

Coordination and Commitment in International Climate Action: Evidence from Palm Oil

Allan Hsiao *

Stanford University

June 13, 2025

([most recent version here](#))

Weak environmental regulation has global consequences. When domestic regulation fails, the international community can target emitters with trade policy. I develop a dynamic empirical framework for evaluating trade policy as a substitute for domestic regulation, and I apply the framework to the market for palm oil, a major driver of deforestation and global CO₂ emissions. Relative to business as usual, a domestic production tax of 50% reduces CO₂ emissions by 7.4 Gt from 1988 to 2016, amounting to 0.26 Gt annually. Coordinated, committed import tariffs of similar magnitude reduce emissions by 5.4 Gt over the same period. The cost of these import tariffs is only \$15 per ton of CO₂, even accounting for compensating transfers that recognize welfare losses for producing countries. Without coordination and commitment, import tariffs have more limited effects. Alternative policies include domestic export taxes, which are fiscally appealing independent of emission concerns, and a carbon border adjustment mechanism, which encourages domestic regulation.

*Email: ajhsiao@stanford.edu. I am grateful to Nikhil Agarwal, Esther Duflo, Ben Olken, and Rob Townsend for their guidance and encouragement throughout this project. I have also benefited from helpful conversations with Daron Acemoglu, David Atkin, Clare Balboni, Abhijit Banerjee, Arnaud Costinot, Dave Donaldson, Namrata Kala, Jing Li, Rachael Meager, Will Rafey, Mar Reguant, Tobias Salz, and Tavneet Suri. I acknowledge generous support from the National Science Foundation Graduate Research Fellowship, Jerry Hausman Graduate Dissertation Fellowship, George and Obie Shultz Fund, and Weiss Fund for Research in Development Economics.

1 Introduction

Carbon emissions have global consequences. The international community may therefore wish to intervene when domestic regulation fails. Indeed, free-riding incentives, political constraints, administrative limits, and potential corruption each undermine domestic regulation (Oates and Portney 2003, Burgess et al. 2012, Oliva 2015). The conventional approach attempts to address these difficulties, such as by improving enforcement (Duflo et al. 2018), but doing so at scale can be infeasible. Trade policy offers an alternative, circumventing these obstacles by targeting the prices that carbon emitters receive in world markets.

How effective is trade policy as a substitute for direct regulation? I develop a dynamic empirical framework to answer this question quantitatively. I highlight two challenges: (1) a coordination problem because of leakage to unregulated markets and (2) a commitment problem because regulation is not statically optimal once emissions are sunk.¹ I apply this framework to studying the palm oil industry, which accounts for 5% of global CO₂ emissions from 1990 to 2016 (figure 1). By comparison, the European Union (EU) accounts for 11% and Russia for 6% over the same period.

Palm oil is an important empirical setting. The industry is a major polluter: land clearing for palm oil plantations threatens carbon-rich peatland forests in Indonesia and Malaysia, which together account for 84% of global palm oil production. At the same time, the industry generates substantial domestic profits that have lifted millions out of poverty (Edwards 2019). This paper informs an active debate on whether foreign governments should intervene with trade policy. The leading example is the EU Regulation on Deforestation-free products (EUDR), which will soon restrict EU imports of palm oil (OJEU 2023). I quantify emission reductions under such trade policy intervention, as well as the losses that Indonesia and Malaysia might claim as payment for ecosystem services.

I characterize palm oil demand with an almost ideal demand system (Deaton and Muellbauer 1980) and annual panel data on vegetable oil consumption by consumer market. The model explicitly captures substitution between palm oil and other veg-

¹ “Leakage” arises under incomplete regulation. Regulation reduces consumption in regulated markets. But in doing so, regulation also reduces world prices and raises consumption in unregulated markets. This response of unregulated markets attenuates the net effect on global consumption.

etable oils in response to price changes. Demand estimation applies iterated linear least squares, as in [Blundell and Robin \(1999\)](#). Prices are endogenous. I instrument for palm oil prices with weather shocks to palm oil production, and I instrument for other vegetable oil prices with weather shocks to other vegetable oil production. These instruments act as supply shifters. I estimate palm oil demand elasticities of 0.7 to 0.9, which indicate relatively inelastic demand for this staple food product. These estimates determine the losses from non-coordination, as leakage is exacerbated by elastic demand in unregulated markets.

I characterize palm oil supply with a dynamic discrete-continuous choice model and granular satellite data on palm oil production over time and space. The model explicitly captures differential responses to long- and short-run price changes. In the model, forward-looking firms invest in mills and plantations to produce palm oil for sale in world markets. I consider two margins of investment. On the extensive margin, firms make a discrete choice to build a mill or not. On the intensive margin, firms make a continuous choice over how much land to deforest and develop into plantations. Deforestation releases carbon emissions.

Supply estimation combines the continuous and discrete Euler methods of [Hall \(1978\)](#) and [Scott \(2013\)](#). Continuation values difference out, and estimation simplifies to linear regression with instruments. That is, I estimate the model without solving it. For identification, I combine variation in world prices over time with variation in yields across space. Revenues are the product of prices and yields. Thus, if supply is elastic, then high-yield plantations respond more strongly to prices than low-yield plantations. If supply is instead inelastic, then high- and low-yield plantations have similarly muted responses. Prices are again endogenous. I instrument for palm oil prices with total vegetable oil consumption and weather shocks to other vegetable oil production. These instruments act as demand shifters. Total consumption raises the category budget for vegetable oils overall, and weather shocks affect residual demand for palm oil. I estimate palm oil supply elasticities of 2.9 in the long run and 1.3 in the short run, reflecting that firms respond weakly to temporary price changes. These estimates determine the losses from non-commitment, as temporary regulation induces temporary price changes with limited effects.

For counterfactuals, I quantify the global impacts of regulation. I simulate direct regulation with production taxes, as well as trade policy with import tariffs, export

taxes, and a carbon border adjustment mechanism. For each policy, I solve the model for equilibrium prices, production, and consumption. Emissions depend on the spatial distribution of plantation development, which I model, and carbon stocks, which I observe. Welfare in each market is the sum of consumer surplus, producer surplus, and government revenue. For a given social cost of carbon, global social welfare is the sum of welfare across markets and the value of global emission reductions. I find that a palm oil production tax of 50% can reduce CO₂ emissions by 7.4 Gt over the study period from 1988 to 2016, relative to business as usual. If feasible, this production tax generates net welfare gains for Indonesia and Malaysia by improving their terms of trade. By comparison, EU-led import tariffs of similar magnitude can reduce emissions by 5.4 Gt over the same period. The cost to the EU is only \$15 per ton of CO₂, even accounting for compensating transfers that recognize profit losses for Indonesia and Malaysia.

Import tariffs rely on an EU that can coordinate across importers and commit to long-run enforcement. Coordination and commitment are difficult. The challenge with coordination is free-riding in two forms. Import tariffs reduce both global emissions and world prices, such that unregulated importers enjoy double benefits without bearing the burden of regulation. The challenge with commitment is the temptation to eliminate import tariffs once emissions are sunk. Present-biased governments may find it difficult to resist this temptation, seeking short-run gain at long-run cost. I find that emission reductions are smaller and substantially costlier when either coordination or commitment fails. Neither is independently sufficient for achieving the largest environmental gains. Both are necessary.

I consider several alternative policies. First, if international coordination fails, the EU can act unilaterally. Unilateral action can still achieve 1 Gt of abatement over the study period, although the cost to the EU rises to \$50 per ton of CO₂. Second, Indonesia and Malaysia can impose export taxes. Export taxes target the same goods as import tariffs and thus achieve the same emission reductions. Export taxes require enforcement only at international ports, unlike direct regulation. And export taxes have fiscal appeal because they generate government revenue at the expense of foreign consumers, while also sparing domestic consumers. Third, importers can implement a carbon border adjustment mechanism, which combines import tariffs with a credit for domestic regulation. This credit strengthens the fiscal incentive for Indonesia and

Malaysia to regulate.

This paper develops a new dynamic empirical framework for assessing green trade policy. I build on a rich literature that studies environmental regulation and trade, where free-riding and leakage motivate carbon coalitions ([Nordhaus 2015](#), [Böhringer et al. 2016](#), [Farrokhi and Lashkaripour 2025](#)) and border adjustment taxes ([Markusen 1975](#), [Copeland and Taylor 1994, 1995](#), [Hoel 1996](#), [Rauscher 1997](#), [Fowlie 2009](#), [Elliott et al. 2010](#), [Fowlie et al. 2016](#), [Kortum and Weisbach 2017, 2024](#)), and where trade policy influences environmental incentives ([Shapiro 2021](#), [Harstad 2024](#)). I also build on a literature studying commitment in environmental regulation ([Marsiliani and Renström 2000](#), [Abrego and Perroni 2002](#), [Helm et al. 2003](#), [Brunner et al. 2012](#), [Harstad 2020](#), [Acemoglu and Rafey 2023](#)). I quantify the challenges of coordination and commitment jointly and in an important empirical setting. By focusing on one industry, I can leverage detailed microdata to capture rich dynamics and fine-grained spatial heterogeneity.

Methodologically, I build on models of industry dynamics in the tradition of [Hopenhayn \(1992\)](#) and [Ericson and Pakes \(1995\)](#). I draw on a growing literature, formalized by [Aguirregabiria and Magesan \(2013\)](#), [Scott \(2013\)](#), and [Kalouptside et al. \(2021\)](#), that develops Euler conditional choice probability methods for estimating dynamic discrete choice models. Using techniques from [Hotz and Miller \(1993\)](#) and [Arcidiacono and Miller \(2011\)](#), this literature adapts classic continuous Euler methods from [Hall \(1978\)](#) and [Hansen and Singleton \(1982\)](#) to the discrete setting. I combine discrete and continuous Euler techniques to estimate a dynamic discrete-continuous choice model of entry and investment. Relative to other such models, including [Blevins \(2014\)](#), [Iskhakov et al. \(2017\)](#), and [Murphy \(2018\)](#), I offer a simple estimation strategy that is computationally light and straightforward to implement.

