manufacturing expenditure X_n^M in region n given by equation (26) and origin expenditure shares β_{ni}^M given by equation (31). Labor demand arises from two sources: agriculture and manufacturing. Agricultural demand for labor depends on each crop's total revenue but also on its labor-intensity (as represented by Cobb-Douglas coefficient γ_k^n). Since manufacturing uses labor as its sole input, its full revenue is paid to labor in the form of wages.

Given the three market clearing conditions above, we can define the model's equilibrium:

Definition 1 (Equilibrium) *Given the model's parameters (R, K, η, ν, ϕ, $\boldsymbol{\sigma}$, $\boldsymbol{\gamma}$) and exogenous variables (F_{RoW}, $\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{T}$, $\boldsymbol{\tau}$, $\boldsymbol{b}$) , an* ***equilibrium*** *is a set of endogenous variables ($\boldsymbol{p^A}$, $\boldsymbol{p^f}$, $\boldsymbol{p^-}$, $\boldsymbol{P}$, $\boldsymbol{w}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$, $\boldsymbol{Q}$) satisfying equations 11, 13, 14, 15, 16, 20, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34.* [46]

5.5 Welfare

In any given equilibrium, we can use the equilibrium variables to express the welfare of various groups of agents. For workers in region i, whose only source of income is wages, welfare V_i^W can be written as:

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

$$V_i^W = \frac{1}{\eta}\Big(\frac{w_i}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big) \tag{35}$$

Worker welfare 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.

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

$$V_i^F \equiv \mathbb{E}(V(R^h, P_i^A, P_i^M)) = \frac{1}{\eta}\Big(\frac{1}{(P_i^A)^{\phi}(P_i^M)^{1-\phi}}\Big)^{\eta} e^{\eta[\mu_i + \ln(\frac{R_i}{L_i})] + \eta^2 \frac{\sigma_L^2}{2}} - \nu \ln\Big(\frac{P_i^A}{P_i^M}\Big) \tag{36}$$

Using equations and and the Law of Iterating Expectations, we can then write the average welfare V_i^{avg} in region i as:

$$V_i^{avg} = \frac{N_i^W}{N_i} V_i^W + \frac{N_i^F}{N_i} V_i^F \tag{37}$$

6 Estimation

To use our quantitative model to study the impacts of the fertilizer import ban on the Sri Lankan economy, we start by calibrating the model to a pre-ban baseline, for which we use year 2019. In this section, we describe the steps involved in such calibration: first, estimating the model's parameters; second, performing a model "inversion" whereby we use data on observable variables and the equilibrium structure of the model to back out the implied unobservable variables, chiefly the regional productivity parameters in agriculture and manufacturing and the crop- and region-specific preference parameters.

6.1 Estimating Parameters

6.1.1 Preference Parameters

Preference-related parameters consist of a group related to the PIGL non-homotheticity (ϕ, η, ν) and another group that includes the elasticities of substitution ($\boldsymbol{\sigma}$) of the nested CES formulation. The parameter values we use are displayed in Table 1.

Within the first group, 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 (23). Therefore, we simply set ϕ to 0.0105, which is the lowest reported food expenditure share in our 2019 household survey data. Parameter ν is set to 0.12, following the literature (Eckert and Peters, 2023).

We estimate the Engel elasticity η using our 2019 household survey data. Rearranging equation (23), taking logs and adding a household h subscript, we obtain:

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

where h indexes households, n indexes regions, ξ^A_{nh} is the share of income that household h spends on food, and y_{nh} is household income.[47] Assuming price indices are do not vary across households within a given region, we can collect terms into a region fixed effect ω_n. Adding an error term ϵ and further controlling for household size $hhsize_{nh}$, the resulting estimating equation becomes:

$$\ln(\xi^A_{nh} - \phi) = -\eta \ln(y_{nh}) + \theta hhsize_{nh} + \omega_n + \epsilon_{nh} \tag{38}$$

To assuage concerns about the potential endogeneity of household income (e.g. due to seasonality in prices and employment opportunities), we estimate equation (38) using an instrumental variables (IV) approach. Specifically, 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.[48] The equation is estimated using Weighted Two-Stage Least Squares (2SLS), with standard errors clustered at the PSU (Primary Sampling Unit) level.[49] The resulting estimate for η is 0.656, which is highly statistically significant.[50]

Table 1: Preference Parameters

Parameter	Value	Standard Error	Source	Description
ϕ	0.0105	.	Household Survey	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	3.627	1.809	Estimated on trade data	EoS across origins (within crop)

To estimate elasticities of substitution across crops (σ_A) and across origins of a given crop (σ_K), we use our 2019 household survey data. To estimate σ_A, we start by rearranging equation (27), taking logs and adding a household h subscript, obtaining the equation:

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

[47] In this specification, regions are defined as Sri Lanka's 25 districts since divisional secretariats are not recorded in the household survey data.

[48] Both exogenous sources of income refer to the previous 12 months.

[49] Observations are weighted using sampling weights available in the household survey data.

[50] For the complete 2SLS estimation results of (η, σ_A, σ_K), including the first stages, see Appendix Section G.

where β^A_{nhk} is the share of crop k within total food expenditures of household h living in region n, and P_{nhk} is the price of crop k faced by household h. To operationalize estimation, we collect crop- and household-specific terms into fixed effects ω_k and ω_h and add an error term ϵ^A_{nhk}, obtaining the estimating equation:

$$\ln(\beta^A_{nhk}) = -(\sigma_A - 1)\ln(P_{nhk}) + \omega_k + \omega_h + \epsilon^A_{nhk} \tag{39}$$

Since there may be concerns about the potential endogeneity of crop prices (e.g. due to unobserved region-specific tastes), we follow the approach of Sotelo (2020) and adopt an IV approach to estimate equation (39). Specifically, price variable $\ln(P_{nhk})$ is instrumented 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. Therefore, our identification of the cross-crop elasticity of substitution σ_A comes from variation in crop prices due to regional shifters of agricultural supply that are driven by climactic and geological endowments, and thus arguably orthogonal to unobserved factors that may potentially affect consumer behavior. The equation is estimated using weighted 2SLS with standard errors clustered at the household level.[51] Our estimate, $\hat{\sigma}_A = 1.714$, is lower than that obtained by Sotelo (2020) (2.39) but higher than those in much of the literature.[52]

To estimate the Armington elasticity of substitution within a crop (σ_K), we leverage the differential incidence of the fertilizer ban across districts due to pre-existing differences in crop specialization. Given the demand system implied by equation (29), we can write the total agricultural exports from district i to the rest of the world ($X^A_{RoW,i}$) as:

$$X^A_{RoW,i} = \sum_k \left(\frac{p_{RoW,i,k}}{P_{Row,k}}\right)^{1-\sigma_K} X^A_{RoW,k}$$

where $X^A_{RoW,k}$ is the total expenditure of the rest of the world on crop k, $p_{RoW,i,k}$ is the price in RoW of crop k produced in region i, and $P_{Row,k}$ is the price index of crop k in the RoW. If we take time differentials of the equation with respect to a baseline period (defined as the first quarter of 2017) and rearrange, it can be shown that we obtain the following estimating equation:

$$\Delta^q \ln(X^A_{Row,i}) = -(\sigma_K - 1)\sum_k sh^0_{ik} \times \Delta^q \ln(p_{RoW,i,k}) + \omega_i + \omega_q + \epsilon^K_{iq} \tag{40}$$

where i indexes districts, q indexes quarters, Δ^q is a time difference operator defined as $\Delta^q Z \equiv Z_q - Z_{Q1,2017}$ for any variable Z, sh^0_{ik} is the initial share of crop k in the export mix of district i in the baseline year (2017), ω_i and ω_k are sets of district and crop fixed effects respectively, and ϵ^K_{iq} is

[51] The crops included in the estimation sample are: rice, maize, groundnut, onion, potato, coconut, and tea.

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

an error term.[53] To assuage concerns about potentially endogenous changes in crop prices (e.g. due to international demand shifts), we estimate equation (40) using an IV approach. Specifically, the main variable of interest $\sum_k sh_{ik}^0 \times \Delta^q \ln(p_{RoW,i,k})$ is instrumented with the following instrumental variable z_i^q:

$$z_i^q \equiv \sum_k sh_{ik}^0 \times \frac{\gamma_k^f}{S_k} \times \Delta^q \ln(p^f)$$

where $\frac{\gamma_k^f}{S_k}$ is the cost share of fertilizer for crop k, and p^f is the average fertilizer price in Sri Lanka. The instrument z_i^q leverages the interaction of pre-existing cross-regional differences in crop specialization (represented by sh_{ik}^0), differences in fertilizer intensity across crops (represented by fertilizer Cobb-Douglas coefficients γ_k^f), and fluctuations over time in fertilizer prices within Sri Lanka (represented by $\Delta^q \ln(p^f)$).[54] Given that fertilizer price changes in Sri Lanka were strongly influenced by the import ban, they are plausibly orthogonal to unobserved factors that may potentially affect crop prices, such as international demand shocks. We estimate equation (40) using 2SLS, obtaining an estimate $\hat{\sigma}_K = 3.627$ for the Armington elasticity of substitution within a crop.

We set the Armington elasticity of substitution within manufacturing to $\sigma_M = 2.528$, which is consistent with the lower range of estimates from Feenstra et al. (2018).[55]

6.1.2 Other Parameters

Our model also includes parameters unrelated to consumer preferences, most of which concern land inequality and agricultural production. Since our focus is on impacts on the agricultural market, it is important to find values for these parameters that are appropriate to our context. Table 2 displays the parameter values we use for agricultural production.

To estimate the fertilizer coefficients ($\{\gamma_k^f\}_k$) in the Cobb-Douglas agricultural production function (8), we leverage 2022 data on fertilizer intensity from the National Fertilizer Secretariat (see Appendix Table A1). For each crop k, we compute the fertilizer coefficient $\{\gamma_k^f\}_k$ as:

$$\gamma_k^f = \frac{M_k^{f,LKA}}{p_k^{LKA} Q_k^{LKA}}, \quad k \in \{1, ..., K\} \tag{41}$$

where $M_k^{f,LKA}$ is the annual USD value of fertilizer used for production of crop k in Sri Lanka, p_k^{LKA} is the average export price of crop k (in USD per kg), and Q_k^{LKA} is Sri Lanka's annual

[53] That is, initial share sh_{ik}^0 is given by: $sh_{ik}^0 = X_{RoW,i,k}^{A,2017} / X_{RoW,i}^{A,2017}$.

[54] Average fertilizer prices in Sri Lanka in each quarter are computed using data on fertilizer imports (see Section 3.2 for details). For details on how we the estimate fertilizer Cobb-Douglas coefficients for each crop, see Section 6.1.2.

[55] 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.

Table 2: Other Parameters

	Cobb-Douglas Coefficients		
Crop	Fertilizer (γ_k^f)	Labor (γ_k^n)	Land (γ_k^l)
Cinnamon	0.053	0.419	0.528
Cloves	0.043	0.424	0.533
Groundnuts	0.023	0.433	0.544
Maize	0.042	0.425	0.533
Onions	0.024	0.433	0.543
Potatoes	0.151	0.384	0.465
Rice	0.086	0.481	0.433

production of crop k (in kg).[56]

Given our estimates for the fertilizer coefficients (γ_k^f), and assuming production functions are constant returns to scale (i.e., $S_k = 1$ for all k), 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 the two coefficients using data from Sri Lanka's Input-Output (IO) matrix, specifically on labor's share of value added:

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

where $CEMP_s$ is total compensation of employees in sector s, VA_s is total value added (at factor cost price) in sector s, and $s(k)$ is the sector to which crop k belongs.[57] The land coefficient γ_k^l is then computed as: $\gamma_k^l = 1 - \gamma_k^f - \gamma_k^n$. Intuitively, we apportion each crop's cost shares (net of γ_k^f) to labor based on labor's share of value added according to the I-O matrix.

We also estimate the parameters (μ_i, σ_{Li}^2) governing land distribution in each district i. Given the properties of the log-normal distribution, we use 2016 household survey data to estimate σ_{Li} as the (weighted) standard deviation of the natural logarithm of landholdings across households.[58] Then, the properties of the log-normal imply that μ_i can be calibrated as:

$$\mu_i = \ln\left(\frac{L_i}{N_i^F}\right) - \frac{\sigma_{Li}^2}{2}$$

[56] Fertilizer usage $M_k^{f,LKA}$ includes the three main types of fertilizer used in Sri Lanka: urea, muriate of potash (MOP), and triple superphosphate (TSP). Average export prices p_k^{LKA} are computed using our customs data. Annual crop production is taken from the 2023 "Economic Statistics of Sri Lanka" report from the Department of Census and Statistics, available at: http://www.statistics.gov.lk/Publication/Economic-Statistics-2023.

[57] The IO table includes "minor export crops", "highland crops", "coconut and toddy", and potatoes as sectors. We classify cinnamon and cloves as "minor export crops", groundnuts, maize, and onions as "highland crops", coconut as "coconut and toddy", and rice as "paddy".

[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 hectares. The standard deviation is computed using only households who own positive landholdings. We use data from the 2016 household survey because landholding data was not available in the more recent 2019 survey.

where L_i is total arable land in district i (in hectares) and N_i^F is the number of farmers in region i.[59]

Finally, we initially set the international fertilizer trade cost parameter to $\tau^f_{i,RoW} = 1$ (for all i) to reflect the situation of relatively unimpeded imports of fertilizer prior to the bans of May 2021, thus allowing fertilizer prices to match international prices in equilibrium. We maintain this assumption during model estimation (see Section 6.2) since it refers to the Sri Lankan economy in the baseline year of 2019. Later, when running counterfactual exercises, we increase the value of this parameter to reflect the fertilizer import ban (see Section 7 for details).