More broadly, trade policy enables regulation in otherwise low-regulation environments. For deforestation, trade policy does not rely on domestic governments that are willing and able to enforce regulation, unlike domestic policies ([Souza-Rodrigues 2019](#), [Assunção et al. 2023](#), [Araujo et al. 2024](#), [Burgess et al. 2024](#), [Domínguez-Iino 2025](#)) or conservation contracting ([Harstad 2012](#), [Harstad and Mideksa 2017](#)). Trade policy also scales readily, unlike direct payments for ecosystem services ([Jayachandran et al. 2017](#), [Edwards et al. 2020](#)). I show that trade policy can greatly reduce emissions in an industry that is crucial in the fight against climate change.

2 Background

Palm oil is a major source of global carbon emissions. Production is concentrated in Indonesia and Malaysia, where slash-and-burn practices have transformed the natural landscape. Sweeping plantations emerge from widespread deforestation, including of the peatland forests prevalent in the region. These forests house vast amounts of carbon in the form of peat, with layers of decomposing organic matter that extend as deep as ten meters belowground.² Palm-driven deforestation is thus particularly consequential, as it destroys both tree biomass and peat deposits. Figure 1 shows that palm emissions account for more CO₂ from 1990 to 2016 than the entire Indian economy, with peat destruction generating the vast majority of emissions.

Palm oil production begins with the planting of oil palm seedlings, which mature into trees. These trees bear fruit after three years and continue to do so over a lifespan of 30 years. Plantations harvest fresh fruit bunches that mills process into palm oil and palm kernel oil, with further processing by refineries. Roughly 90% of the oil in palm fruit is extracted from the flesh as palm oil, while the remaining 10% is extracted from the seed as palm kernel oil. These oils are exported widely. Indonesia and Malaysia account for 84% of global production and 89% of exports (table 1). Production is unconcentrated at the firm level, with the largest firm accounting for only 4% of global production (POA 2017).

Plantations and mills operate in tandem, as unmilled fruit decays within one day of harvest and is not consumed directly. For industrial plantations, which are 60% of production, vertical integration links plantations and mills directly. For smallholder plantations, which are 40% of production, vertical contracting creates similar links. Smallholders receive investment support from mills, which are nearly all industrial, in exchange for exclusive contracting (Cramb and McCarthy 2016). Mills exercise market power in setting contract terms, as smallholders face credit constraints and crop perishability that limit their bargaining power. Mills thus extract rents from plantations. If mills extract rents fully, then vertical integration and contracting coincide.³

² Converting peatlands to croplands involves draining peatlands and clearing the land with fire. Even without clearing by fire, unsubmerged peat releases carbon as it decomposes, and dried-out peat is likely to ignite from slash-and-burn activity in surrounding areas.

³ Indeed, market power over smallholder farmers is common in agricultural value chains (Bergquist

Figure 1: Emissions




Figure 1a compares Indonesian and Malaysian palm oil to the top emitters, accounting for land-use change. Palm emissions are 5.45% of global emissions from 1990 to 2016. Figure 1b separates palm emissions from tree biomass and peat deposits. Emissions are in gigatons of CO₂.

Consumption takes many forms, as palm oil is among the most widely used plant products in the world. Its uses range from cooking and baking to cosmetics and biofuels, and this ubiquity has driven continued growth in palm oil production and emissions. In 2016, palm oil expenditures were \$45 billion and 32% of total vegetable oil expenditures – more than any other vegetable oil. Substitutes include coconut, olive, rapeseed, soybean, and sunflower oils, but versatility in use and a low price point have helped palm oil maintain its market share.⁴ Firms trade palm oil in competitive global commodity markets, with the largest firm accounting for only 2% of global consumption (WWF 2016). At the country level, the EU, China, and India account for 33% of global consumption and 48% of imports, while Indonesia and Malaysia consume 23% of the world’s palm oil domestically (table 1).

Significant palm emissions motivate regulation, but domestic regulation faces challenges. Palm oil profits limit incentives to pass regulation, and weak enforcement hampers regulation that does pass. In 2010, Norway pledged \$1 billion to Indonesia in cash incentives for domestic forest regulation, prompting the Indonesian government to issue a moratorium on new forest concessions in 2011. But the moratorium had

and Dinerstein 2020, Chatterjee 2023, Rubens 2023, Zavala 2024).

⁴ For the EU, biofuels have driven an important part of palm oil demand. I abstract from substitution between palm oil and fossil fuels because of EU biofuel targets. For example, 14% of fuel for transportation must be renewable by 2030. Where binding, these targets prevent increased fossil fuel use and thus encourage substitution from palm oil to other vegetable oils.

Table 1: Production, consumption, and trade

	Production	Exports	Consumption	Imports
	%	%	%	%
Indonesia	44	41	15	0
Malaysia	40	48	8	3
European Union	0	0	12	18
China	0	0	10	15
India	0	0	11	15
Rest of world	16	10	44	49

Each column sums to 100% and covers 1988 to 2016. I pool palm oil and palm kernel oil by volume.

little effect, failing to curb deforestation within existing concessions or otherwise, including in protected areas (Busch et al. 2015). Some policies even promote palm oil production rather than restricting it: for transportation, Indonesia and Malaysia mandate that fossil fuels be blended with palm-based biofuels at rates of 30% and 20%, respectively (USDA 2019a, 2019b).

Consequently, European policymakers have discussed intervening with regulation. The EU is set to eliminate green subsidies for palm-based biofuels, cap production, and achieve a complete phase-out by 2030. Palm-based biofuels face the further loss of green tax incentives in France and an outright ban in Norway. French parliament debated a “Nutella tax” in 2016, highlighting the copious use of palm oil in Nutella and other food products. Each policy uses European purchasing power to target emissions abroad. This paper considers the impacts of such policy.


3 Data

I construct annual panel data on palm oil prices, consumption, and production from 1988 to 2016. Appendix A details data sources and construction.

3.1 Demand

I measure annual prices and consumption of vegetable oils. The data span the study period from 1988 to 2016 and cover all major vegetable oils: palm, palm kernel, coconut, olive, rapeseed, soybean, and sunflower. World price data come from the

Figure 2: Prices and quantities



World prices are in nominal USD per ton, and world quantities are in megatons. Each aggregates over palm oil and palm kernel oil.

International Monetary Fund. Palm oil prices derive from forward contract prices at Bursa Malaysia Derivatives Berhad, which is the primary global exchange market for palm oil futures. I use consumer price index data from the World Bank to adjust for inflation and denominate prices in year-2000 dollars. Consumption data by country come from the USDA Foreign Agricultural Service. I compute total vegetable oil expenditures from these prices and quantities.

I aggregate along two margins. First, I aggregate countries into four consumer markets: the EU, China and India, Indonesia and Malaysia, and the rest of the world. For each consumer market, I compute total quantities and expenditure-weighted average inflation.⁵ Second, I aggregate individual vegetable oils into two product groups: palm oils and other oils. Palm oils include palm oil and palm kernel oil, while other oils include coconut, olive, rapeseed, soybean, and sunflower oils. For each product group, I compute total quantities and expenditure-weighted average prices.⁶ From here, I will use “palm oil” in reference to the palm oils product group. Figure 2 shows that palm oil prices have risen over time despite a seven-fold increase in quantities traded. Concurrent growth in prices and quantities indicates an outward shift of the aggregate demand curve, and indeed palm oil was adopted widely for use in food products, consumer goods, and biofuels during this period.

⁵ For inflation, I aggregate over the countries in each market. I average over the consumer price index data, weighting by each country’s household final consumption expenditures.

⁶ For prices, I aggregate over individual oils o within each product group. I use Stone price index $\ln p_t = \sum_o \omega_{ot} \ln p_{ot}$ for years t , world expenditure shares ω_{ot} , and world prices p_{ot} .

3.2 Supply

I measure Indonesian and Malaysian palm oil production by site and year. I define sites as groupings of plantations and mills, and I treat sites as firms. Sites choose to invest in plantations and mills, subject to state variables that affect profits.

Choices

I capture plantations and mills with satellite-based measures. The study area is Sumatra and Kalimantan of Indonesia and all of Malaysia. For plantations, [Xu et al. \(2020\)](#) use PALSAR and MODIS satellite data to map palm oil plantations in the study area at 1 km resolution from 2001 to 2016. I extend their measure back to 1988 with data on tree cover loss as a proxy for plantation development. I obtain these data from [Song et al. \(2018\)](#), who construct tree cover loss from 1988 to 2016 with Landsat and MODIS satellite data. I estimate the relationship between tree cover loss and plantation development in the overlapping period from 2001 to 2016, and I find that tree cover loss is strongly predictive of plantation development. I then apply the estimated relationship to extend the plantation development data into the non-overlapping period from 1988 to 2000. For mills, data from the World Resources Institute and the Center for International Forestry Research record palm oil mill locations for all of Indonesia and Malaysia in 2018. With historical satellite data from Google Earth, I confirm each location and identify 1,526 mills. I drop the 29 mills that lie outside of the study area.

I use the plantation and mill data to divide the study area into independent plots, which I call “sites.” Active sites have one mill with nearby plantations. Potential sites have no mills or plantations, but they represent potential entrants. In the data, the provinces with the highest density of palm oil production contain one mill per 535 km² of land area in 2016. I treat this ratio as a target density. For each province, I obtain site boundaries by k -means clustering on geographic coordinates, where the number of clusters k is given by land area divided by the target density. I impose that no cluster contain more than one observed mill and that observed plantations be assigned to clusters with an observed mill. I obtain 2,050 contiguous sites.

I overlay plantations, mills, and site boundaries to construct a panel by site and year. I use the plantation data to identify the timing of mill construction by assuming

Table 2: Site statistics

Variable	Mean	SD	Min	Max	N
Mill	0.72	0.45	0	1	2,050
Plantations, ha	9,694	12,047	0	165,986	2,050
Yields, t/ha	3.37	0.57	2.01	5.22	2,050
Road distance, km	48	50	0	267	2,050
Port distance, km	191	100	7	468	2,050
Urban distance, km	125	90	0	417	2,050
Tree biomass CO ₂ , t/ha	386	157	26	753	2,050
Peat deposit CO ₂ , t/ha	1,240	2,079	0	16,217	2,050

Each observation is an Indonesian or Malaysian site in year 2016. Plantations are in hectares, and palm oil yields are in tons per hectare per year. Distances in kilometers are to major roads, major ports, and administrative cities (Indonesia) or federal territories (Malaysia). Carbon stock densities, in tons per hectare, include aboveground tree biomass and belowground peat deposits.

that sites build mills alongside their first plantations. I drop the 0.3% of plantations without mills and the 1% of mills without plantations observed by 2016. I assume zero exit for both plantations and mills, and indeed exit is limited when observable. [Xu et al. \(2020\)](#) measure cumulative plantation exit of only 4.6% between 2007 and 2016, perhaps because oil palm is a perennial crop with steady profits once planted. I compare the cleaned data to government statistics, and I find that the data align well. Appendix A shows that I match the large growth in plantation area over time, as well as the distribution of mills across space.

The top rows of table 2 summarize site choices by 2016. Of 2,050 total sites, 72% have an observed mill. The average plantation is large, at nearly 10,000 hectares in area. Over time, I observe plantation acreage increasing substantially from 2.4 Mha in 1988 to 19.9 Mha in 2016, relative to a study area of 134 Mha. That is, 15% of total land is developed into palm oil plantations. Roughly half of the study area is too mountainous for agriculture, and so the proportion of arable land developed is even higher. At the site level, 2.6% of sites without a mill choose to construct a new mill in an average year. Sites with a mill choose to develop an average of 464 ha of new plantation each year. Consistent with interior solutions, which I will assume for estimation, new plantation development is non-zero for 99.6% of site-year observations and never exceeds the available land area.

States

Palm oil profits depend on prices and yields. I use the same palm oil prices described previously for demand, and I compute palm oil yields over time with an agronomic model and government statistics. The PALMSIM model of [Hoffmann et al. \(2014\)](#) predicts potential yields under optimal growing conditions as a function of exogenous climate conditions.⁷ I run the model with WorldClim data on solar radiation and precipitation to obtain potential yields by site. Government statistics from the Indonesian Ministry of Agriculture and the Malaysian Palm Oil Board record actual yields by province-year. I calculate yield gaps as one minus the ratio of actual to potential yields. I assume that sites within a province-year share a common yield gap, which I multiply by potential yields to obtain actual yields by site-year. Yields vary across space, reflecting climate conditions that can differ by site. Yields also vary over time, reflecting technological progress that can differ by province.

I also consider cross-sectional variation in covariates that potentially affect production costs. I calculate distance to markets as the sum of Euclidean distances to the nearest major road, port, and urban area. These distances proxy for transport costs. I compute carbon stocks from geospatial data on tree biomass and peat deposits ([Zarin et al. 2016](#), [Gumbricht et al. 2017](#)), which allow me to link plantation development to emissions. Administrative boundaries delineate the four major producing regions: Sumatra, Kalimantan, Peninsular Malaysia, and East Malaysia.

The bottom rows of table 2 summarize the state variables. Yields are high at 3.37 tons per hectare per year for the average site. Average annual revenues are therefore \$1,840 per hectare at an average price of \$546 per ton, among years plotted in figure 2. Carbon externalities are also large. The average site stores 1,626 tons of CO₂ per hectare, with 386 tons from tree biomass and a much larger 1,240 tons from peat deposits. Even with recurring revenue, carbon damages outweigh revenues for any social cost of carbon that exceeds \$12 per ton.⁸ Carbon damages are most severe for peat-rich sites, where carbon stores can exceed 10,000 tons per hectare.

⁷ The model includes plant growth and radiation modules, which simulate fresh fruit bunch production as the outcome of frond, trunk, root, and flower growth. [Hoffmann et al. \(2014\)](#) validate the model with observed yields under optimal conditions from 13 sites that span my study area.

⁸ For discount factor $\beta = 0.9$, annual revenue of \$1,840 has a net present value of \$18,400, ignoring production costs. For $SCC = \$12$, carbon stores of 1,626 tons imply \$19,512 in carbon damages.

Shifters

For supply shifters, I consider crop yields by vegetable oil. The direct measure for palm oil combines agronomic modeling and government statistics, but it is difficult to replicate this approach for every vegetable oil. Thus, I instead construct an indirect measure that isolates weather shocks to oil crop production. I collect daily rainfall and temperature data at 0.25° resolution from the Global Meteorological Forcing Dataset, which I combine with crop-specific optimal growing conditions from the FAO Ecocrop Database, as well as province-specific production from the USDA Foreign Agricultural Service. For each year, crop, and province, I compute weather shocks as total absolute deviations from optimal levels during the growing season. Then, for each year and vegetable oil, I aggregate over crops and provinces while weighting by production. These weather shocks proxy for yields.

4 Demand

I model consumers that demand palm oil and other vegetable oils. I use iterative methods for estimation, which simplifies to linear regression with instruments. I describe demand estimates by consumer market, and I discuss the implications for coordination.

4.1 Model

Consumers choose between palm and other vegetable oils. I model demand in product space with an almost ideal demand system, which allows me to capture cross-product substitution patterns flexibly (Deaton and Muellbauer 1980).⁹ For markets k , years t , and vegetable oils $o \in \{1, 2\} = \{\text{palm}, \text{other}\}$, demand is given by

$$\omega_{okt} = \sum_{\hat{o}} \alpha_{o\hat{o}k} \ln p_{\hat{o}t} + \gamma_{ok}^0 + \gamma_{ok}^1 t + \delta_{ok} \ln \left(\frac{X_{kt}}{P_{kt}} \right) + \varepsilon_{okt}, \quad (1)$$

$$\ln P_{kt} = \frac{1}{2} \sum_o \sum_{\hat{o}} \alpha_{o\hat{o}k} \ln p_{ot} \ln p_{\hat{o}t} + \sum_o (\gamma_{ok}^0 + \gamma_{ok}^1 t) \ln p_{ot}. \quad (2)$$

⁹ The characteristic-space approach of Berry et al. (1995) restricts patterns of substitution to operate through product characteristics. It also requires specifying the product characteristics that consumers value. But unlike the product-space approach, it is tractable with many products.

In equation 1, expenditure shares ω_{okt} depend on world prices p_{ot} for both palm and other oils, fixed effects γ_{ok}^0 and time trends γ_{ok}^1 that capture unobserved heterogeneity by market, total vegetable oil expenditures X_{kt} , price index P_{kt} , and shocks ε_{okt} . Unobservables accommodate existing tariffs, which in any case are limited.¹⁰ Own- and cross-price coefficients $\alpha_{o\hat{o}k}$ allow for flexible patterns of substitution. In equation 2, translog price index P_{kt} aggregates over individual oil prices p_{ot} .¹¹ By definition of expenditure shares $\omega_{okt} = q_{okt}p_{ot}/X_{kt}$, quantities demanded are

$$q_{okt}^D = \frac{\omega_{okt}X_{kt}}{p_{ot}}. \quad (3)$$

4.2 Estimation

I estimate the model by iterated linear least squares (Blundell and Robin 1999). The challenge is that equations 1 and 2 call for nonlinear estimation, as demand parameters enter nonlinearly through price index P_{kt} . But for fixed price index values P_{kt}^0 , equation 1 is a linear regression equation.

$$\omega_{okt} = \sum_{\hat{o}} \alpha_{o\hat{o}k} \ln p_{\hat{o}t} + \gamma_{ok}^0 + \gamma_{ok}^1 t + \delta_{ok} \ln \left(\frac{X_{kt}}{P_{kt}^0} \right) + \varepsilon_{okt} \quad (4)$$

First, I compute initial price index values $\ln P_{kt}^0 = \ln X_{kt} - \ln Q_{kt}$ from data on total expenditures X_{kt} and quantities $Q_{kt} = \sum_o q_{okt}$. Second, I estimate equation 4 taking these price index values as given. I do so on palm oil expenditure shares alone, noting that other oil shares are collinear because shares sum to one, and I impose the standard adding-up, homogeneity, and symmetry restrictions.¹² Regression coefficients identify demand parameters. Third, I use estimated demand parameters to compute price index values by equation 2. Fourth, I repeat from step two until convergence.

I estimate equation 4 for each market separately. Prices p_{ot} are endogenous, as

¹⁰ EU tariffs are only 3.8% for crude palm oil (WTO 2023a). Unobservables also absorb physical trade costs, including shipping costs for palm oil from Indonesia and Malaysia.

¹¹ Price index P_{kt} depends on market-specific parameters, and so it varies by market even though world prices p_{ot} do not.