6.2 Recovering Unobservables

Equipped with estimates for the model's parameters, we proceed to recover the remaining unobservable variables that are needed to fully characterize the 2019 baseline economy. To perform this task, we use the process of 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. More formally:

Definition 2 (Inversion) *Given the model's parameters, observable exogenous variables ($\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{\tau}$), and a specific value for a vector of observable endogenous variables ($\boldsymbol{w}$, $\boldsymbol{Q}$, $\boldsymbol{p^A}$, $\boldsymbol{p^f}$), an* ***inversion*** *is a set of unobservable exogenous variables ($\boldsymbol{T}$, $\boldsymbol{b}$, F_{RoW}) and unobservable endogenous variables ($\boldsymbol{p}^-$, $\boldsymbol{P}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$) such that the resulting set of endogenous variables is an equilibrium of the economy with the resulting set of exogenous variables.*

Appendix D provides details on the inversion procedure. We use data from 2019, which we choose as our baseline year because it is prior to the 2021 fertilizer ban (as well as other unusual events such as the Covid-19 pandemic) while also exploiting rich data from the 2019 household survey. The "regions" in the model correspond to Sri Lanka's 25 districts ($I = 25$). We include in the analysis the seven crops ($K = 7$) for which there is data on fertilizer intensity and district-level producer prices, namely: cinnamon, cloves, groundnut, maize, onions, potatoes, and rice.[60]

7 Counterfactuals

In this section, we use the estimated model to investigate the effects of Sri Lanka's fertilizer import ban. Starting from our estimated baseline economy, we compute a counterfactual equilibrium in

[59] Appendix Section C.2 provides more details on the estimation of (μ_i, σ^2_{Li}) for each district.

[60] In our current implementation of the model, we assume agricultural and manufacturing trade is cost-free, that is: $\tau^M_{ni} = 1$ and $\tau^A_{ni,k} = 1$ for all (n,i,k). The next version of the paper will feature positive trade costs in these two sectors.

which the key fertilizer trade cost parameter ($\{\tau^f_{i,RoW}\}^I_{i=1}$) is increased (so as to represent the fertilizer import ban) while all other exogenous variables (F_{RoW}, $\boldsymbol{N}$, $\boldsymbol{L}$, $\boldsymbol{T}$, $\boldsymbol{\tau^A}$, $\boldsymbol{\tau^M}$, $\boldsymbol{b}$) 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 fertilizer import ban shock in general equilibrium. Since fertilizer can only be sourced from abroad, the effects of the ban can also be equivalently interpreted as the value of access to fertilizer for the economy.

7.1 Setting Counterfactual Parameters

There are three practical considerations when implementing the counterfactual. First, since the equilibrium system is homogeneous of degree zero in the set of prices, we must choose one price to be normalized. We set the average wage in the rest of the world (w_{RoW}) to its actual value in the pre-ban period.[61] For convenience of exposition, all results are displayed in US dollars (USD).

Second, we must choose a specific value for the counterfactual trade cost parameter ($\{\tau^f_{i,RoW}\}^I_{i=1}$). We set this parameter to $\tau^f_{i,RoW} = 1.875$, for $i \in \{1, ..., I\}$, because it leads to a counterfactual equilibrium in which the fertilizer price in Sri Lanka ($\{p^f_i\}^I_{i=1}$) is \$0.58/kg, the average price observed in the post-ban year of 2022.[62] In other words, the change in fertilizer trade costs is calibrated to generate the same increase in Sri Lankan fertilizer prices as was observed in the data from the pre- to the post-ban period.

Third, we need a welfare metric that is able 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 (22), it is difficult to interpret the meaning of such metrics.[63] 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: workers, three types of farmer (median, average, and representative), and the representative agent.[64] Appendix E provides more detailed

[61] We use the world's income per capita in 2019, which was \$11,330.5, equivalent to approximately 4.08 million LKR (using the 2022 exchange rate).

[62] The average fertilizer price in Sri Lanka in 2022 was actually \$0.6458/kg. However, part of the price increase between 2019 and 2022 may be due to dollar inflation, so we deflate the 2022 price by the cumulative 2019-2022 inflation (approximately 11.1%), leading to an adjusted price of \$0.5813/kg. Inflation data is from the Federal Reserve Economic Data (FRED) database maintained by the Federal Reserve Bank of St. Louis, which is available at: https://fred.stlouisfed.org/series/FPCPITOTLZGUSA

[63] 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% in magnitude, for example.

[64] 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 i is the farmer whose utility is equal to V^F_i, the district's average farmer utility. Similarly, the representative agent is defined as an hypothetical agent whose income

definitions of equivalent variation and agent types.

7.2 Results

In this section, we present estimated causal effects of the fertilizer import ban on various outcomes of interest, starting in Table 3 with the set of outcomes that are measured at the national level. Following our estimation strategy, the primary effect of the fertilizer ban consists of an increase in the price of fertilizer by \$0.27/kg (or 87%) between the baseline and the counterfactual equilibrium. As a result, the quantity of imported fertilizer decreases by half (or about 322 million kilograms), leading to a drop of \$11.5MM in the value of fertilizer imports.

Table 3: Ban Effects (Country-Level)

Variable	% Effect	Effect	Description
M^f_{LKA}	-5.7%	-\$11.5 MM	Fertilizer imports (\$)
f_{LKA}	-49.7%	-321.8 MM kg	Fertilizer imports (kg)
p^f_{LKA}	+87.2%	+\$0.270	Fertilizer price (\$/kg)
EXP^A_{LKA}	-5.5%	-\$137.7 MM	Agricultural exports
$NEXP^A_{LKA}$	-6.7 %	-\$134.3 MM	Net agricultural exports

Notes: The table shows absolute and relative (%) changes in several variables between the baseline equilibrium and the counterfactual equilibrium. Fertilizer imports (M^f_{LKA}), measured in millions of US dollars (\$ MM) and defined as the total expenditure of the Sri Lankan agricultural sector on fertilizer imported from the Rest of the World (RoW), are given by: $\sum_{i=1}^{I}\sum_k p_{ik}Q_{ik}\gamma^f_k/S_k$. Imported fertilizer quantity (f_{LKA}), measured in kilograms (kg) and defined as the total quantity of fertilizer imported by the Sri Lankan agricultural sector from RoW, is given by: $\sum_{i=1}^{I}\sum_k p_{ik}Q_{ik}\gamma^f_k/(S_k p^f_i)$. Average fertilizer price in Sri Lanka (p^f_{LKA}), measured in US dollars per kilogram (\$/kg) and defined as the ratio of total fertilizer expenditures to total fertilizer usage by the Sri Lankan agricultural sector, is given by M^f_{LKA}/f_{LKA}. Gross agricultural exports (EXP^A_{LKA}), measured in millions of US dollars (\$ MM) and defined as the sum across all Sri Lankan districts of district-level agricultural sales to RoW, are given by: $\sum_{i=1}^{I}\sum_k X^A_{RoW}\beta^A_{RoW,k}\beta^A_{RoW,i,k}$. Net agricultural exports ($NEXP^A_{LKA}$) are defined as gross agricultural exports minus gross agricultural imports, with the latter defined as the sum across all Sri Lankan districts of district-level agricultural purchases from RoW and computed as: $\sum_{n=1}^{I}\sum_k X^A_n\beta^A_{nk}\beta^A_{n,RoW,k}$.

The resulting lack of fertilizer leads to declines in Sri Lankan crop yields, which ultimately result in falling agricultural exports to the tune of \$138MM. The magnitude of the export losses relative to the reduction in fertilizer imports indicates that the ban's implicit motivation of saving foreign exchange may have been jeopardized by the substantial decline in crop exports. To better understand the ban's effects on agriculture, Table 4 shows crop-by-crop results.

provides him a level of utility equal to V_i^{avg}, the district's average utility.

Table 4: Ban Effects (Crop-Level)

Crop (k)	**Fertilizer intensity** (γ_k^f)	**Production** in MT	**in %**	**Gross Exports in \$ MM**	**in %**	**Net Exports in \$ MM**
Cinnamon	0.053	-1,083	-4.7%	-7.7	-3.5%	-7.7
Clove	0.043	-263	-3.9%	-1.1	-3.0%	-1.1
Groundnuts	0.023	-491	-1.3%	-0.4	-1.2%	-0.2
Maize	0.042	-10,268	-3.3%	-2.2	-2.5%	-1.0
Onions	0.024	-695	-2.0%	-0.5	-1.6%	-0.1
Potatoes	0.151	-4,988	-14.3%	-2.8	-10.7%	-2.2
Rice	0.086	-402,865	-7.9%	-123.1	-5.9%	-122.0

Notes: For each crop, the table shows its fertilizer intensity and the change in Sri Lanka's crop production and crop exports between the baseline equilibrium and the counterfactual equilibrium. The fertilizer intensity of each crop k is defined as the fertilizer coefficient in its Cobb-Douglas production function (γ_k^f). Production of crop k, measured in metric tons (MT), is computed by adding district-level production across all Sri Lankan districts ($\sum_{i=1}^{I} Q_{ik}$). Gross exports of crop k, measured in millions of US dollars (\$ MM), are defined as the sum across all Sri Lankan districts of district-level sales of crop k to RoW, which can therefore be computed as: $\sum_{i=1}^{I} X_{RoW}^{A} \beta_{RoW,k}^{A} \beta_{RoW,i,k}^{A}$. Net exports of crop k are defined as gross exports minus gross imports of crop k, where gross imports are defined as the sum across all Sri Lankan districts of district-level purchases of crop k from RoW, which are computed as: $\sum_{n=1}^{I} X_{n}^{A} \beta_{nk}^{A} \beta_{n,RoW,k}^{A}$.

As a result of the ban, physical yields fell by 1-14%, 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, potato, saw yields falling by 14.3% and exports by 10.7%. On the other hand, less fertilizer-intensive crops like groundnuts saw decreases of only 1.3% in yields and 1.2% in exports. Further illustration is provided by Figure 8, which plots each crop's fertilizer coefficient γ_n^f (which captures the crop's fertilizer intensity) against the ban's effect on the value of the crops. It clearly shows a strong negative correlation of -0.999 between the two variables, with the two most fertilizer-intensive crops, rice and potatoes, suffering the strongest percentage exports losses by far. The decrease in rice yields predicted by our counterfactual exercise is 7.9%, as shown in Table 4. This explains one quarter of the 2022 drop in rice yields observed in the data (relative to the 2013-2021 average).

Appendix Table A6 displays the effects of the ban across Sri Lankan districts on the welfare of workers, farmer, and the representative agent. Welfare effects are negative for all agents in all districts, but the size of the losses varies considerably by occupation and geography. Farmer are harmed disproportionately because they are attached to the agricultural sector, in contrast to workers, who have the option of switching employment to the manufacturing sector even if they were initially working in agriculture. Due to this composition effect, the welfare of the representative agent decreases more in regions that have a relatively high fraction of farmers in their population,

Figure 8: Ban Effects on Exports, by crop

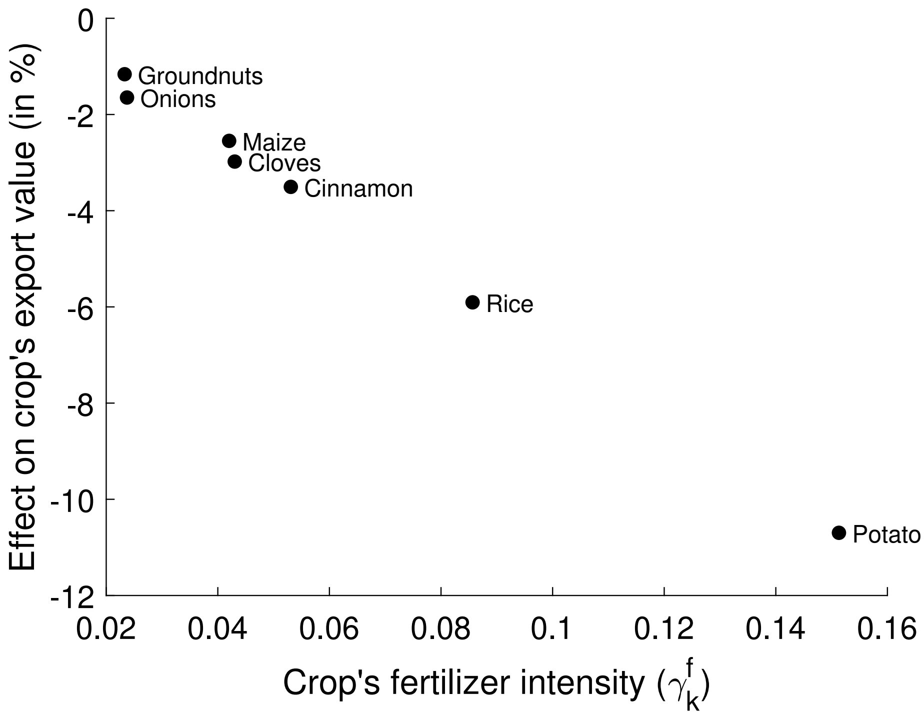

Notes: The figure plots the fertilizer intensity of each crop against the % change between the baseline equilibrium and the counterfactual equilibrium in the value of the crop's exports from Sri Lanka to the Rest of the World (RoW). The fertilizer intensity of crop k is defined as the fertilizer coefficient in its Cobb-Douglas production function (γ_k^f). The export value of crop k is defined as the sum across all Sri Lankan districts of district-level sales of crop k to RoW, which can therefore be computed as: $\sum_{i=1}^{I} X_{RoW}^{A} \beta_{RoW,k}^{A} \beta_{RoW,i,k}^{A}$. The correlation coefficient between fertilizer intensity and % change in export value is -0.9987.

Figure 9: Ban Effects on Welfare, by district (representative agent)

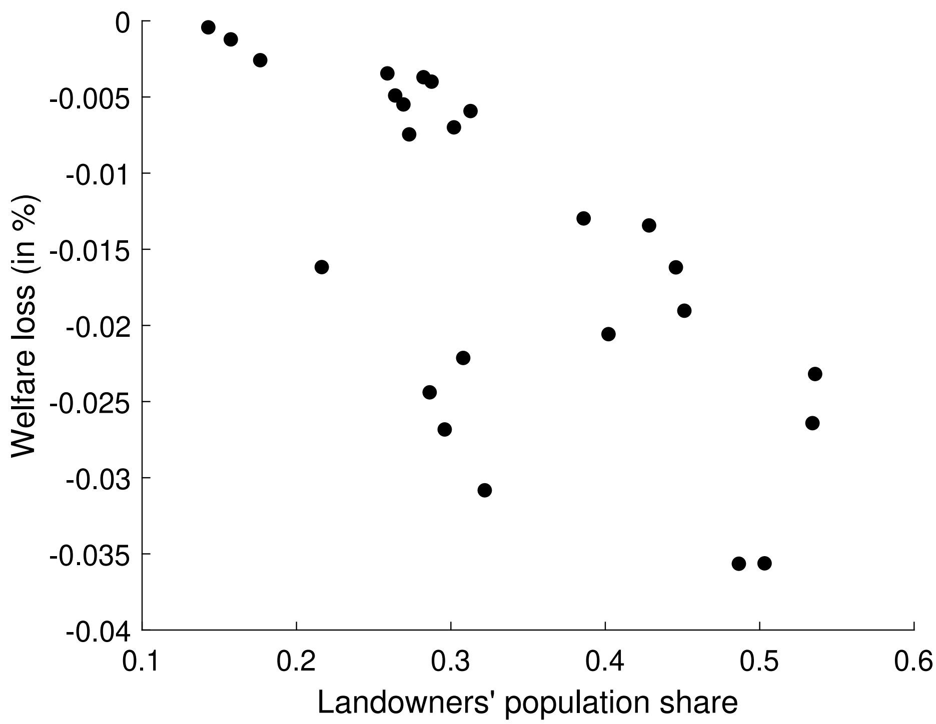

Notes: The figure plots the farmer population share of each district against the % change in the welfare of the district's representative agent between the baseline equilibrium and the counterfactual equilibrium. Farmers' population share in district i is defined as N_i^F/N_i, where N_i^F is the number of farmers and N_i is the total population of the district. Welfare % change is defined as an Equivalent Variation (EV): starting from the representative agent's baseline income, the EV is the % change in income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels. The correlation coefficient between farmers' population share and % welfare change is -0.700.

as shown in Figure 9. For example, in the districts where farmers make up less than 20% of the population, the harm to the representative agent is equivalent to an income loss of less than 0.5%. In contrast, wherever farmers represent more than half of the population, the corresponding damages are equivalent to a 2.0%-3.5% income reduction.

Figure 10 shows how welfare losses vary across space. For farmers (panel 10a), this spatial heterogeneity can be explained by geographic differences in the relative importance of specific crops within agriculture. As shown in Figure 11, the damage on farmers (and on the representative agent more generally) was worse in districts whose agricultural sector specializes in relatively fertilizer-intensive crops, and who are thereby more negatively affected by the fertilizer ban.[65] For example, in districts with a 6% fertilizer cost share, the harm to representative farmers was

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

Figure 10: Spatial Heterogeneity of Ban Effects on Welfare

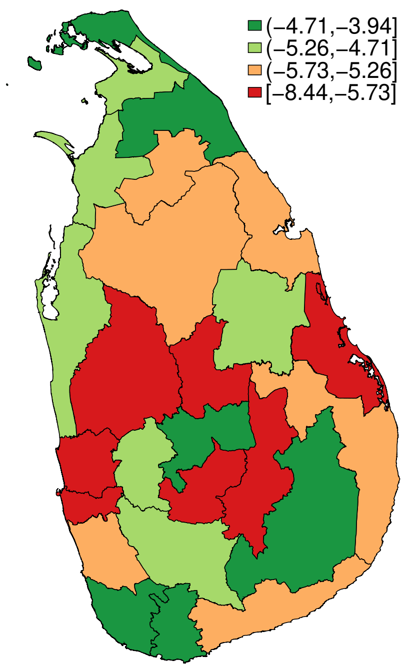

(a) Representative Farmer

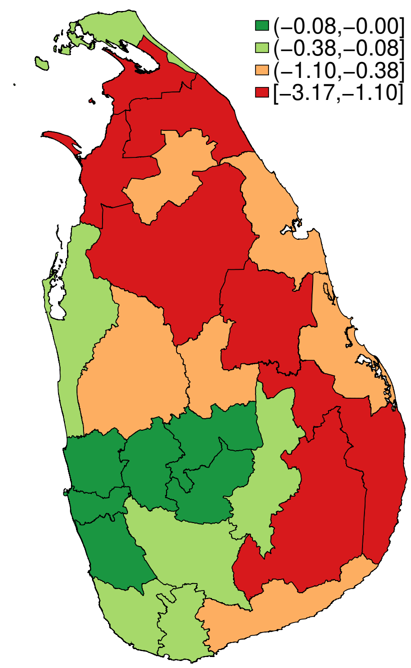

(b) Worker

Notes: The figures show the % change in the welfare of each district's representative farmer (panel 10a) and worker (panel 10b) between the baseline equilibrium and the counterfactual equilibrium. Welfare % change is defined as an Equivalent Variation (EV): starting from an agent's baseline income, the EV is the % change in income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels.

Figure 11: Ban Effects on Welfare, by district

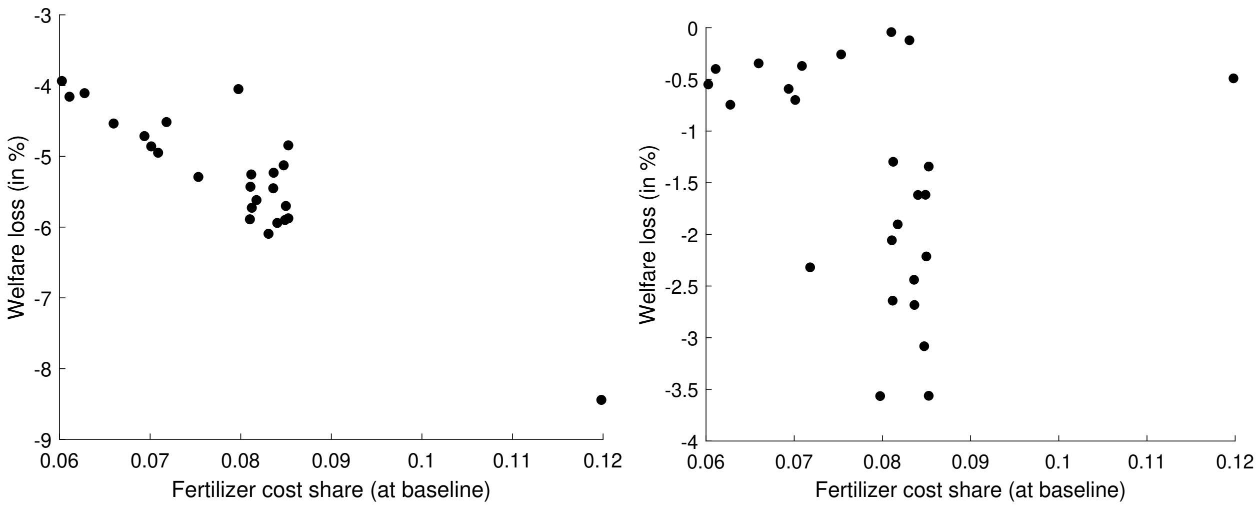

(a) Representative Farmer

(b) Representative Agent

Notes: The figure plots each district's fertilizer cost share in the baseline equilibrium against the $\%$ change in the welfare of the district's representative farmer (Panel 11a) or representative agent (Panel 11b) between the baseline equilibrium and the counterfactual equilibrium. Fertilizer cost share in district i is defined as the ratio of the district's total fertilizer expenditure ($\sum_k p_{ik} Q_{ik} \gamma_k^f / S_k$) to its total agricultural revenues ($\sum_k p_{ik} Q_{ik}$). Welfare $\%$ changes are defined as Equivalent Variations (EV): starting from the baseline income of the representative farmer or representative agent, the EV is the $\%$ change in income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels. The correlation coefficients between fertilizer cost share and $\%$ welfare change are -0.8963 (Panel 11a) and -0.2560 (Panel 11b).

equivalent to a 4% income loss, whereas the corresponding damage to their counterparts in districts with a 8%-9% fertilizer cost share was as high as 6%. This correlation can also be informally illustrated by comparing Figures 5 and 10a, noticing that some of the worst farmer losses were in high-rice yield areas, which tend to be specialized in rice, a fertilizer-intensive crop.

On the other hand, the spatial heterogeneity in the welfare losses of workers (panel 10b) can be explained by a slightly different type of geographic variation, namely the size of the agricultural sector as a whole.[66] Figure 12 shows that workers lost more welfare in districts with a high agricultural share of employment. Therefore, although workers suffered less than farmers everywhere, the workers living in heavily agricultural regions still faced losses equivalent to as much as a 3% income reduction because their regions did not have a vibrant manufacturing sector that could serve as a buffer and easily absorb the excess labor from the adversely-affected agricultural sector.

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 value of fertilizer in agricultural production and trade 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 agricultural production, fertilizer imports, and exports of fertilizer-intensive agricultural products.

To better understand the mechanisms and welfare implications of the fertilizer ban's effects, we propose and estimate a quantitative spatial model of trade and agriculture that reflects key features of the Sri Lankan agricultural economy. The model implies that the welfare effects from the ban differed across regions in Sri Lanka, depending on the region's propensity to grow fertilizer-intensive crops. We find average losses equivalent to a 1.5% income decline due to fertilizer scarcity, with stronger incidence on farmers (whose income is tied to agriculture) relative to workers. The ban lowered crop yields by 1.3%-14.3%, leading to losses of $137.7 million in foregone agricultural exports, which more than offsets the foreign exchange savings from reduced fertilizer imports. Our findings quantify the value of fertilizer for agriculture in general equilibrium, an estimate that helps inform fertilizer-related policy (such as fertilizer subsidies) and the wider debate on environmental regulation.

[66] The fact that spatial heterogeneity in farmer welfare effects and in worker welfare effects are explained by different factors is also suggested by the substantial differences between Figures 10a and 10b in the spatial pattern of welfare losses.

Figure 12: Ban Effects on Welfare, by district (workers)

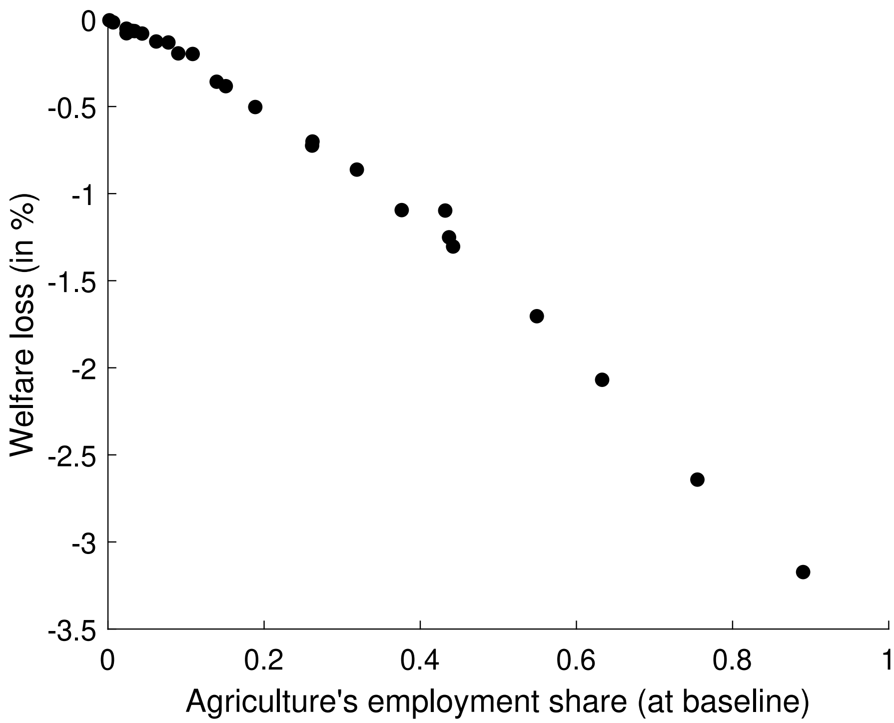

Notes: The figure plots each district's agricultural share of employment in the baseline equilibrium against the % change in the district's worker welfare between the baseline equilibrium and the counterfactual equilibrium. Agricultural employment share in district i is defined as $N_i^{A,W}/N_i^W$, where N_i^{AW} is agricultural employment and N_i^W is the total number of workers in the district. Welfare % change is defined as an Equivalent Variation (EV): starting from the worker's baseline income, the EV is the % change in income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels. The correlation coefficient between agricultural employment share and % welfare change is -0.992.

References

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

Amarasinghe, D. and J. S. Bandara (2005). Input output tables for sri lanka – 2000.

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

Atkin, D., J. Blaum, P. Fajgelbaum, and A. Ospital (2022). Protectionism unchained: Determinants and consequences of discretionary trade policy in argentina. Technical report, Mimeo.

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

Behrman, J. R. and A. Deolalikar (1989). Is variety the spice of life? implications for calorie intake. *The review of economics and statistics*, 666–672.

Bergquist, L. F., B. Faber, T. Fally, M. Hoelzlein, E. Miguel, and A. Rodríguez-Clare (2022). Scaling agricultural policy interventions. Technical report, National Bureau of Economic Research.

Bernini, F. and E. García-Lembergman (2022). Multi-destination exporters, market power and the elasticity of markups across destinations. Technical report, Mimeo.

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.

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

Callaway, B. (2022). Difference-in-differences for policy evaluation. *arXiv preprint arXiv:2203.15646*.

Carleton, T., L. Crews, and I. Nath (2023). Agriculture, trade, and the spatial efficiency of global water use. *Working Paper*.

Carter, M., R. Laajaj, and D. Yang (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.

Costinot, A., D. Donaldson, and C. Smith (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 X. d'Haultfoeuille (2022). Difference-in-differences estimators of intertemporal treatment effects. Technical report, National Bureau of Economic Research.

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

Dong, J., X. Xiao, M. A. Menarguez, G. Zhang, Y. Qin, D. Thau, C. Biradar, and B. Moore III (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.

Duflo, E., M. Kremer, and J. Robinson (2008). How high are rates of return to fertilizer? evidence from field experiments in kenya. *American economic review 98*(2), 482–488.