¹² With more products, estimation can apply seemingly unrelated regression to a system of equations. Under adding-up, $\sum_o \alpha_{o\hat{o}k} = 0$ for all \hat{o} , $\sum_o \gamma_{ok}^0 = 1$, $\sum_o \gamma_{ok}^1 = 0$, and $\sum_o \delta_{ok} = 0$. It is automatically satisfied if $\sum_o \omega_{okt} = 1$. Under homogeneity, $\sum_{\hat{o}} \alpha_{o\hat{o}k} = 0$ for all o , such that demand is unaffected by scaling prices and expenditures. Under symmetry, $\alpha_{o\hat{o}k} = \alpha_{\hat{o}ok}$ for all o, \hat{o} . With two products, imposing homogeneity imposes symmetry and vice versa.

unobserved shocks ε_{okt} affect prices by increasing demand. Thus, I instrument for prices with crop yields as a supply shifter. I use weather shocks to vegetable oil production as a measure of yields, as I can construct these shocks for every vegetable oil.¹³ Greater shocks correspond to lower yields, which lower supply and increase prices. The exclusion restriction is that these shocks affect vegetable oil demand only through their impact on prices, and not through their impact on income or expenditures more broadly. To this end, I isolate the weather shocks that are most relevant to vegetable oil production: deviations from optimal weather conditions for oil crops, specifically in the provinces and states that produce these crops, and only in the months of the growing season. Appendix B tests for and rules out income and expenditure effects. Moreover, unobserved shocks ε_{okt} may be correlated over time, and so I account for serial correlation with Newey-West standard errors.

With the estimated parameters, I can compute demand elasticities. I raise palm oil prices by 1% in each year, holding all else constant. I then compute quantities demanded with equations 1, 2, and 3. I report percentage changes in total consumption over the study period, with standard errors given by the delta method.

4.3 Estimates

Table 3 presents demand parameter estimates. Interpretation is indirect because equation 1 is specified in expenditure shares and not in quantities. For palm oil, own-price coefficients α suggest that expenditure shares do not react strongly to prices. When palm oil prices rise by 1%, EU palm oil expenditure shares rise by only 0.041 percentage points. This modest effect on expenditure shares implies that quantities fall as prices rise. Intercepts γ^0 and time trends γ^1 capture observed differences in palm oil consumption across markets. Indonesia and Malaysia have a large, positive intercept of 1.060, which rationalizes observed expenditure shares for palm oil that exceed 90%. Indonesia and Malaysia are major palm oil producers, and these high expenditure shares for palm oil are consistent with home bias. All markets have positive time trends, rationalizing rising consumption in spite of rising prices, as in figure 2. Expenditure coefficients δ govern how consumption responds as expenditures rise. Other importers shift toward higher palm oil shares, while the remaining markets

¹³ If I were to restrict demand estimation to palm oil alone, then I could directly apply the detailed yields that I obtain for palm oil. These yields enter the supply model as y_{it} .

Table 3: Demand parameters

θ^d	European Union		China/India		Other importers		Indonesia/Malaysia	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
α	0.041	(0.030)	0.045	(0.043)	-0.001	(0.038)	0.047	(0.032)
γ^0	0.082	(0.283)	-0.269	(0.294)	-0.444*	(0.244)	1.060***	(0.188)
γ^1	0.004***	(0.001)	0.002	(0.002)	0.004***	(0.001)	0.012***	(0.002)
δ	0.008	(0.028)	0.048	(0.029)	0.061***	(0.022)	-0.023	(0.021)

Each pair of columns shows parameters for a consumer market: α_{11k} , γ_{1k}^0 , γ_{1k}^1 , and δ_{1k} . Parameters for other oils follow from the adding-up, homogeneity, and symmetry restrictions. Demand estimation draws on annual data covering coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

respond more neutrally.

Table 4 presents demand elasticities for palm oil by consumer market. I estimate elasticities that are roughly similar across markets and all less than one, such that demand is relatively inelastic. Inelastic demand reduces leakage concerns and thus the losses from a failure to coordinate. However, leakage concerns remain as long as demand is less than perfectly inelastic. Without price instruments, I obtain estimates with upward bias, particularly for the EU, China, and India. This upward bias arises because prices are positively correlated with unobserved demand shocks. Appendix B shows the strong first stage for weather shocks as instruments and presents demand elasticities for other oils, which I find are similar in magnitude to demand elasticities for palm oil.

I model demand as static, which simplifies estimation at the cost of potential bias. The bias can go in either direction. If switching among vegetable oils is a gradual process that involves new recipes and suppliers, then contemporaneous price responses will be attenuated. I will underestimate demand elasticities and understate leakage concerns. If consumers stockpile to take advantage of temporary price drops, then contemporaneous price responses will be exaggerated. I will overestimate demand elasticities and overstate leakage concerns. In both cases, the underlying issue is that estimation relies on annual variation in prices, but consumers may not be responding to short-run prices. Appendix B evaluates these concerns with price lags and leads, as well as rolling variation over decadal horizons. I obtain similar estimates across specifications.

Table 4: Demand elasticities

	IV		OLS	
	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.209*	(0.109)
China/India	-0.692***	(0.168)	-0.271	(0.763)
Other importers	-0.876***	(0.128)	-0.521	(0.646)
Indonesia/Malaysia	-0.925***	(0.046)	-0.905***	(0.151)

Each pair of columns shows own-price elasticities for palm oil by consumer market. I report elasticities of total consumption with respect to a 1% increase in prices from 1988 to 2016. Demand estimation draws on annual data covering coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. IV estimation instruments for prices with weather shocks to vegetable oil production, while OLS estimation does not. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5 Supply

I model producers that supply palm oil by investing in mills and plantations. World prices clear markets in equilibrium. I use Euler methods for estimation, which simplifies to linear regression with instruments. I describe supply estimates over the short and long run, and I discuss the implications for commitment.

5.1 Model

Sites produce palm oil with mills and plantations. These sites are small, independent, and forward-looking with rational expectations. Long-lived owners manage sites without exit or scrappage. I model dynamics explicitly, as I seek to connect short-run responses in the data to long-run responses in counterfactuals.

Choices and states

Sites i make choices $\{m_{it}, n_{it}\}$. In each year t , sites without mills choose whether to construct a mill, and then sites with a mill choose how much land to develop into plantations. Mill construction m_{it} is a binary, extensive-margin choice to enter into production or not, while plantation development n_{it} is a continuous, intensive-margin choice over the scale of production.

Observed states $\{M_{it}, N_{it}\}$ track choices $\{m_{it}, n_{it}\}$. Each is within the control of

individual sites. Mill stock M_{it} and plantation acreage N_{it} follow laws of motion

$$M_{it+1} = M_{it} + m_{it}, \quad N_{it+3} = N_{it+2} + n_{it}.$$

Plantation acreage N_{it} tracks mature, fruit-bearing plantations. Newly planted crops require three years to bear fruit, and so plantation acreage grows with a three-year lag. Sites face three constraints. First, each site supports no more than one mill, such that $M_{it}, m_{it} \in \{0, 1\}$. Second, sites must develop plantations within their own lands, such that $N_{it} \in [0, L_i]$ and $n_{it} \in [0, L_i - N_{it+2}]$ for land area L_i . Third, plantations cannot operate without mills, such that $N_{it} = 0$ if $M_{it} = 0$.

Observed states $\{p_t, y_{it}, x_i, g_i\}$ affect choices $\{m_{it}, n_{it}\}$. Sites take each as given. Individual sites are price takers for world palm oil prices p_t , where $p_t = p_{1t}$ of equations 1 and 2. Yields y_{it} depend on climatic conditions that sites cannot change. These prices and yields determine revenues. Cost factors x_i include distance to markets and carbon stocks. Distance to markets sums over distances to major roads, ports, and urban areas, none of which target individual sites.¹⁴ I will estimate the extent to which this distance raises transport costs. Carbon stocks are predetermined and increase emissions, which sites may or may not internalize. Region g_i encodes the four regions of study – Sumatra, Kalimantan, Peninsular Malaysia, and East Malaysia – to allow for regional unobserved heterogeneity. Regional boundaries are fixed.

Unobserved states $\{\bar{v}_{it}, \bar{\varepsilon}_{it}, \varepsilon_{it}\}$ also affect choices $\{m_{it}, n_{it}\}$. Mill shocks \bar{v}_{it} are logit-distributed and IID. Unobserved mill and plantation costs $\{\bar{\varepsilon}_{it}, \varepsilon_{it}\}$ are more flexible: they are uncorrelated with each other, but individually can be correlated across sites and over time. I collect states with the notation

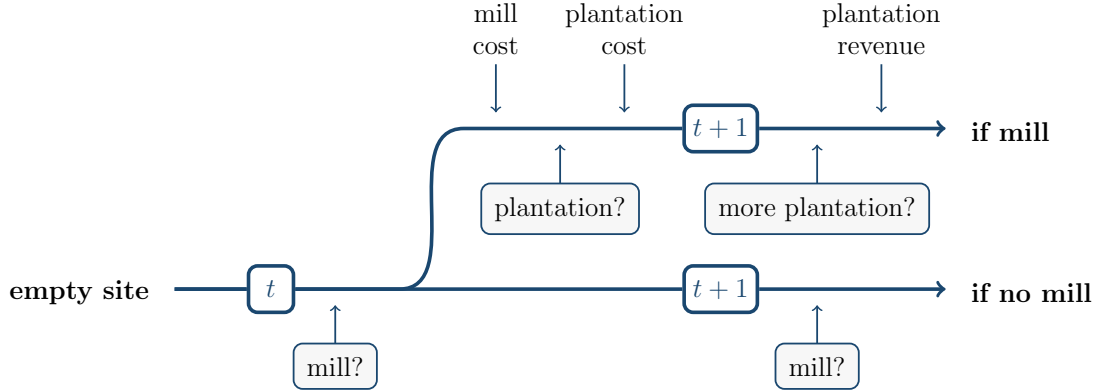
$$s_{it} = \{p_t, y_{it}, x_i, g_i, \bar{\varepsilon}_{it}, \varepsilon_{it}\}.$$

Timing and production

Each year, sites realize state s_{it} and then proceed in two stages. Figure 3 illustrates. In the first stage, sites construct mills. Sites with an existing mill do not face a choice, as sites can support only one mill. If $M_{it} = 1$, then $m_{it} = 0$. Sites otherwise

¹⁴ Major roads exclude small roads built for plantations, major ports predate plantations, and major urban areas exclude palm oil settlements.