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

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

Ederington, J. and M. Ruta (2016). *Nontariff Measures and the World Trading System*, Chapter 5, pp. 211–277. Elsevier.

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.

Farrokhi, F. and H. S. Pellegrina (2022). Trade, technology, and agricultural productivity. *Working Paper*.

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

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

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

Foster, A. D. and M. R. Rosenzweig (2008). Economic development and the decline of agricultural employment. In T. P. Schultz and J. A. Strauss (Eds.), *Handbook of Development Economics*, Volume 4, Chapter 47, pp. 3051–83. Elsevier.

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

Hebebrand, C. and J. W. Glauber (2023). The russia-ukraine war after a year: Impacts on fertilizer production, prices, and trade flows. *IFPRI book chapters*, 1–8.

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

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

Nath, I. (2023). Climate change, the food problem, and the challenge of adaptation through sectoral reallocation. *Working Paper*.

Ozdogan, M., S. Wang, D. Ghose, E. Fraga, A. Fernandes, and G. Varela (2023). Paddy rice area and yield mapping in sri lanka using remote sensing. *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 D. Ord (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.

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

Wijesinghe, A., C. Yogarajah, and A. Rathnayake (2023). Import controls in the pandemic: Is sri lanka back to protectionism? Technical report, Institute of Policy Studies of Sri Lanka.

Zereyesus, Y. A., L. Cardell, C. Valdes, K. Ajewole, W. Zeng, J. Beckman, M. Ivanic, R. N. Hashad, J. Jelliffe, and J. Kee (2022). International food security assessment, 2022–32. Technical report.

Appendices

A Data Appendix

A.1 Panjiva trade data cleaning

The names of firms in Sri Lanka engaged in the export or import transactions are noisy. Names appear with multiple spellings or with spelling errors and some entries under firm names are not proper names but are instead a number or an address. 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 the actual firm names a list of stop words or geographic stop words and any symbol (that is not a letter nor a number). The list of stop words was selected based on text analysis indicating the most common words appearing in firm company names, such as ‘corporation’, ‘international’, ‘limited’, ‘llc’. We also remove from the firm names prefixes and suffixes for example ‘co’, ‘sa’ or ‘pvc’.

3. Based on the list of pre-processed firm names obtained after 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. The 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 (i.e., a measure of the similarity between two strings). Cutoffs to determine what is a “small” distance between pairs 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 as identified in Step 4 are then sorted alphabetically so as to identify neighbor names (names that appear in consecutive rows after sorting) and correct any potentially similarity that may have been missed in Step 4.

6. The algorithms in Step 4 above generate some “big groups” of firm names considered to be 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 so two separate groups of similar firm names result .

7. After similar firm names are identified from all steps above and are assigned a temporary unique numeric identifier, firm names are sorted alphabetically but also based on the firm

address. If two consecutive names have similar addresses and similar names and did not have the same unique numeric identifier they will be joined in a a new final unique numeric identifier.

The steps above are applied to clean names of firms in Sri Lanka. 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 the total trade value is negligible respect we dropped all observations whose names are of courier firms, namely 'dhl express parcel delivery', 'd h l', 'fits aviation', 'cargo logistics worldwide aramex courier freight "non commercial" "non commecial" xpress "by air freight" Skynet shipping setrans "sky international". Finally, we also drop from the sample observations for which the firm name corresponds to individuals making personal imports under a permit, such firm names have the structure 'Person Name+Institution+PERMIT' e.g., 'KADAWARAGE C V WING COMMANDER (DENTAL DOCTOR) SRI LANKA AIR FORCE MIN.OF DEFENCE PERMIT'. Such observations account for a very small total trade value.

A.2 Fertilizer Intensity

Table A1 displays crop-level usage of different kinds of fertilizer in Sri Lanka. The last column shows each crop's resulting fertilizer intensity (measured in kilograms per hectare). As explained in Section 4.3.1 above, this fertilizer intensity metric is used in our main analysis to build a firm-level metric of fertilizer usage (see equation 2).

Table A1: Fertilizer Usage and Intensity in Sri Lanka, by crop

Crop	**Urea**	**TSP**	**MOP**	**Other Fertilizers**	**Total**	**Fertilizer Intensity (MT/ha)**
Chili pepper	3800	3800	3800	0	11400	.300
Cinnamon	15114		7558	0	22672	.662
Cloves	1359		2039	226	3624	.454
Cocoa	178		222	44	444	.243
Coconut	20147	2865	40295	15000	78307	.173
Cowpea	1008	1550	1203	0	3761	.243
Fruits	48359	35353	80067	0	163779	1.279
Green gram	1495	2300	1725	0	5520	.240
Horse gram	3	6	2	0	11	.073
Kurakkan	2345	605	908	0	3858	.351
Maize	32300	14000	7000	0	53300	.381
Nutmeg	484		727	80	1291	.454
Onion	2926	1500	1126	0	5552	.370
Paddy	247000	61000	74000	0	382000	.294
Peanuts	1203	1850	1388	0	4441	.240
Pepper	14562		10922	3640	29124	.710
Potatoes	2310	1890	1750	0	5950	.850
Rubber	6767		7151	4391	18309	.146
Sesame	1595	1740	870	0	4205	.290
Soybeans	.	.	.	.	.	.067
Tea	96050		36864	59206	192120	.959
Undu	1755	2700	2025	0	6480	.240
Vegetables	22744	21297	14920	0	58961	.640

Source: National Fertilizer Secretariat (except for Soybeans, whose fertilizer intensity comes from the "Asia" entries in the databases of the International Fertilizer Development Center and the International Fertilizer Association). All fertilizers are measured in metric tons (MT).

The data on Table A1 comes from Sri Lanka's National Fertilizer Secretariat. However, the data on soybeans was deemed unreliable by agriculture specialists, so we complement our data using the databases of 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. Figure A1 shows that this alternative fertilizer intensity metric has a close linear relationship with our original metric from the National Fertilizer Secretariat (with the exception of soybeans, which are a clear outlier in the graph).

Figure A1: Fertilizer Intensity: Sri Lanka vs Asia

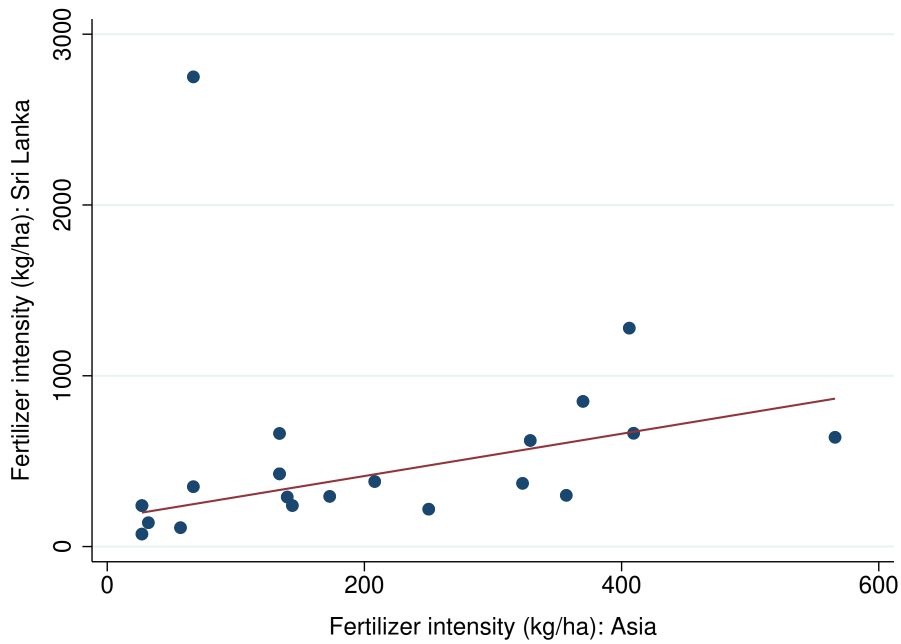

Notes: The figure shows the fertilizer intensity of 22 crops (see Section 3 for the full list), measured in kg/ha, from two alternative data sources: Sri Lanka's National Fertilizer Secretariat (on the vertical axis) and the entries corresponding to "Asia" in the databases of the International Fertilizer Development Center and the International Fertilizer Association (on the horizontal axis). 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.

A.3 Data on Rest of World (RoW)

Because the Rest of the World (RoW) features as one of the regions in the 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 RoW are assumed to be equal to the world's GDP per capita (in current US$) in 2019, which was $11,330.5 according to the World Bank's World Development Indicators.[67]

We set the number of workers in 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 RoW in two steps. First, we multiply the number of workers in RoW by the agricultural share of global employment (26%, also from ILO data), thus obtaining the number of agricultural workers in RoW.[68] 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.[69] The resulting estimate for the number of farmers in RoW is approximately 260 million.

Our data on agricultural production, cultivated area, and producer prices in RoW comes from the Food and Agriculture Organization of the United Nations (FAO).[70] We download data for each country's production quantities (in tonnes), harvested area (in hectares), and producer price (in USD per tonne) of seven crops in year 2019, then sum across countries to obtain production and cultivated area in RoW.[71] Each crop's producer price in RoW is set to the median price across all available countries.

Finally, it is sometimes necessary in our analysis to convert currency denominations between US dollars and Sri Lankan Rupees (LKR). We use the exchange rate $E_{2019} = 176$ LKR/USD for year 2019 and $E_{2022} = 360$ LKR/USD for year 2022, both of which were obtained from Google Finance.[72]

[67] We use data from 2019 because this is the year we use as the baseline period in model estimation (see Section 6). The data portal can be accessed at: https://databank.worldbank.org/indicator/NY.GDP.PCAP.CD/1ff4a498/Popular-Indicators

[68] 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

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

[70] 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

[71] The crops we use are listed in FAO's 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".

[72] The exact dates from which the exchange rates were extracted are June 28 2019 and Jun 24 2022.

A.4 Weather Variables

Since a wide array of weather variables have been shown to be important drivers of crop yields across a broad range of crops grown in Sri Lanka, it is important to control for weather in our analysis of the effects of the fertilizer import ban on rice yields in Section C.3. In this section, we briefly describe our sources of weather data.

For rainfall data, we use the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), a quasi-global gridded rainfall time series data set based on a combination of high-resolution satellite imagery and in-situ station data which covers latitudes between 50°S and 50°N, all longitudes, and all years from 1981 to the near-present.[73] For each agricultural growing season (Maha and Yala) and year between 1999 and 2022, we obtain three variables from CHIRPS: total rainfall (in millimeters 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).[74] 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.[75] For each growing season and year between 1999 and 2022, we obtain four variables from ERA5-Land: average incoming shortwave solar radiation (in watts per square meter); total vapor pressure deficit (in hectopascals); and average minimum and maximum temperatures (in degrees Celsius).[76] As with rainfall, we aggregate grid-cell level data up to the district level.

[73] For more information and to access CHIRPS, see the website: https://www.chc.ucsb.edu/data/chirps.

[74] The Maha season is defined as starting on September 1 and ending on March 31st every year. The Yala season is defined as starting on May 1st and ending on August 31st of every year.

[75] For more information and to access ERA5-Land, see the website: https://www.ecmwf.int/en/era5-land.

[76] Vapor pressure deficit (VPD) is the difference between saturated vapor pressure (temperature dependent) and the actual vapor pressure. Higher VPD values indicate higher dryness of the air 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.

B Characterization of the Import Bans

Table A2: Import share and total imports of HS 8-digit products subject to import bans

Year	Import share	Imports (in Million USD)
2017	17.17%	2,720
2018	19.01%	3,090
2019	16.47%	2,250
2020	13.23%	1,510
2021	10.59%	1,310

Source: authors' calculations based on data compiled from Sri Lanka's Extraordinary Gazettes and from the Panjiva trade data platform.

Table A3: List of HS 8-digit Fertilizer Products Subject to Import Bans

HS8	Description
31024000	Mixtures of ammonium nitrate with calcium carbonate or other inorganic non-fertilising 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 fertilisers containing the three fertilising 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 fertilisers containing the two fertilising elements phosphorus and potassium
31059000	Other

Source: Sri Lanka's Extraordinary Gazettes on Imports and Exports.

Figure A2: Timing of Fertilizer Import Bans

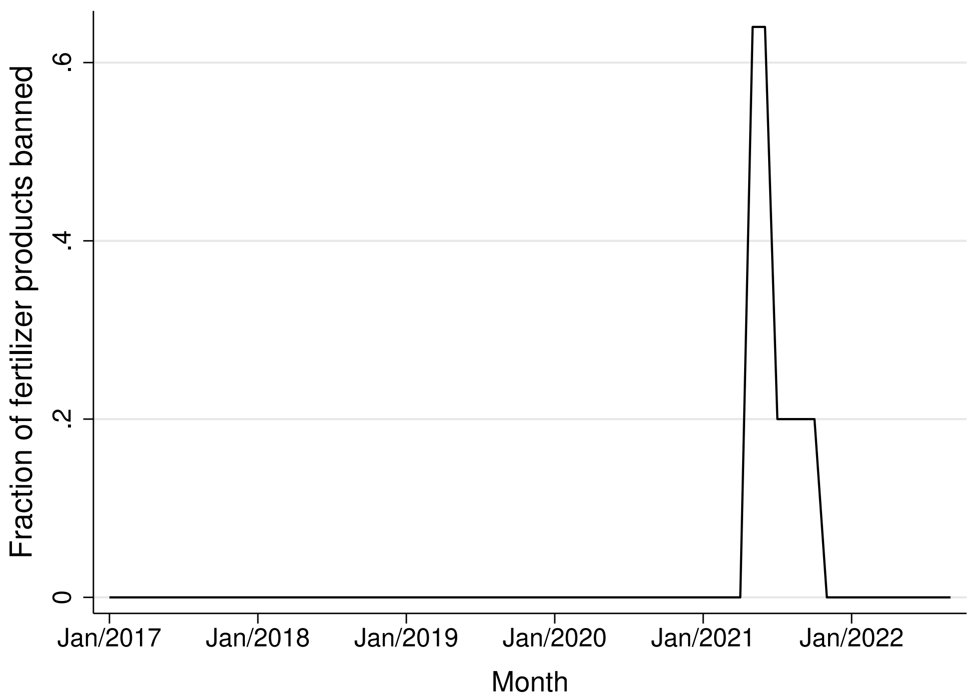

Notes: For each month between January 2017 and September 2022, the figure shows the fraction of HS8 fertilizer products that were subject to an import ban. The denominator of the fraction is 25, which is the number of products that were imported by Sri Lanka at any point during the 2017-2022. A fertilizer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

Figure A3: Dynamic Ban Effects on Imports: Quantity vs. Price

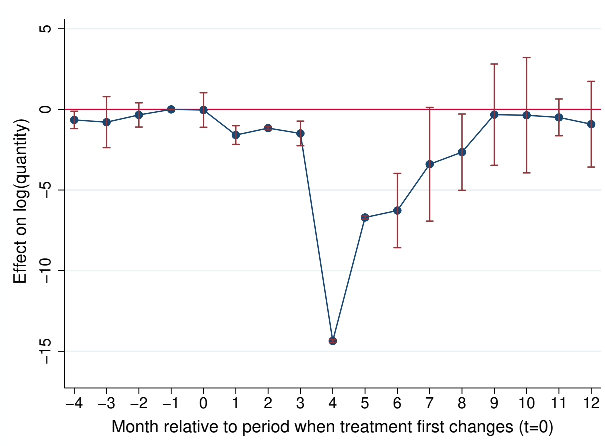

(a) Quantity Effect

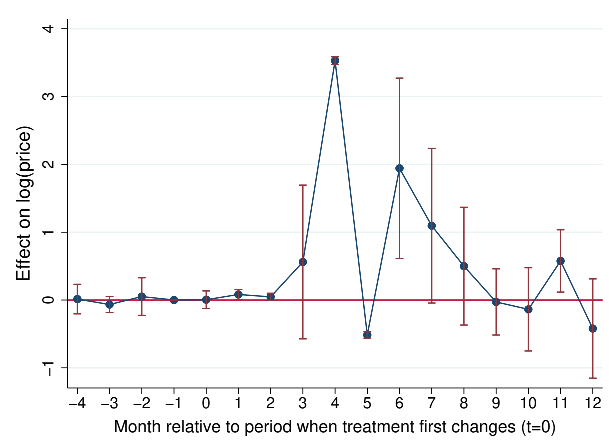

(b) Price Effect

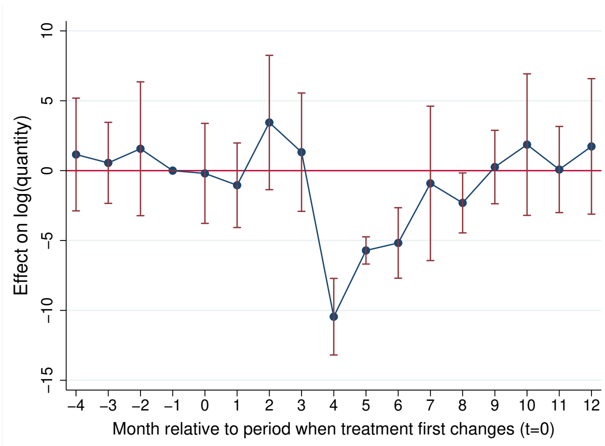

(c) Quantity Effect, fertilizer only

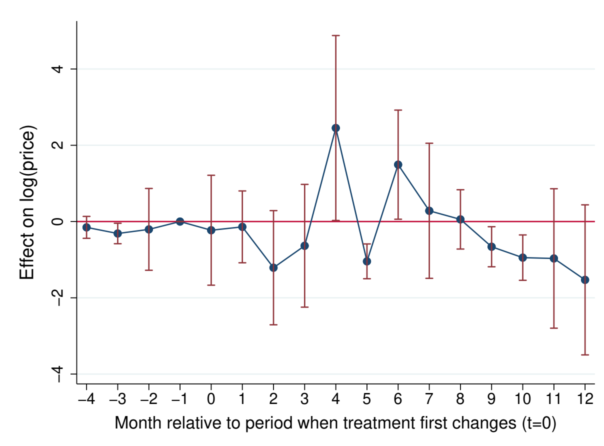

(d) Price Effect, fertilizer only

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of import bans on the quantity and average price of imports, relative to the last month before the ban takes effect, using an HS8 product-level panel and methods described in section 4.1.1 to isolate the ban effects. The sample excludes non-fertilizer HS8 products whose imports were ever banned. Additionally, in panels A3c and A3d, the sample also excludes all non-fertilizer products, even the ones who were never banned. The treatment variable is a dummy ban_{ct} indicating whether product c was banned in 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 quantity of imports (panels A3a and A3c) or on the log average price of imports (panels A3b and A3d), with negative values correspond to decreases in quantities/prices. The p-values of the joint significance tests for the placebo estimators were 0.02 (panel A3a), 0.50 (panel A3b), 0.92 (panel A3c), and 0.12 (panel A3d).

Figure A4: Dynamic Ban Effects on Firms' Agro Exports: Quantity vs. Price (High vs Low Fertilizer Intensity)

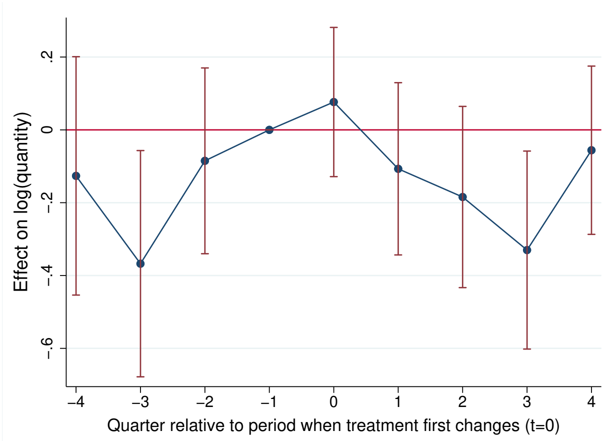

(a) Quantity Effect

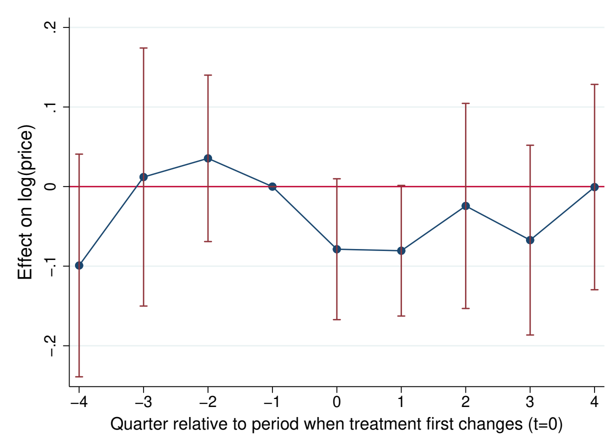

(b) Price Effect

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of fertilizer import bans on firms' agricultural exports, relative to the last quarter before the ban takes effect, using a firm-level panel. The sample covers the period between quarter 1 of 2019 and quarter 3 of 2022 (but coefficients are only reported from quarter 2 of 2020 to quarter 2 of 2022) and includes all firms who exported at least one of 23 crops for which we have fertilizer intensity data (in kilograms per hectare) during the 2017-2019 period. We compute each firm's fertilizer usage U_f as a weighted average of crops' fertilizer intensities with weights given by each crop's share in firm f's 2017-2019 export record. The treatment variable is a dummy T_{ft} which takes value one if there were any fertilizer import bans in place in quarter t *and* firm f has fertilizer intensity U_f above the 75th percentile. Firms below the 75th percentile serve as the implicit control group. The y-axis shows the effects of fertilizer import bans on the log quantity (panel A4a) or log average price (A4b) of agricultural exports, where negative values correspond to decreases in quantities/prices. The p-values of the joint significance tests for the placebo estimators are 0.07 (panel A4a) and 0.15 (panel A4b).

Figure A5: Dynamic Ban Effects on Agro Exports: Quantity vs. Price (IO linkage)

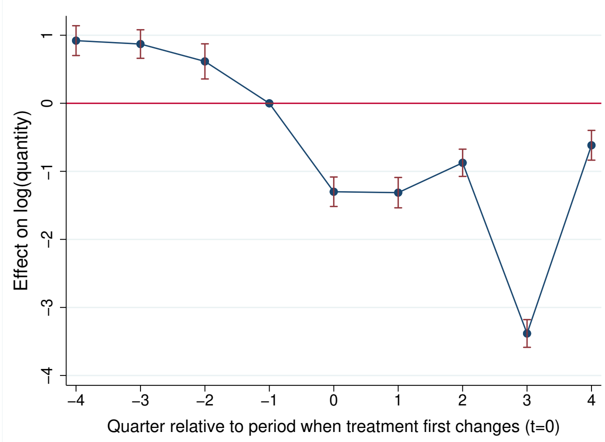

(a) Quantity Effect

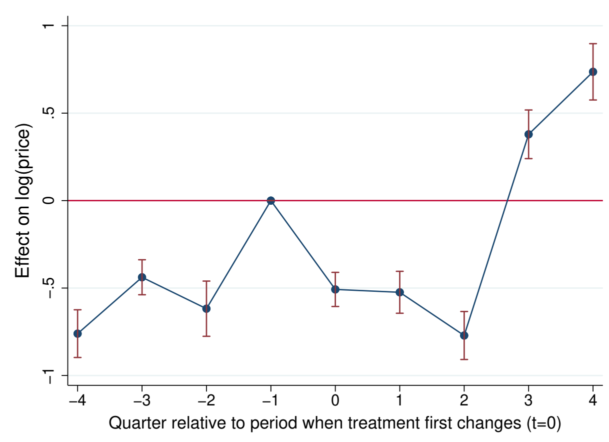

(b) Price Effect

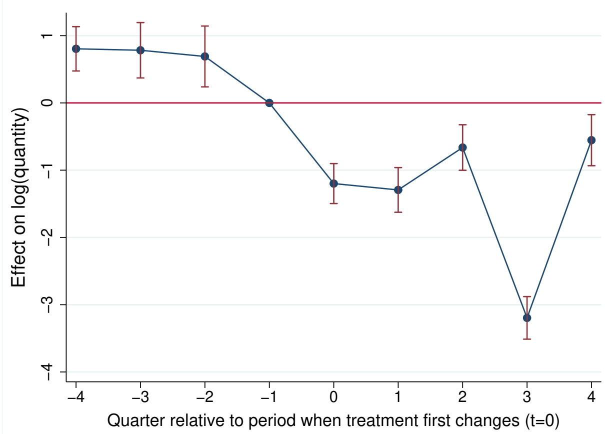

(c) Quantity Effect, excluding products with banned non-fertilizer inputs

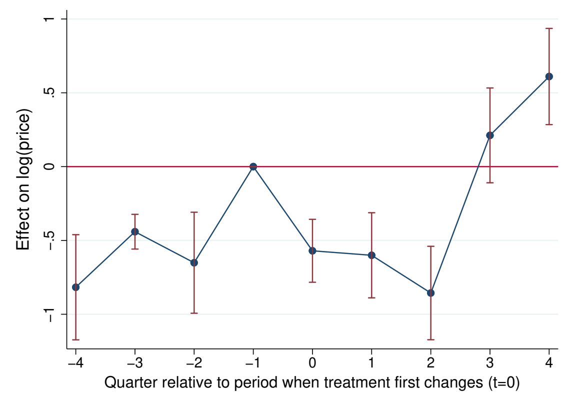

(d) Price Effect, excluding products with banned non-fertilizer inputs

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of import bans on agricultural exports, relative to the last quarter before the ban takes effect, using an HS8 product-level panel and methods described in section 4.1.1 to isolate the ban effects. The sample covers the period between quarter 1 of 2019 and quarter 3 of 2022 (but coefficients are only reported from quarter 2 of 2020 to quarter 2 of 2022). In panels A5c and A5d, the sample excludes products that had some non-fertilizer inputs (but no fertilizer inputs) banned during the period. The treatment variable is a dummy $sevfban_{ct}$ indicating whether import bans on the fertilizer inputs required by product c in quarter t were above the 75th percentile of severity, and the not-yet-treated products serve as the control group. The severity of the fertilizer import bans faced by a product c is defined as the (adjusted) count of HS8 fertilizer products that were banned in quarter t, with adjustment terms given by each fertilizer's unit requirements in the production of product c. The y-axis shows the effects of import bans on the log quantity (panels A5a and panels A5c) and log average price (panel A5b and A5d) of exports, where negative values correspond to a decrease in quantities/prices. The p-values of the joint significance tests for the placebo estimators are 0.00 (panel A5a), 0.00 (panel A5b), 0.00002 (panel A5c), and 0.00 (panel A5d).