Figure 3: Supply model timeline



An empty site makes a binary choice to construct a mill or not. If not, then the site faces the same choice in the next period. If so, then the site makes a continuous choice over how much land to develop into plantations. The site can then expand its plantation in future years.

face a choice. If $M_{it} = 0$, then they realize logit shock \bar{v}_{it} and choose mill construction $m_{it} \in \{0, 1\}$. For sites i and years t , the ex-ante value function is

$$\bar{V}(s_{it}) = \mathbb{E} \left[\max_{m_{it}} \{ \beta \bar{V}(s_{it+1}), -\bar{c}(s_{it}) + V(0, s_{it}) - \bar{v}_{it} \} \mid s_{it} \right]. \quad (5)$$

Sites that choose $m_{it} = 0$ receive next-year value $\bar{V}(s_{it+1})$. They do not construct a mill, and they face the same choice the next year. Sites that choose $m_{it} = 1$ incur mill cost $\bar{c}(s_{it})$ for plantation value $V(0, s_{it})$, starting from $N_{it} = 0$. That is, they construct a mill and begin to develop plantations, which eventually generate revenues. The outside option is to never construct a mill, with utility normalized to zero.

In the second stage, sites develop plantations. Sites without an existing or new mill do not face a choice, as plantations require mills. If $M_{it} + m_{it} = 0$, then $n_{it} = 0$. Sites otherwise face a choice. If $M_{it} + m_{it} = 1$, then they choose plantation development n_{it} . I assume interior solutions $n_{it} \in (0, L_i - N_{it+2})$ for land area L_i . For sites i and years t , the ex-ante value function is

$$V(N_{it}, s_{it}) = \mathbb{E} \left[\max_{n_{it}} \{ r(N_{it}, s_{it}) - c(n_{it}, s_{it}) + \beta V(N_{it+1}, s_{it+1}) \} \mid N_{it}, s_{it} \right]. \quad (6)$$

Mature plantations N_{it} generate revenues $r(N_{it}, s_{it})$, while plantation development n_{it} incurs costs $c(n_{it}, s_{it})$. Next-year value $V(N_{it+1}, s_{it+1})$ captures future profits, including the option value of future plantation development.

I specify revenues and costs as follows. For plantations, linear revenues and convex costs ensure unique optima.

$$r(N_{it}, s_{it}) = \alpha p_t y_{it} N_{it}, \quad c(n_{it}, s_{it}) = \left(\gamma_g^0 + \gamma_g^1 t + x_i \delta + \varepsilon_{it} + \frac{1}{2} \psi n_{it} \right) n_{it}$$

Revenues reflect prices p_t , yields y_{it} , and plantation acreage N_{it} . Parameter α governs how strongly development responds to higher revenues.¹⁵ Costs depend on fixed effects γ_g^0 and time trends γ_g^1 that capture unobserved heterogeneity by region, cost factors x_i that capture observed heterogeneity by site, and unobserved costs ε_{it} by site.¹⁶ Unobservables accommodate existing regulation, which in any case is limited and weakly enforced (Busch et al. 2015). Quadratic costs ψ encourage plantation development over time, capturing credit constraints and local factor market congestion. For mills, there are no direct revenues. Costs are

$$\bar{c}(s_{it}) = \bar{\gamma}_g^0 + \bar{\gamma}_g^1 t + x_i \bar{\delta} + \bar{\varepsilon}_{it}.$$

They again depend on fixed effects $\bar{\gamma}_g^0$, time trends $\bar{\gamma}_g^1$, observed costs x_i , and unobserved costs $\bar{\varepsilon}_{it}$ that capture regional and site heterogeneity. I interpret costs as upfront costs.¹⁷

Production depends on yields y_{it} and plantation acreage N_{it} . Quantities supplied are

$$q_{it}^S = y_{it} N_{it}. \tag{7}$$

Because of dynamics, quantities in one year depend on states in every year. By the laws of motion, current acreage N_{it} is a stock that depends on all past choices. And by equations 5 and 6, these past choices are forward-looking and in turn depend on states in every future year. Thus, to solve the model, I will need to specify the full expected path of states over time. Appendix C details this calculation.

¹⁵ It is equivalent to set $\alpha = 1$, treat revenues as numeraire, and estimate a logit scale parameter.

¹⁶ To identify unobserved heterogeneity by site, I must observe multiple choices per site. But multiple plantation choices are observed only for early sites, and multiple mill choices are ruled out because sites support only one mill each. I instead estimate regional effects, effectively pooling in the cross section rather than over time.

¹⁷ In practice, costs combine upfront costs, flow costs, and scrap values. Limited exit in the data suggests that upfront costs are relatively large. If upfront costs were small relative to flow costs or scrap values, then I would instead observe entry followed by exit at higher rates. Separating flow costs and scrap values is difficult without additional data.

Equilibrium

For terminal year T and vegetable oils $o \in \{1, 2\} = \{\text{palm}, \text{other}\}$, a dynamic competitive equilibrium is defined by prices $p^* = \{p_{11}^*, p_{21}^*, \dots, p_{1T}^*, p_{2T}^*\}$ such that:

1. World demand for palm and other oils is given by equations 1, 2, and 3. Demand depends on contemporaneous prices $\{p_{1t}, p_{2t}\}$ and total expenditures X_{kt} . Summing over markets k , world demand is $D_{ot}(p_{1t}, p_{2t}) = \sum_k q_{okt}^D(p_{1t}, p_{2t}; X_{kt})$.
2. World supply of palm oil is given by equations 5, 6, and 7. Supply depends on all prices $p_1^T = \{p_{11}, \dots, p_{1T}\}$ and yields $y_i^T = \{y_{i1}, \dots, y_{iT}\}$. Sites are price takers individually, but they affect world prices collectively. Summing over sites i , world supply is $S_{1t}(p_1^T) = \sum_i q_{it}^S(p_1^T; y_i^T)$.
3. World supply of other oils is given inelastically by quantities $\{S_{21}, \dots, S_{2T}\}$.¹⁸
4. World markets clear. For world demand and supply defined above,

$$D_{1t}(p_{1t}, p_{2t}) = S_{1t}(p_1^T), \quad D_{2t}(p_{1t}, p_{2t}) = S_{2t} \quad \forall t. \quad (8)$$

5.2 Estimation

I use Euler methods to estimate the model without the need to compute continuation values (Hall 1978, Scott 2013). I obtain two linear regression equations that I stack and estimate jointly. Estimation is straightforward and computationally light.

$$n_{it} - \beta n_{it+1} = \frac{\alpha \beta^3}{\psi} p_{t+3} y_{it+3} + \frac{\beta}{\psi} \gamma_g^1 - \frac{1 - \beta}{\psi} (\gamma_g^0 + \gamma_g^1 t + x_i \delta) + \mu_{it} + \eta_{it} \quad (9)$$

$$\ln \left(\frac{\pi_{it}}{1 - \pi_{it}} \right) - \beta \ln \pi_{it+1} = \frac{1}{2} \psi n_{it}^2 + \beta \bar{\gamma}_g^1 - (1 - \beta) (\bar{\gamma}_g^0 + \bar{\gamma}_g^1 t + x_i \bar{\delta}) + \bar{\mu}_{it} + \bar{\eta}_{it} \quad (10)$$

The dependent variables include plantation development $\{n_{it}, n_{it+1}\}$ and conditional choice probabilities $\{\pi_{it}, \pi_{it+1}\}$, which are the probabilities of mill construction. The residuals include structural errors $\{\mu_{it}, \bar{\mu}_{it}\}$ and expectational errors $\{\eta_{it}, \bar{\eta}_{it}\}$, where structural errors reflect unobserved costs $\{\varepsilon_{it}, \bar{\varepsilon}_{it}\}$.

$$\mu_{it} = -\frac{1}{\psi} \varepsilon_{it} + \frac{\beta}{\psi} \varepsilon_{it+1}, \quad \bar{\mu}_{it} = -\bar{\varepsilon}_{it} + \beta \bar{\varepsilon}_{it+1}$$

¹⁸ I can alternatively assume perfectly elastic supply of other oils and treat prices p_{2t} as fixed. It is more difficult to estimate a model of other oils alongside the present model of palm oil. An intermediate option is to calibrate the elasticity of supply of other oils.