Figure A6: Dynamic Ban Effects: Deseasonalized Imports

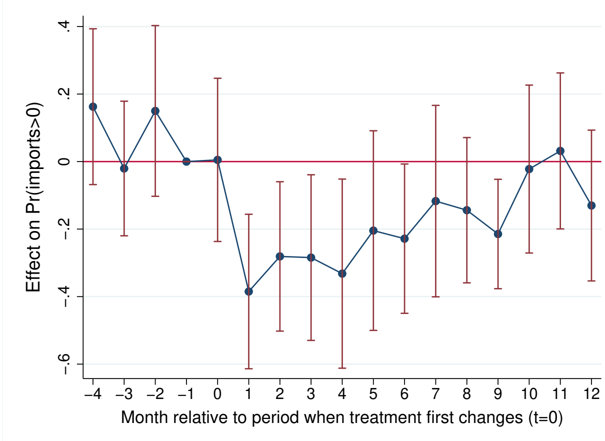

(a) Extensive Margin

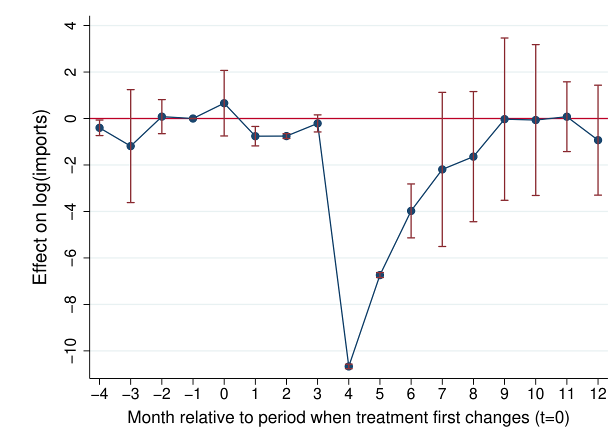

(b) Intensive Margin

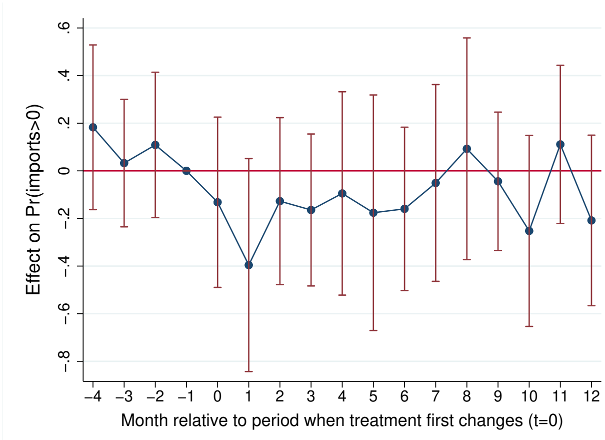

(c) Extensive Margin, fertilizer only

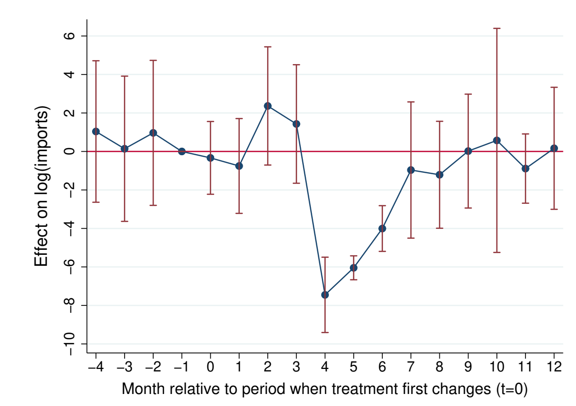

(d) Intensive Margin, fertilizer only

Notes: The figures show De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of import bans on (deseasonalized) imports, relative to the last month before the ban takes effect, using an HS8 product-level panel and methods described in section 4.1.1 to isolate the ban effects. The sample covers the period between January 2019 and September 2022 (but coefficients are only reported from January 2021 to May 2022) and excludes non-fertilizer HS8 products whose imports were ever banned. Additionally, in panels A6c and A6d, the sample also excludes all non-fertilizer products, even the ones who were never banned. For each product c and month $m \in \{1, ...12\}$, we compute the (log of the) historical average import value ($\log(M^h_{c,m})$) and the historical probability of nonzero imports ($P^h_{c,m}$) for that month m using years 2017-2019; then, for every product c and period t, we subtract the relevant historical average from log import value $\log(M_{ct})$ and from the indicator for nonzero imports $\mathbb{1}[M_{ct} > 0]$, thus obtaining deseasonalized versions of these two outcome variables. The treatment variable is a dummy ban_{ct} indicating whether product c was banned in 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), with negative values correspond to decreases in probabilities/values. The p-values of the joint significance tests for the placebo estimators were 0.454 (panel A6a), 0.03 (panel A6b), 0.774 (panel A6c), and 0.959 (panel A6d).

Figure A7: Dynamic Ban Effects on Firms: Deseasonalized Agricultural Exports (High vs Low Fertilizer Intensity)

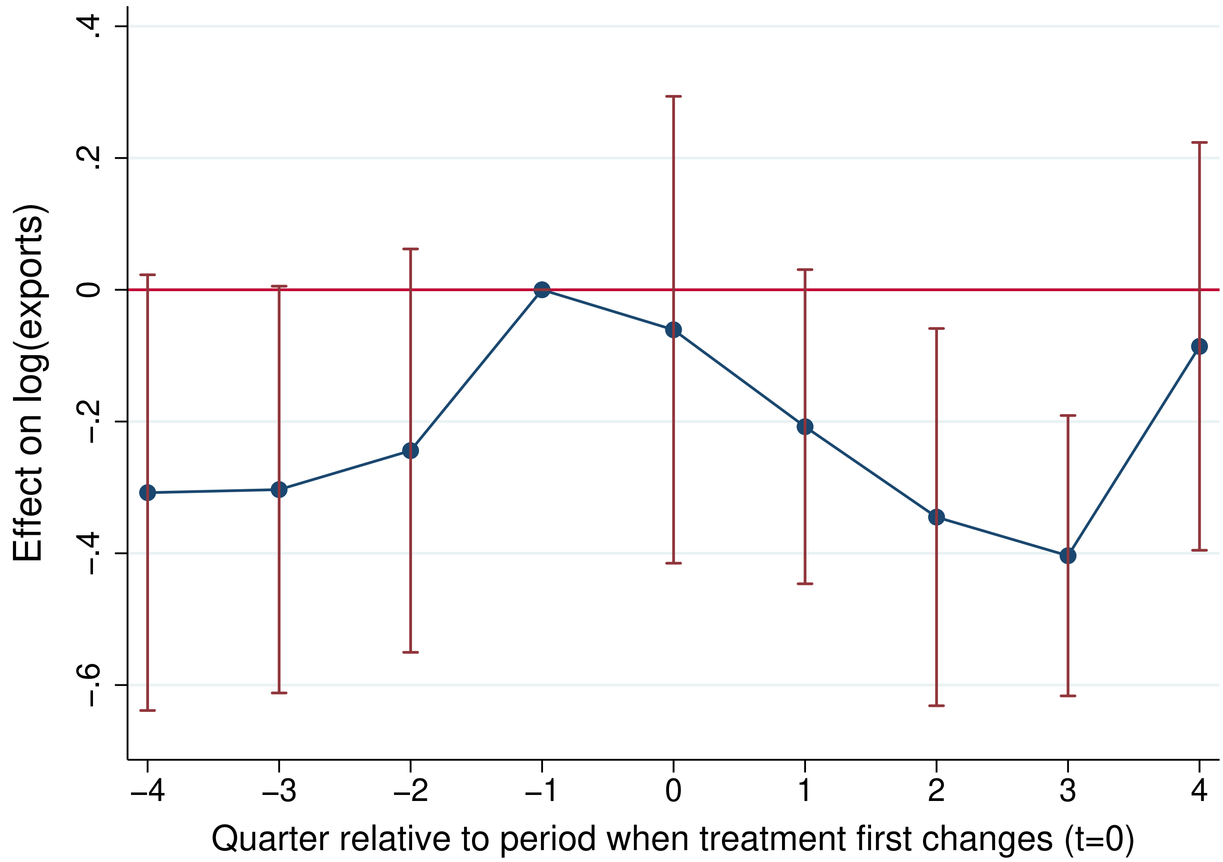

Notes: The figure shows De Chaisemartin and d'Haultfoeuille (2022) estimates of the effects of fertilizer import bans on the (deseasonalized) value of firms' agricultural exports, relative to the last quarter before the ban takes effect, using a firm-level panel. The sample covers the period between quarter 1 of 2019 and quarter 3 of 2022 (but coefficients are only reported from quarter 2 of 2020 to quarter 2 of 2022) and includes all firms who exported at least one of 23 crops for which we have fertilizer intensity data (in kilograms per hectare) during the 2017-2019 period. We compute each firm's fertilizer usage U_f as a weighted average of crops' fertilizer intensities with weights given by each crop's share in firm f's 2017-2019 export record. The treatment variable is a dummy T_{ft} which takes value one if there were any fertilizer import bans in place in quarter t *and* firm f has fertilizer intensity U_f above the 75th percentile. Firms below the 75th percentile serve as the implicit control group. For each firm f and quarter $q \in \{1, 2, 3, 4\}$, we compute the (log of the) historical average export value ($\log(X^h_{f,q})$) for that quarter q using years 2017-2019; then, for every firm f and period t, we subtract the relevant historical average from the log export value $\log(X_{ft})$, thus obtaining a deseasonalized version of the outcome variables. 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 0.123.

Figure A8: Event Study: Imports and the Arrival of Covid-19

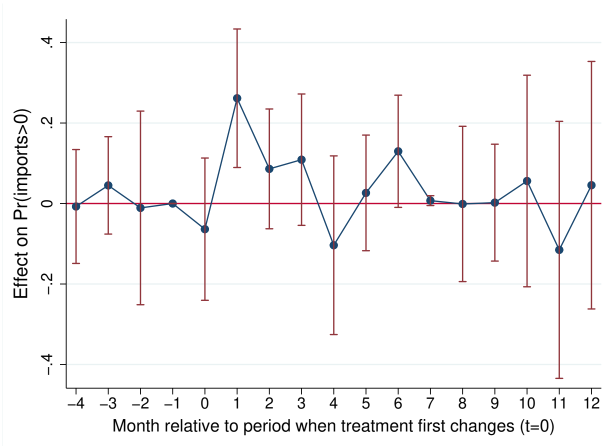

(a) Extensive Margin

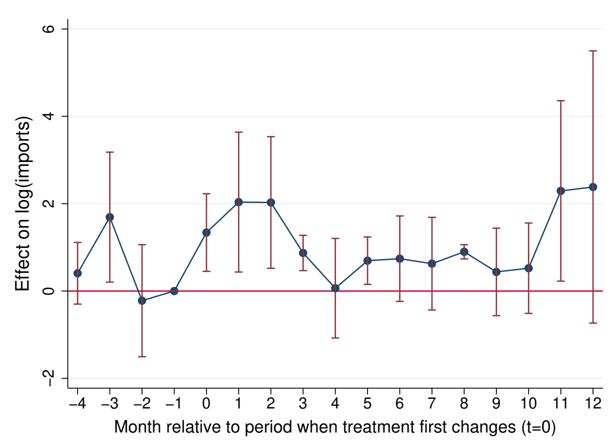

(b) Intensive Margin

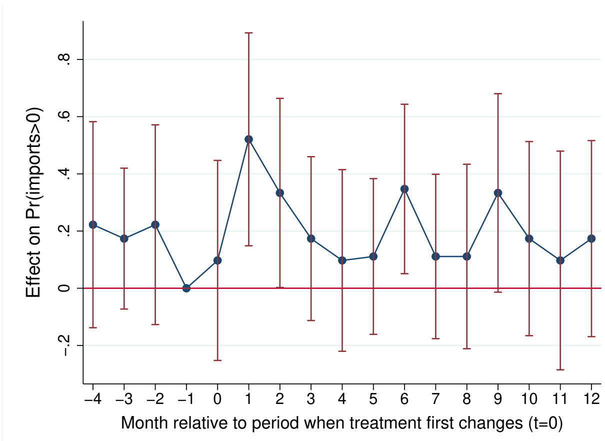

(c) Extensive Margin, fertilizer only

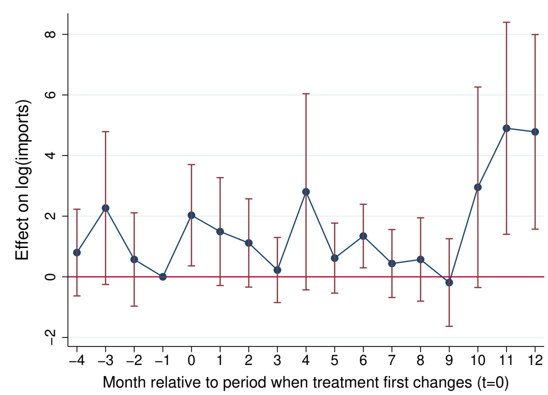

(d) Intensive Margin, fertilizer only

Notes: The figure shows De Chaisemartin and d'Haultfoeuille (2022) estimates of the differential effects of the Covid-19 pandemic on fertilizer imports, relative to the last month before the start of Covid-19, using an HS8 product-level panel and methods described in section 4.1.1 to isolate the effects. The sample covers the period between January 2019 and April 2021 (but coefficients are only reported from November 2019 to March 2021), and it excludes non-fertilizer HS8 products whose imports were ever banned. Additionally, in panels A8c and A8d, 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 A8a and A8c) or on the log value of imports (panels A8b and A8d), with negative values corresponding to decreases in probabilities/values. The p-values of the joint significance tests for the placebo estimators are 0.433 (panel A8a), 0.097 (panel A8b), 0.47 (panel A8c), and 0.357 (panel A8d).

Figure A9: Distribution of Food Expenditure Shares

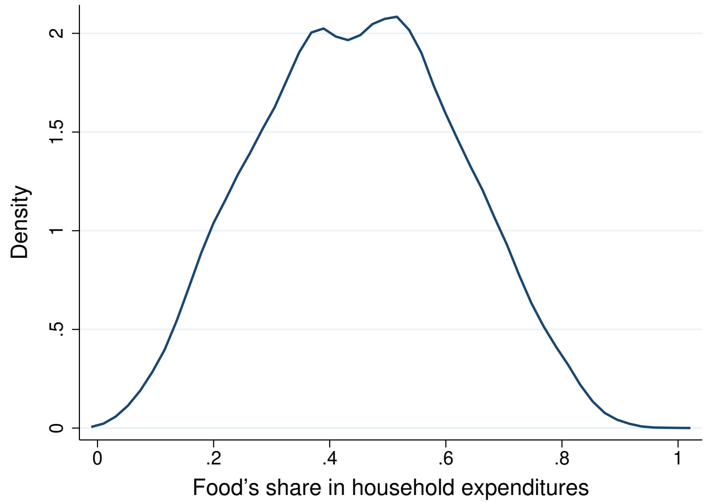

Notes: The figure shows kernel density estimates for the share of total household expenditure that is spent on food. A unit of observation is a household. The underlying data comes from the 2019 edition of Sri Lanka's Household Income and Expenditure Survey (HIES). Estimation uses an Epanechnikov kernel with optimal bandwidth (0.0213).