Regression coefficients identify supply parameters. The coefficients of equation 9 identify parameters $\{\frac{\alpha}{\psi}, \frac{\gamma_g^0}{\psi}, \frac{\gamma_g^1}{\psi}, \frac{\delta}{\psi}\}$ if discount factor β is known. The discount factor is not identified, as is typical of dynamic discrete choice models (Magnac and Thesmar 2002), and so I set $\beta = 0.9$. The coefficients of equation 10 identify $\{\psi, \bar{\gamma}_g^0, \bar{\gamma}_g^1, \bar{\delta}\}$, thereby isolating ψ and giving $\{\alpha, \gamma_g^0, \gamma_g^1, \delta\}$ in levels. The main parameter of interest is revenue coefficient α , which captures the elasticity of development with respect to prices. I note that price and yield variation jointly identify α . That is, I benefit from granular spatial variation in yields, rather than relying solely on time-series variation in world prices. Intuitively, high-yield sites benefit more from high prices than low-yield sites, as revenues reflect both prices and yields. If supply is elastic, then high-yield sites develop more aggressively than low-yield sites when prices rise.¹⁹

I derive the regression equations from Euler equations, which compare investment in years t and $t + 1$. Appendix C presents derivations. On the intensive margin, I differentiate equation 6 with respect to plantation development n_{it} and n_{it+1} . Continuation values align and difference out by the envelope theorem. I obtain equation 9, which captures an intertemporal trade-off: earlier plantation development n_{it} brings added revenue $p_{t+3}y_{it+3}$ and avoids rising cost trends γ_g^1 , while later development n_{it+1} delays costs $(\gamma_g^0 + \gamma_g^1 t + x_i \delta)$ and discounts them. On the extensive margin, I difference equation 5 with respect to mill construction m_{it} and m_{it+1} . Continuation values align and difference out by finite dependence, which holds because mill construction and plantation development are terminal actions that lead to common future states and payoffs (Arcidiacono and Miller 2011). Whether sites invest in year t or $t + 1$, mills are operational, and plantations have matured by year $t + 4$. I obtain equation 10, which also captures a trade-off: earlier mill construction brings added plantation profits, as embodied by n_{it}^2 , while later mill construction delays costs. Conditional choice probabilities $\{\pi_{it}, \pi_{it+1}\}$ are the probabilities of earlier and later mill construction. Observed choices capture future payoffs and stand in for continuation values, echoing the typical intuition for conditional choice probability estimation.

I estimate equation 9 on the sample of sites with a new or existing mill ($M_{it} + m_{it} = 1$). It is these sites that face a plantation development decision. There are three problems. First, future revenue $p_{t+3}y_{it+3}$ may be correlated with structural error μ_{it} ,

¹⁹ This high- to low-yield comparison gives identification only in relative terms. But zero-yield sites offer a natural normalization, as they receive zero benefit from price increases.

which includes unobserved costs ε_{it} . These unobserved costs may contain aggregate shocks that affect future supply and thus future prices p_{t+3} . The structural error remains uncorrelated with observed states $\{y_{it}, x_i, g_i\}$, which are site fundamentals. Second, future revenue is correlated with expectational error η_{it} . This expectational error is the difference between unobserved expectations and observed realizations. It includes $\mathbb{E}[p_{t+3}y_{it+3}|s_{it}] - p_{t+3}y_{it+3}$, which is mechanically correlated with $p_{t+3}y_{it+3}$. Third, the structural error is autocorrelated. It is correlated over time because both μ_{it} and μ_{it+1} contain ε_{it+1} , and furthermore ε_{it+1} may be correlated across sites.

I address these problems by instrumenting and clustering. For the first problem, I instrument with demand shifters. In particular, I instrument for $p_{t+3}y_{it+3}$ with $Z_t y_{it}$, where Z_t includes total vegetable oil consumption and weather shocks to other vegetable oil production. Total consumption raises the category budget and thus demand for palm oil. I focus on total consumption outside of Indonesia and Malaysia, as Indonesian and Malaysian consumption may not be excluded from palm oil production. Weather shocks to other oils affect the supply of other oils and thus residual demand for palm oil. Other oils are produced outside of Indonesia and Malaysia, and so these weather shocks are foreign and arguably excluded from palm oil production. For the second problem, I use lagged instruments. Under rational expectations, sites condition on all information known at time t , such that $Z_t y_{it}$ is orthogonal to η_{it} . For the third problem, I cluster standard errors by district to accommodate autocorrelation, at least in some form.

I estimate equation 10 on the sample of sites without mills ($M_{it} = 0$). It is these sites that face a mill construction decision. First, I discuss the choice terms, which I must compute from data. I compute conditional choice probabilities π_{it} non-parametrically, smoothing spatially over observed choices with cubic splines in latitude, longitude, cost factors, and time. I also compute plantation development n_{it} . I must do so because I estimate equation 10 for sites without mills, but I observe development only for sites with a mill. I smooth over observed choices, assuming that unobserved mill and plantation costs are uncorrelated. Second, I discuss the error terms, which motivate clustering. Structural error $\bar{\mu}_{it}$ is uncorrelated with observed states x_i but is correlated over time, and so I cluster standard errors by district. The structural error is uncorrelated with n_{it} , again assuming that unobserved mill and plantation costs are uncorrelated. Expectational error $\bar{\eta}_{it}$ is uncorrelated with x_i and

Table 5: Supply parameters

		Intensive				Extensive			
		θ^s	Unit	Estimate	SE	$\bar{\theta}^s$	Unit	Estimate	SE
Revenue	Price	α	10^{-8}	3.018**	(1.479)				
Cost	Median	γ^0/α	\$1K	8.037***	(0.345)	$\bar{\gamma}^0/\alpha$	\$1M	90.43**	(39.88)
	Trend	γ^1/α	\$1K	-0.401***	(0.037)	$\bar{\gamma}^1/\alpha$	\$1M	1.341***	(0.376)
	Distance	δ/α	\$1K	0.001	(0.001)	$\bar{\delta}/\alpha$	\$1M	0.343**	(0.165)
	Carbon	δ/α	\$1K	-0.000**	(0.000)	$\bar{\delta}/\alpha$	\$1M	0.002	(0.001)
	Convexity	ψ/α	\$1	6.508***	(0.759)				

The revenue row is the price coefficient. The cost rows divide cost parameters by the price coefficient, such that magnitudes are interpretable as inflation-adjusted, year-2000 USD. Costs describe median costs for sites with observed construction, as well as annual cost trends across regions, costs of market distance and carbon stocks, and cost convexities. Market distance sums over distances to major roads, ports, and urban areas, while carbon stocks sum over above- and belowground carbon stocks. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

n_{it} by rational expectations.


With the estimated parameters, I can compute supply elasticities. I raise palm oil prices by 1% over shorter and longer periods within the study period, holding all else constant. Small price changes within the study period do not affect plantation development n_{it} or mill construction probabilities π_{it} at the end of the study period, and so I can read these values from data. I then compute quantities supplied directly from regression equations 9 and 10, as described in appendix C. I do so instead of solving the model and computing quantities from equations 5 and 6, which require specifying expectations beyond the study period. I report percentage changes in total production over the study period, with standard errors given by the delta method.

5.3 Estimates

Table 5 presents supply parameter estimates, which I compute from regression coefficients. A positive price coefficient gives an upward-sloping supply curve, and dividing parameters by this coefficient gives magnitudes in dollar terms. I estimate relatively high median costs of \$8,037 per hectare of plantation and \$90.43 million per mill, with an additional \$1,510 per hectare from cost convexity.²⁰ Accounting esti-

²⁰ Estimated convexity is \$6.51 per hectare (times $\frac{1}{2}$), multiplied by an average n_{it} of 464 hectares.

Figure 4: Supply elasticities



I plot elasticities of total production with respect to a 1% increase in prices from 1988 to the years shown on the x axis. That is, the small dot marks how total production from 1988 to 2016 responds to a 1% increase in prices from 1988 to 2001. The large dot marks how total production from 1988 to 2016 responds to a 1% increase in prices from 1988 to 2016. The top curve is computed from IV estimates, and the bottom curve from OLS estimates. I plot 95% confidence bands.

mates are smaller at \$7,000 and \$20 million, respectively (Fairhurst and McLaughlin 2009, Man and Baharum 2011). These accounting estimates include planting and operating costs but abstract from capital and land acquisition, which my estimates capture. My estimates also capture other lifetime costs, including replanting and capital replacement, as well as constraints to expansion, including urban boundaries, that I do not model explicitly. Plantation costs fall by a meaningful 5% of median costs per year, while mill costs rise by 1.5%. Appendix C presents regional costs.

I find that producers internalize their private transport costs, but not their emission externalities. Distance from markets increases costs on the extensive margin, given transport costs from mills to markets. These distance costs are large: an additional kilometer from a major road, port, or urban area increases mill costs by 0.38% of median costs. If an additional kilometer of remoteness increases road, port, and urban distances simultaneously, then mill costs increase by 1.14%. At the same time, distance to markets has no impact on the intensive margin. Once a mill has been constructed, plantation development proceeds unhindered. Carbon stocks also have no impact on production. If anything, carbon stocks decrease costs on the intensive margin, as forests and peat may proxy for a lack of competing land claims. However, on both margins, the effects of carbon stocks are small in magnitude.

Figure 4 presents supply elasticities for palm oil. I report elasticities of total

production over the study period. Dots mark the short- and long-run price changes that I will consider in counterfactuals. Price changes sustained from 1988 to 2001 give a short-run elasticity of 1.3, and those from 1988 to 2016 give a long-run elasticity of 2.9. Very short-run price changes from 1988 to 1990 have no effect because I take 1988 as the initial year, and production responds with a three-year lag based on the time between planting and bearing fruit. My long-run estimate is consistent with those from the Amazon, where others have estimated long-run price elasticities of 4.1 and 6.3 (Sant’Anna 2024, Araujo et al. 2024). Each far exceeds the Scott (2013) estimate of 0.3 for the US. Large long-run elasticities highlight the need to commit to long-run policy, as forward-looking sites consider revenues over time. Without instruments, I obtain estimates with downward bias. This downward bias arises because revenues are negatively correlated with the error terms in equation 9. Revenues $p_{t+3}y_{it+3}$ enter expectational error η_{it} negatively, and unobserved costs ε_{it} raise prices p_{t+3} and enter structural error μ_{it} negatively. Appendix C shows the strong first stage for demand shifters as instruments.