C Additional Empirical Findings

C.1 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 using data from the 2019 edition of the Household Income and Expenditure Survey to display the empirical distribution of these food expenditure shares across Sri Lankan households in 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 food's share of expenditures. This can be seen in Figure A10, which plots household income against food's share of household expenditures, showing a strong and clear negative correlation between the two variables.[77]

C.2 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 we can see in Figure A11a,

[77]The two variables plotted in Figure A10 are "residual" versions, i.e. deviations from their mean value in the household's district of residence. See the figure's bottom notes for details.

Figure A10: Association Between Food Share and Income

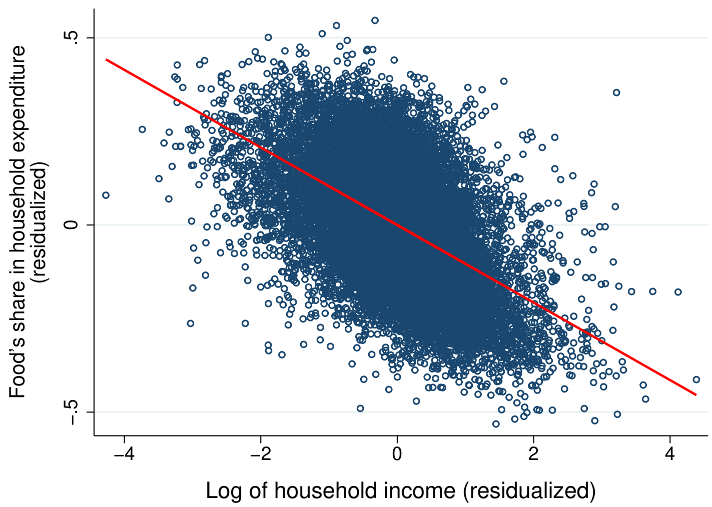

Notes: The figure plots a residualized version of the natural logarithm of household income against a residualized version of the share of total household expenditure that is spent on food. A unit of observation is a household. The underlying data comes from the 2019 edition of Sri Lanka's Household Income and Expenditure Survey (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 survey. A red regression line is overlaid on the graph representing the regression of residualized food share on residualized log household income (also estimated by WLS).

Figure A11: Distribution of Landholding Sizes

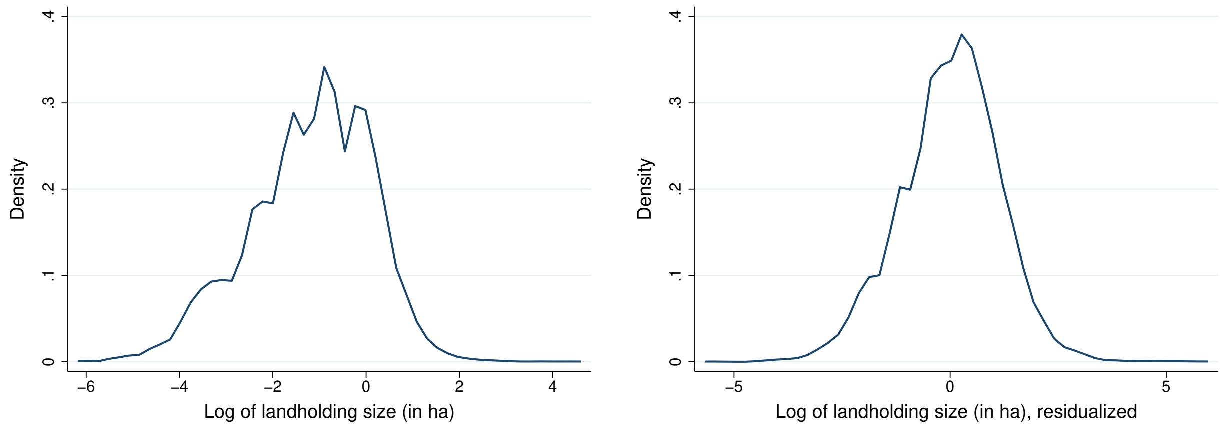

(a) Unadjusted

(b) Adjusted for Region

Notes: The figure shows kernel density estimates for the log of the size of a household's landholdings (measured in hectares). A unit of observation is a household. The underlying data comes from the 2016 edition of Sri Lanka's Household Income and Expenditure Survey (HIES). Panel A11a shows unadjusted estimates. Panel A11b 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 survey. Estimation uses Epanechnikov kernels with optimal bandwidth (0.2021 in Panel A11a and 0.1647 in Panel A11b).

which shows the distribution of (the natural logarithm of) landholding size. Note that this variation is not solely driven by differences across regions in average landholding size, as dispersion remains high after controlling for region (Figure A11b).

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 6.1.2), which we do using data from the 2016 edition of the Household Income and Expenditure Survey.[78]

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 $l\bar{L}_n$:

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

where h indexes the households in the sample, w_h is the household's statistical weight, and L^h is

[78] See Section 3.5 for details on the Household Survey data.

the number of hectares of owned land. Only households with positive landholdings are included in the sample.[79] We then estimate the dispersion parameter σ^2_{Ln} for district n as:

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

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 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 hectares) and N^F_n is the number of farmers in district n.[80]

C.3 Evolution of Rice Agriculture over Time

In this section, we focus on rice, Sri Lanka's main food crop. We combine cutting-edge remote sensing data (described in Section 3.4) with detailed weather data (described in Section A.4) to show the evolution of rice yields and cultivated area over time. Specifically, for each of the two growing seasons (Maha and Yala) we run the following regression:

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

where s indexes seasons, d indexes divisional secretariats (DSs), y indexes years, A^s_{dy} is the cultivated rice area in district d during season s of year y, θ^s_y and ω^s_d are sets of year and DS fixed effects (respectively), $\mathbf{W^s_{dy}}$ is a vector of weather variables, and ϵ^s_{dy} is an error term. We estimate regression (42) by Ordinary Least Squares with standard error clustered at the DS level.[81] We then repeat the process replacing cultivated area with yields as the dependent variable. The results are shown in Figure A12.

Because we control for weather and DS fixed effects, the year coefficients displayed on Figure A12 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

[79] We define landholdings to include only land that is owned by the household (as opposed to leased, rented, etc) and that does not have housing built on top of it.

[80] For details on the data on arable land and the number of farmers, see Sections 3.6 and 3.5, respectively.

[81] The vector of weather controls consists of: rainfall, number of rainy days, dry spells, solar radiation, VPD, average minimum and maximum temperatures. See Appendix Section A.4 for definitions and further details.

Figure A12: Rice Yields and Cultivated Area, by year

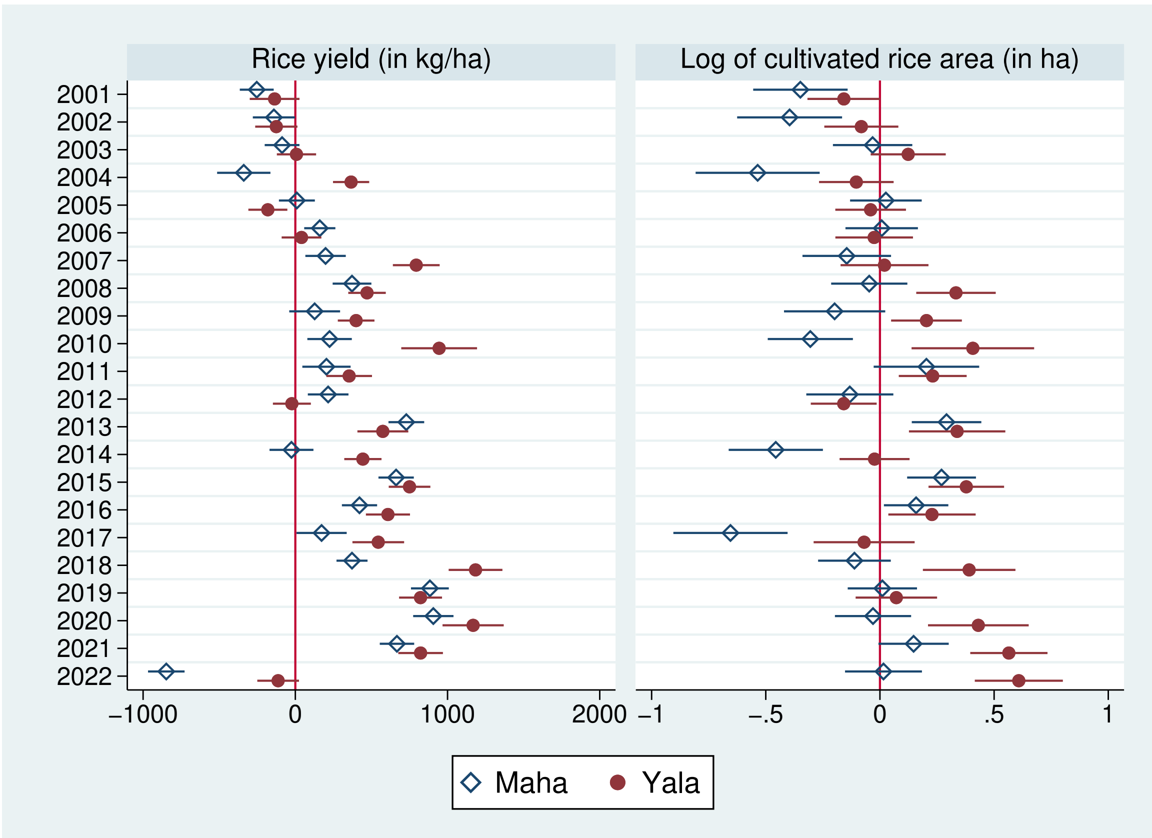

Notes: The figure plots coefficients and 95% confidence intervals from a series of regressions at the Divisional Secretariat (DS) level in which we regress rice yields (left panel) and the log of cultivated rice area (right panel) on a set of year dummies. Regressions are run separately for the Maha season (diamonds) and the Yala season (circles). Yields are measured in kilograms per hectare, and area is measured in hectares. Standard errors are 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). See Appendix Section A.4 for definitions of weather variables and further details.

after the fertilizer import ban) were 848kg/ha below average, with the null hypothesis of zero deviation from the average being strongly rejected at conventional significance levels. In other words, the first rice growing season after the ban had, by far, the worst productivity in at least 20 years.

On the other hand, 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 the assumption in our model 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 during 2022 Maha supports our hypothesis that the fertilizer ban had a major negative impact on Sri Lankan agriculture.

C.4 Relationship between Land Inequality and Agricultural Expenditures

Equation (24) in our model predicts a negative relationship at the region level between aggregate agricultural expenditure (X_n^A) and land inequality (σ_{Ln}^2) when holding aggregate income (E_n) and 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 if we return to the homothetic case by setting $\eta = 1$.

In this section, we perform a loose test of this prediction using data from the 2019 edition of the Household Income and Expenditure Survey.[82] Specifically, we use household survey 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 Table A4.

The coefficient on log total expenditures is 0.897 in Column (1) of Table A4. 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 -0.053, pointing to a negative association between inequality and agricultural expenditure, as predicted by equation (24). 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 (24), with very similar results.[83]

[82] See Section 3.5 for details on the Household Survey data.

[83] 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

Table A4: Food Expenditure and Land Inequality

	Dependent variable: Log of food expenditures	
	(1)	(2)
Land inequality ($\hat{\sigma}^2_{Ln}$)	-0.053	-0.055
	(0.048)	(0.049)
Log of total expenditures ($\ln(E_n)$)	0.897***	0.894***
	(0.026)	(0.029)
Constant	1.672***	1.715**
	(0.624)	(0.640)
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. The regressions are estimated by Ordinary Least Squares (OLS) using data from the 2019 edition of the Household Income and Expenditure Survey. 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 (24), 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 $\hat{\phi} = 0.0105$.

A literal interpretation of equation (24) would predict not only the qualitative relationship between inequality and agricultural expenditure but also the specific magnitude of the coefficient on σ^2_{Ln}, which would be $-\eta(1-\eta)/2$, or -0.113 if we use our $\hat{\eta}$ estimate from Table 1. Interestingly, this value is in the 95% confidence interval of the estimated inequality coefficient in both Columns (1) and (2) of Table A4. However, putting excessive credence in this similarity may not be warranted given the large standard errors of the inequality coefficient.

D Inversion Details

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

estimate of ϕ as displayed on Table 1 and explained in Section 6.1.1.

D.1 Inversion procedure

It is possible to rewrite the model's three market clearing equations (32, 33, 34) to obtain the following system of equations:[84]

$$F_{RoW} = \frac{1}{p^f_{RoW}} \sum_{i=1}^{R} \sum_k \frac{\gamma^f_k}{S_k} p_{ik} Q_{ik} \tag{43}$$

$$Q_{ik} = \frac{1}{p_{ik}} \sum_n X^A_n b_k (\frac{P_{nk}}{P^A_n})^{1-\sigma_A} b_{ik} (\frac{\tau^A_{ni,k} p_{ik}}{P_{nk}})^{1-\sigma_K} \tag{44}$$

$$w_i N^W_i = (\sum_k \frac{\gamma^n_k}{S_k} p_{ik} Q_{ik}) + \sum_n (E_n - X^A_n)(\frac{\tau^M_{ni} w_i}{T^M_i P^M_n})^{1-\sigma_M} \tag{45}$$

In these three 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.[85] We solve this system of equations numerically to recover all unobserved variables ($\boldsymbol{T}^M$, $\boldsymbol{b}$, F_{RoW}, $\boldsymbol{p}^-$, $\boldsymbol{P}$, $\boldsymbol{R}$, $\boldsymbol{E}$, $\boldsymbol{X}$, $\boldsymbol{\beta}$), and then recover agricultural productivities ($\boldsymbol{T}^A$) through equation (14) which links observed crop production to underlying productivity, crop prices, and input prices.

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

D.2 Data Requirements for Inversion

To solve the system of equations in Appendix Section D.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 basic 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 seven crops (i.e. $K = 7$) for which we have data on both fertilizer intensity and producer prices, namely: cinnamon, cloves, groundnut, maize, onions,

[84] More specifically, to obtain the system of equations substitute equation (20) into fertilizer market clearing condition (32), equations (15), (27), and (29) into crop market clearing condition (33), and equations (26) and (31) into labor clearing condition (34).

[85] Blue variables can be recovered from black and red variables by using equations (13), (16), (24), (25), (28), (30), and (31).

potatoes, and rice.[86] We choose to use year 2019 for inversion because it is prior to the 2021 fertilizer import ban (and to other potentially influential events as the Covid-19 pandemic and the Russian Federation-Ukraine war) while also offering rich household survey data.

Preference parameters (σ_A, σ_K, σ_M) are taken from Table 1 and production parameters ($\{\gamma_k^f$, γ_k^n, $\gamma_k^l\}_k$) from Table 2. Crop prices ($\{p_{ik}\}_{i,k}$) for each crop and district in 2019, measured in Sri Lankan Rupees per kilogram (LKR/kg), are taken from the 2016-2019 Bulletin of Prices (see Section 3.7 for details). The 2019 crop-district-level production ($\{Q_{ik}\}_{i,k}$), in kilograms, and cultivated area ($\{L_{ik}\}_{i,k}$), in hectares, are from the Department of Census and Statistics (see Section 3.6). District-level annual wages ($\{w_i\}_i$), in LKR, and numbers of workers and farmers ($\{N_i^W$, $N_i^F\}_i$) are computed from the 2019 edition of the Household Income and Expenditure Survey (see Section 3.5). Transportation costs for agriculture ($\{\tau_{ni,k}^A\}_{n,i,k}$) and manufacturing ($\{\tau_{ni}^M\}_{n,i}$) are set to one.[87]

For the Rest of the World (RoW), data on 2019 crop prices ($\{p_{RoW,k}\}_k$), in US dollars per kilogram, production ($\{Q_{RoW,k}\}_k$), in tonnes, and cultivated area ($\{L_{RoW,k}\}_k$), in hectares, are from the Food and Agriculture Organization of the United Nations (FAO). The wage in RoW (w_{RoW}) is set to the world's income per capita in 2019 (in USD), which we take from the World Bank's World Development Indicators. The numbers of workers (N_{RoW}^W) and farmers (N_{RoW}^F) in RoW are estimated by combining data from the International Labour Organization (ILO) with data from the Brazilian Agricultural Census (whose ratio of farmers to agricultural workers we assume holds for the world as a whole). All monetary variables expressed in USD are converted into LKR using the 2019 exchange rate. Section A.3 provides more details on all RoW variables.

We assume the transportation cost of fertilizer is one ($\{\tau_{ni}^f\} = 1$ for all i, n) in the baseline 2019 equilibrium.[88] The international price of fertilizer, p_{RoW}^f =\$0.31/kg, comes from our trade data (see Section 3.2 and Figure 1), and is converted into LKR using the 2019 exchange rate.

E Equivalent Variation

Since we are interested in the welfare effects of the fertilizer import ban, we use our estimated quantitative model 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 (22) allows us to see that indirect utility can be negative for some combinations of prices and income.[89] As a result, it is not straightforward to express utility as "real income"

[86] See Sections 3.3 and 3.7 for details on fertilizer intensity and producer price data, respectively.

[87] The next version of this paper will feature positive trade costs in these two sectors.

[88] This transportation cost is increased to higher values in the counterfactual exercises of Section 7 to study the effects of the fertilizer import ban.

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

(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 equivalent variation (EV).

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

$$EV = V^{-1}(v_1; P^A_{n,0}, P^M_{n,0}) - y_0 \tag{46}$$

where v_1 is the agent's utility in the final equilibrium, y_0 is her income in the initial equilibrium, $(P^A_{n,0}, P^M_{n,0})$ are the prices indices in the initial equilibrium, and V^{-1} is the inverse of the indirect utility function (22) with respect to the income argument (y).[90] 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 as well as across farmers with different landholding sizes.

Thus, to assess the welfare changes faced by different agents, we define five 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):[91]

1. **Worker**: all workers in the district are identical and earn the same wage w_n. Their EV is obtained by substituting $v_1 = V^W_{n,1}$ and $y_0 = w_{n,0}$ in equation (46):

$$EV^W_n = V^{-1}(V^W_{n,1}; P^A_{n,0}, P^M_{n,0}) - w_{n,0}$$

2. **Median farmer**: defined as the farmer with a landholding size equal to the median landholding size in the district, which is $\exp(\mu_n)$ due to the log-normality of the land distribution (see Section 5.3.1). Her initial income level is thus $y_0 = \exp(\mu_n)R_{n,0}/L_i$, and her final utility v_1 is obtained by evaluating indirect utility function (22) at final income $y_1 = \exp(\mu_n)R_{n,1}/L_i$ and final prices $(P^A_{n,1}, P^M_{n,1})$. Her EV can thus be expressed as:

$$EV^{medF}_n = V^{-1}(V^{medF}_{n,1}; P^A_{n,0}, P^M_{n,0}) - \exp(\mu_n)R_{n,0}/L_i,$$

[90] Simple algebric 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 $(P^A_{n,0}, P^M_{n,0})$, see Sections 5.3.3 and 5.3.4.

[91] Some EV expressions also use welfare variables $(V^W_n, V^F_n, V^{avg}_n)$ as inputs. For the definitions of these variables, see Section 5.5.

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

3. **Average farmer**: defined as the farmer with a landholding size equal to the average landholding size in the district, 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 farmer's EV is:

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

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

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 so that relative welfare losses are more comparable across regions and agent types. This relative EV can be computed by dividing the EV by the corresponding initial income level: $\frac{EV}{y_0}$.

F Proofs

F.1 Aggregate Agricultural Expenditure

In Section 5.3.2, we claim that agricultural expenditure share in region n is given by equation (24):

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

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, and E_n is aggregate income.

To derive this equation, we first separately derive aggregate agricultural expenditure by workers ($X_n^{A,W}$) and by farmers ($X_n^{A,F}$), then add them up. Aggregate expenditure by workers ($X_n^{A,W}$) is given by multiplying the number of workers (N_n^W), earnings per worker (i.e. the wage w_n), and the fraction of income spent on agricultural goods (given by equation 23 with income y evaluated at wages w_n):

$$X_n^{A,W} = N_n^W w_n \xi^A(w_n, P_n^A, P_n^M) = \phi N_n^W w_n + \nu((P_n^A)^\phi (P_n^M)^{1-\phi})^\eta N_n^W w_n^{1-\eta} \tag{47}$$

For an individual farmer h who lives in region n and earns land rent R^h, expenditure on agriculture is $R^h \xi^A(R^h, P_n^A, P_n^M)$, the product of her income R^h and her agricultural expenditure share (see equation 23). Substituting equation (23) 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:

$$\bar{X}_n^{A,F} = \phi r_n + \nu((P_n^A)^\phi (P_n^M)^{1-\phi})^\eta \mathbb{E}[(R^h)^{1-\eta}] \tag{48}$$

From equation (21), 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 is also log-normal:

$$(R_n^h)^a \sim \text{log-N}(a[\mu_n + \ln(R_n) - \ln(L_n)], a^2\sigma_{Ln}^2) \tag{49}$$

for any constant $a \neq 0$. Evaluating equation (49) with $a = 1-\eta$ and using the relationship between the parameters of a log-normal distribution and its average, we conclude that:[92]

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

By using the properties of the exponential and rearranging terms, this expression can be rewritten as follows:[93]

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

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 $\frac{L_n}{N_n^F}$.[94] Using this fact, cancelling terms, and applying the definition of r_n, we can rewrite the previous equation as follows:

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

[93] 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.

[94] 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.

$$\mathbb{E}[(R^h)^{1-\eta}] = r_n^{1-\eta} \exp\Big(-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}\Big) \tag{50}$$

Substituting (50) into equation (48) 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:

$$X_n^{A,F} = \bar{X}_n^{A,F} \times N_n^F = \phi r_n N_n^F + \nu((P^A)^\phi (P^M)^{1-\phi})^\eta N_n^F r_n^{1-\eta} e^{-\eta(1-\eta)\frac{\sigma_{Ln}^2}{2}} \tag{51}$$

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

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

Using equations (51) and (47) to substitute for $X_n^{A,F}$ and $X_n^{A,W}$, rearranging terms, and noting that $R_n = r_n N_n^F$ and $E_n = N_n^W w_n + R_n$ (see equation 25), we obtain equation (24).

F.2 Average Farmer Welfare

In Section 5.5, we claim that average farmer welfare in region i is given by equation (36):

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

To prove this equation, consider that the indirect utility of an individual farmer h who lives in region i and earns land rent R^h is given by $V(R^h, P_i^A, P_i^M)$. Using equation (22), taking the expectation operation across farmers, and noting that P_i^A and P_i^M are the same for all the region's farmers, average farmer indirect utility can be written as:

$$V_i^F \equiv \mathbb{E}(V(R^h, P_i^A, P_i^M)) = \frac{1}{\eta}\Big(\frac{1}{(P^A)^\phi (P^M)^{1-\phi}}\Big)^\eta \mathbb{E}[(R^h)^\eta] - \nu \ln\Big(\frac{P^A}{P^M}\Big) \tag{52}$$

So our remaining task is to find an expression for $\mathbb{E}[(R^h)^\eta]$. From equation (49) 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:

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

Using this expression to substitute $\mathbb{E}[(R^h)^\eta]$ into equation (52) gives us equation (36).

G Parameter Estimates

In Section 6.1.1, we use Two-Stage Least Squares (2SLS) approaches to estimate the parameters (η, σ_A, σ_K) that govern consumer preferences. Appendix Table A5 shows complete estimation results, including the estimated coefficients, standard errors, and first-stage results.

Table A5: 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)	$\Delta^q \ln(X^A_{Row,i})$ (3)
Log of household income $[\hat{\eta}]$	-0.656***		
	(0.129)		
Log of crop price $[1-\hat{\sigma}_A]$		-0.714***	
		(0.153)	
$\sum_k sh^0_{ik} \times \Delta^q \ln(p_{RoW,i,k})$			-2.627
$[1-\hat{\sigma}_K]$			(1.809)
Household size	0.135***		
	(0.029)		

	Panel B: First Stage		
	Dependent variable:		
	Log of household income (1)	Log of crop price (2)	$\sum_k sh^0_{ik} \times \Delta^q \ln(p_{RoW,i,k})$ (3)
Lottery/ad hoc income	8.81e-07***		
	(1.37e-07)		
Disaster/other relief payments	8.16e-07***		
	(1.42e-07)		
Regional potential yield > 0		-.062***	
		(.0044)	
z^q_i			2.461**
			(5.761)
F-statistic	37.2***	199.4***	4.68**
p-value	0.0000	0.0000	0.0313
Fixed Effects?	PSU	Crop, HH	District, quarter
Clustered SEs?	by PSU	by HH	No
Weighted?	Yes	Yes	No
Observations	19,910	91,416	379

Notes: *** denotes significance at the 1% level, ** at the 5% level. Columns (1), (2), and (3) displays results for the estimation of equations (38), (39), and (40), respectively, by 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 Household Income and Expenditure Survey. See Section 6.1.1 for details on data sources, identification strategies, and variable definitions.

Table A6: Regional Ban Effects

District	Worker	Farmer Median	Farmer Average	Farmer Rep	Rep Agent
Ampara	-1.30%	-5.45%	-5.45%	-5.45%	-2.44%
Anuradhapura	-1.25%	-5.26%	-5.26%	-5.26%	-2.64%
Baddulla	-0.36%	-5.88%	-5.88%	-5.88%	-1.34%
Batticaloa	-0.72%	-5.90%	-5.90%	-5.90%	-1.62%
Colombo	-0.005%	-5.90%	-5.89%	-5.89%	-0.04%
Galle	-0.13%	-3.94%	-3.93%	-3.94%	-0.55%
Gampaha	-0.02%	-6.10%	-6.09%	-6.09%	-0.12%
Hambantota	-0.86%	-5.43%	-5.43%	-5.43%	-2.06%
Jaffna	-0.08%	-4.16%	-4.16%	-4.16%	-0.40%
Kaluthara	-0.05%	-5.29%	-5.29%	-5.29%	-0.29%
Kandy	-0.07%	-4.54%	-4.54%	-4.54%	-0.34%
Kegalle	-0.07%	-4.95%	-4.95%	-4.95%	-0.37%
Killinochchi	-2.07%	-5.13%	-5.13%	-5.13%	-3.08%
Kurunagala	-0.50%	-5.94%	-5.94%	-5.94%	-1.62%
Mannar	-1.70%	-5.23%	-5.23%	-5.23%	-2.68%
Matale	-0.38%	-5.73%	-5.73%	-5.73%	-1.30%
Matara	-0.20%	-4.11%	-4.11%	-4.11%	-0.75%
Monaragala	-1.10%	-4.52%	-4.52%	-4.52%	-2.32%
Mullaittivu	-3.17%	-4.05%	-4.05%	-4.05%	-2.56%
Nuwareliya	-0.08%	-8.44%	-8.44%	-8.44%	-0.49%
Polannaruwa	-2.64%	-4.84%	-4.84%	-4.84%	-3.56%
Puttalam	-0.19%	-4.86%	-4.86%	-4.86%	-0.70%
Rattnapura	-0.13%	-4.71%	-4.71%	-4.71%	-0.59%
Trincomalee	-1.09%	-5.70%	-5.70%	-5.70%	-2.21%
Vavuniya	-0.70%	-5.62%	-5.62%	-5.62%	-1.90%
Average	-0.76%	-5.27%	-5.27%	-5.27%	-1.48%

Notes: for each district, the table shows the % welfare changes of five types of agents between the baseline equilibrium and the counterfactual equilibrium. The median (average) farmer is defined as a farmer whose landholding size is equal to the median (average) in 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 an hypothetical agent whose income is just enough to provide her with the district's average utility (V_i^{avg}). Welfare changes are defined as Equivalent Variations (EV): starting from the agent's baseline income, the EV is the % change in income that would leave her with the same utility level as she has in the counterfactual equilibrium, keeping prices fixed at baseline levels.