Appendix C also estimates a static version of the model and finds elasticities that are smaller in magnitude and negative. Static estimation regresses on current prices, which are noisy measures of future prices. This noise biases estimates toward zero. Furthermore, investment can slow in response to short-run price spikes if expectations are mean-reverting, such that high prices today prompt expectations of lower prices tomorrow. For robustness, appendix C presents additional specifications with disaggregated cost factors and alternative basis functions for smoothing. It also evaluates the potential selection bias from assuming that unobserved mill costs are uncorrelated with unobserved plantation costs. I obtain similar estimates across specifications.

Euler estimation has important advantages.²¹ I avoid the need to compute continuation values, which greatly simplifies computation. Because estimation reduces to linear regression, I can address endogeneity and autocorrelation concerns with standard tools. And although I need to assume rational expectations, I do not need to specify expectations more precisely. I do not need to assume perfect foresight, and I note that regional terms $\gamma = \{\gamma_g^0, \gamma_g^1, \bar{\gamma}_g^0, \bar{\gamma}_g^1\}$ accommodate common expectational bias. By comparison, the full-solution approach requires computing continuation val-

²¹ Other discrete Euler applications include Diamond et al. (2019), De Groote and Verboven (2019), Traiberman (2019), and Almagro and Domínguez-Iino (2024). Hsiao (2025) develops an alternative approach with similar advantages, appealing to price data in place of finite dependence.

ues in every iteration. It also requires explicitly specifying long-run expectations, which involves stronger assumptions than rational expectations.

At the same time, estimation relies on several assumptions. First, sites consider investing today or tomorrow. Weak property rights may encourage land grabbing and bias toward investing today, although regional terms γ help by absorbing some variation in property rights. Second, sites are independent and atomistic. Otherwise, finite dependence does not hold: if large sites delay investment, then competitors respond and alter the evolution of the economy, such that continuation values do not align. It helps that world production is unconcentrated, with the largest producer accounting for 4% and the largest ten for 21% (POA 2017). But I must rule out local spatial interaction, which makes estimation intractable. Third, the age of mills and plantations does not affect profits. Otherwise, delayed investment affects profits in all future years, and finite dependence again does not hold.

6 Counterfactuals


I solve the model for counterfactuals, which compare direct regulation with domestic policy to indirect regulation with trade policy. I quantify impacts on emissions, consumer and producer surplus, and government revenue. I close with general lessons.

6.1 Solving the model

I solve for equilibrium prices and quantities with conditions 8. I set discount factor $\beta = 0.9$, unobservables $\bar{\varepsilon}_{it} = \varepsilon_{it} = \varepsilon_{okt} = 0$, and expectational errors $\bar{\eta}_{it} = \eta_{it} = 0$, and I assume palm oil production in Indonesia and Malaysia is proportional to global production.²² In solving the model, I compute supply in levels with equations 5 and 6. I do so by specifying long-run expectations and computing continuation values until terminal year $T = 2050$. Beyond the study period, I assume linear

²² Indonesia and Malaysia account for 84% of palm oil production during the study period (table 1), which limits bias from not directly modeling production elsewhere. I treat Indonesia and Malaysia as representative producers, and I scale their production accordingly. I compute multiplicative adjustment factors Ω_t , such that $D_{1t} = S_{1t}\Omega_t$, and I apply these adjustments when solving equilibrium conditions 8. During the study period, I observe world demand D_{1t} and Indonesian and Malaysian supply S_{1t} . Beyond the study period, I apply $\Omega_{2016} = 1.1$ based on the last year of the study period. Alternatively, additive adjustment factors treat palm oil production elsewhere as fixed, with very similar results in terms of model fit.

Figure 5: Model fit



I plot equilibrium prices, quantities, and cumulative emissions for palm oil, comparing model-implied values to observed data from 1988 to 2016. Prices are in nominal USD per ton, quantities in megatons, and emissions in gigatons of CO_2 .

growth in palm oil yields and the supply of other oils at the rates observed during the study period, as well as annual inflation of 2%. I also assume annual growth in total vegetable oil expenditures X_{kt} at a rate of λ , such that $X_{kt+1} = (1 + \lambda)X_{kt}$. I interpret λ as expected growth in aggregate demand over time.

I recover demand growth λ by matching the data in levels. Having already estimated demand and supply parameters $\hat{\theta} = \{\hat{\theta}^d, \hat{\theta}^s\}$, I choose a candidate value for λ and solve for equilibrium prices. Intuitively, demand growth affects future prices, which in turn affect current entry and thus current prices. I repeat to find the candidate value that best fits the palm oil prices in the data. I obtain $\hat{\lambda} = 0.105$.²³ This procedure effectively inverts the model to recover an implied measure of long-run expectations. I did not need to work with long-run expectations when estimating parameters θ , as estimation relied instead on equations 9 and 10. These estimating equations difference out long-run expectations and avoid solving the model, but they match the data only in changes. I must specify expectations and solve the model to match the data in levels, as is needed for counterfactuals.

Figure 5 assesses model fit by comparing model-implied values to observed data during the study period. The model matches the data well, noting that prices are

²³ Other long-run expectational assumptions, such as linear growth in palm oil yields, will be confounded with this demand growth. But the goal is simply to match the data in levels, rather than to isolate the precise nature of long-run expectations.

directly targeted. Prices and quantities are linked by equilibrium conditions 8, and I match data on each in both levels and trends. Quantities and emissions are linked by plantation development choices. However, quantities depend on yields, while emissions depend instead on carbon stocks. These values need not align. I miss the particularly large emission episodes of the late 2000s, but I otherwise capture the trajectory of emissions over the study period. Counterfactuals will impose regulation and study changes in outcomes relative to this model-implied baseline.

6.2 Policy evaluation

I evaluate policy in the form of palm oil taxes on the supply side, the demand side, and both in combination. Domestic regulation imposes production taxes $\tau_{gt}^S > 0$, which can vary by producing region g_i . Import tariffs and export taxes impose consumption taxes $\tau_{kt}^D > 0$, which can vary by consumer market k . Carbon border adjustments combine import tariffs with credits for domestic regulation. For ad valorem taxes $\{\tau_{kt}^D, \tau_{gt}^S\}$, equilibrium conditions 8 for palm oil become

$$\sum_k q_{1kt}^D ((1 + \tau_{kt}^D) p_{1t}, p_{2t}) = \sum_i q_{1it}^S ((1 - \tau_{g1}^S) p_{11}, \dots, (1 - \tau_{gT}^S) p_{1T}) \quad \forall t.$$

I assess the value of coordination and commitment. I simulate coordination by applying taxes across regions and markets. Complete regulation across regions prevents supply-side leakage, by which production shifts toward unregulated regions. Complete regulation across markets prevents demand-side leakage, by which consumption shifts. I simulate commitment by applying taxes over time. Commitment resists the static incentive to set taxes to zero, given sunk investment and time to build. This static incentive arises because taxes today are costly, but they do not prevent emissions. That is, taxes today do not prevent existing development, which is sunk, or new development, which does not yet produce taxable output.²⁴

²⁴ New development responds only to taxes tomorrow – after time to build has elapsed. But new development today becomes sunk development tomorrow. Without commitment, taxes tomorrow are again set to zero. Only commitment to non-zero taxes tomorrow can prevent development today. I note that these taxes are output taxes. The regulator could instead tax land development itself by imposing an immediate fine, rather than taxing output over time. But the regulator must still commit to enforcing the fine, which will be large if it imposes the full cost of emissions. A large fine may prompt legal challenges and lobbying that complicate commitment to enforcement.

I quantify impacts on emissions, consumer and producer surplus, and government revenue. Emissions depend on carbon stock density, which I observe, and the extent of plantation development, which I model. Surplus and revenue depend on equilibrium prices and quantities, which I solve for. I focus on impacts within the study period from 1988 to 2016, as those beyond the study period depend more heavily on the expectational assumptions required for solving the model. Consumer surplus is the compensating variation needed to maintain baseline utility, producer surplus is revenue net of costs, and government revenue is the product of tax rates, prices, and quantities. Appendix D provides formal expressions for each. I report welfare effects for each market as the sum of consumer surplus, producer surplus, and government revenue. A global social planner evaluates regulation by asking whether the benefits of emission reductions exceed the costs for welfare across markets.

Several restrictions simplify computation. First, tax rates are announced at the outset and taken as given. I abstract from the dynamic game between policymakers and producers. Second, tax rates are constant during an initial commitment period, then lapse to zero afterwards – as is statically optimal. More complex paths are more computationally intensive to evaluate and more difficult to administer in practice. Third, I tax palm oil uniformly. Palm emissions are not uniform, but heterogeneous taxes would require monitoring production and tracking sales.²⁵ Fourth, plantation development releases carbon stocks fully. Trees must be cut to make space for plantations, and the peat layer must be cleared to access the underlying soil. Fifth, I focus on palm emissions. I ignore emissions from demand substitution to other oils or supply substitution to other deforesting activities. Appendix D argues that the resulting bias is limited: other oils involve limited or non-peat deforestation, and other deforesting activities are much less profitable than palm oil production.

6.3 Domestic regulation

Domestic regulation taxes production directly. Table 6 simulates production taxes of 50%. I consider coordinated taxes by Indonesia and Malaysia and unilateral taxes by either alone, as well as commitment to long-run taxes from 1988 to 2016 and

²⁵ Heterogeneous taxes also require commitment not to “greenwash” palm oil produced with sunk deforestation. Moreover, uniform taxes avoid reshuffling concerns. Taxing dirty palm oil alone pushes dirty palm oil to unregulated markets and clean palm oil to regulated markets. With sufficient unregulated demand, the result is pure reallocation and zero decrease in dirty production.

Table 6: Emissions (Gt), welfare (\$1B), and abatement costs (\$/t)

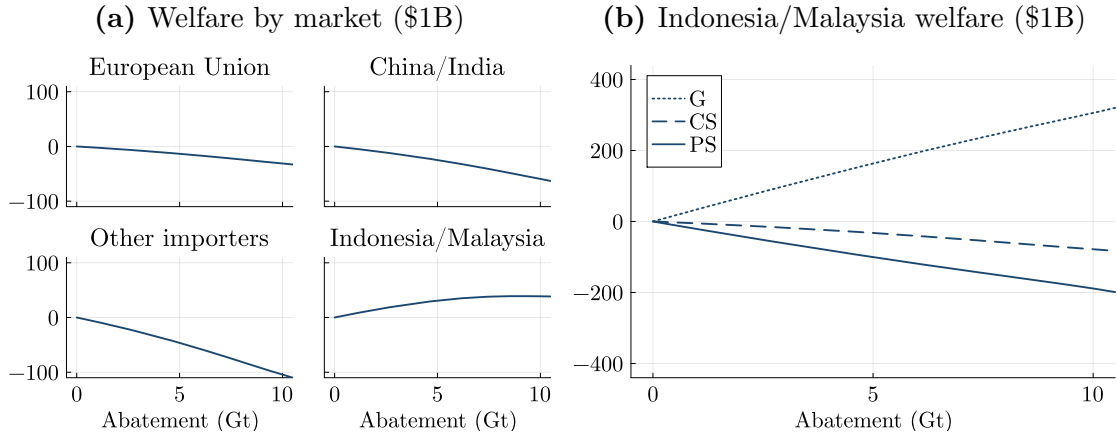
	τ	ΔE		ΔW^{EU}		ΔW^{CI}		ΔW^{OI}		ΔW^{IM}		$\Delta W/\Delta E$	
	%	2016	2001	2016	2001	2016	2001	2016	2001	2016	2001	2016	2001
Production taxes													
All exporters	50	-7.4	-0.8	-21	-5.5	-40	-9.0	-74	-22	38	6.7	13	38
Indonesia only	50	-4.7	-0.9	-6.8	-1.7	-12	-2.8	-22	-6.7	13	0.8	5.9	11
Malaysia only	50	-1.7	-0.1	-3.7	-1.2	-6.9	-1.9	-12	-4.7	7.3	0.9	9.2	60
Import tariffs													
All importers	100	-5.4	-0.5	8.6	2.1	9.4	1.2	13	5.9	-88	-38	10	60
EU, China, India	100	-2.1	-0.1	-1.6	-0.1	-6.9	-0.6	27	7.8	-32	-11	6.5	29
EU only	100	-0.7	-0.1	-9.7	-1.6	4.9	0.9	9.6	3.5	-13	-6.3	12	53

Each row is one counterfactual. I compute total changes in global emissions and market-specific welfare from 1988 to 2016, relative to business as usual, for the European Union (EU), China and India (CI), other importers (OI), and Indonesia and Malaysia (IM). Emissions are in gigatons of CO₂, and welfare is in billions of inflation-adjusted, year-2000 USD. Indonesia and Malaysia are exporters, and welfare includes consumer surplus, producer surplus, and government revenue. Other countries are importers, and welfare includes consumer surplus and government revenue. Abatement costs divide welfare costs, summed across markets, by emission reductions. The units are USD per ton of CO₂. Production taxes of 50% target some or all production, while import tariffs of 100% target some or all imports. Taxes are upheld from 1988 to 2016 or from 1988 to 2001.

short-run taxes from 1988 to 2001. I find that coordinated, long-run taxes reduce CO₂ emissions by 7.4 Gt from 1988 to 2016. The cost to global consumer and producer surplus, net of government revenue, is \$13 per ton of CO₂. That is, this policy improves global social welfare for any social cost of carbon that exceeds \$13 per ton. Unilateral and short-run taxes have smaller effects on emissions. For long-run action, emissions fall by 4.7 Gt when Indonesia acts alone and by 1.7 Gt when Malaysia acts alone. Unilateral Malaysian action is prone to leakage, as elastic Indonesian supply increases rapidly when Malaysian taxes drive up world prices. Unilateral Indonesian action is more effective, as it pushes production toward Malaysia, where higher yields increase efficiency. For short-run action, this compositional shift even leads to slightly larger emission reductions for unilateral Indonesian taxes relative to coordinated taxes. More efficient production in Malaysia also explains higher abatement costs for Malaysia-only regulation.

Domestic regulation reduces welfare for the EU, China, India, and other importers. Production taxes raise world prices and lower consumer surplus in these markets. Losses for other importers are twice as large as losses for China and India,

Figure 6: Production taxes



I simulate coordinated, long-run production taxes of increasing intensity by Indonesia and Malaysia. I plot effects on market-specific welfare and total abatement from 1988 to 2016. Abatement is palm emission reductions. Welfare is in billions of inflation-adjusted, year-2000 USD, and abatement is in gigatons of CO₂. Welfare for Indonesia and Malaysia includes consumer surplus, producer surplus, and government revenue. Welfare elsewhere includes consumer surplus.

which in turn are twice as large as losses for the EU. At the same time, production taxes increase welfare for Indonesia and Malaysia. These countries can manipulate their terms of trade, as their producer market power allows them to elevate world prices and raise tax revenue at the expense of foreign consumers. That is, production taxes simultaneously reduce emissions and raise welfare for Indonesia and Malaysia. If enforceable, domestic regulation is fiscally appealing even absent international pressures to abate.

Figure 6 plots welfare against abatement for production taxes of varying intensity, focusing on coordinated, long-run taxes. Figure 6a shows that increasing abatement also increases welfare losses for the EU, China, India, and other importers, where consumer surplus falls as world prices rise. Indonesia and Malaysia experience welfare gains as they exercise market power, with welfare maximized at 9.3 Gt of abatement. Abatement at this level corresponds to a 60% production tax. Figure 6b shows that welfare gains for Indonesia and Malaysia come from substantial government revenue, collected in part from foreign consumers. This revenue offsets consumer surplus losses from higher prices, as well as producer surplus losses that amount to hundreds of billions of dollars. Indonesian and Malaysian producers suffer losses that far exceed Norway's \$1 billion in cash compensation for forest regulation.

6.4 Import tariffs

Import tariffs tax traded consumption. Table 6 simulates import tariffs of 100%. I thus match the production taxes above: consumers pay twice the amount that producers receive when taxing demand at 100%, and the same holds when taxing supply at 50%. I consider coordinated tariffs by all importers, multilateral tariffs by an EU-China-India coalition, and unilateral tariffs by the EU alone, as well as commitment to long-run tariffs from 1988 to 2016, medium-run tariffs from 1988 to 2011, and short-run tariffs from 1988 to 2001. I note that precedents exist for import tariffs of 100%. For the EU, tariffs of 162% on sugar, 88% on beef, and 62% on milk aim to protect domestic agriculture, while tariffs of 257% on cigarettes act as excise duties on harmful goods.²⁶ More broadly, Hsiao et al. (2025) document that governments commonly use trade policy to intervene in agricultural markets.

I find that coordinated, long-run tariffs reduce CO₂ emissions by 5.4 Gt from 1988 to 2016, amounting to 0.19 Gt annually. This reduction is smaller than the 7.4 Gt achieved by production taxes, as import tariffs fail to regulate Indonesian and Malaysian consumers. But import tariffs still have large effects. Average annual palm emissions were 1.6 Gt from 1990 to 2016, relative to 5.0 Gt for the US, 5.0 Gt for China, 3.5 Gt for the EU, and 1.1 Gt for India (figure 1). Coordinated, long-run tariffs therefore reduce emissions by an amount equal to 17% of Indian emissions annually. However, unilateral and short-run import tariffs have smaller effects. For long-run action, emissions fall by 2.1 Gt under an EU-China-India coalition and by 0.7 Gt when the EU acts alone. Each is prone to leakage, as unregulated demand rises when import tariffs drive down world prices. For short-run action, temporary tariffs do little to dissuade palm oil production, as producers look toward high prices in post-tariff years. Emissions fall by no more than 0.5 Gt.

Import tariffs can increase welfare for the EU, China, India, and other importers.

²⁶ For sugar, beef, and milk, I compute ad valorem equivalents by combining non-ad valorem rates with primary commodity prices for 2020 (WTO 2023a, IMF 2023). I choose 2020 to capture ad valorem equivalents before the recent inflationary period. For cigarettes, EU legislation requires that “the overall excise duty on cigarettes shall represent at least 60% of the weighted average retail selling price of cigarettes released for consumption” since 2014 (OJEU 2011). The European Commission offers the following sample calculation: a pre-tax price of 0.70 EUR, an excise duty of 1.80 EUR, and a post-duty 20% VAT of 0.50 EUR together yield a retail price of 3.00 EUR. The excise duty is 60% of the retail price and 257% of the pre-tax price.


These countries can manipulate their terms of trade, as their consumer market power allows them to lower world prices and raise tax revenue at the expense of foreign producers. This market power is strongest when importers act together, as coordinated tariffs – both long- and short-run – raise welfare for all importers. Smaller tariff coalitions have less market power, and so coalition importers suffer welfare losses because government revenue does not offset the direct consumer surplus losses from import tariffs. But non-coalition importers still enjoy welfare gains because import tariffs lead to lower world prices.

At the same time, Indonesia and Malaysia suffer large welfare losses across import tariff scenarios. To this end, I consider compensating transfers in the spirit of payments for ecosystem services. First, these transfers promote equity. Palm oil fuels economic development in Indonesia and Malaysia, especially in poor, rural communities. In curbing emissions, these countries forgo local profits for global benefit. Second, these transfers help to navigate legal and diplomatic concerns. Indonesia and Malaysia have criticized EU trade policy for palm oil, arguing that it penalizes palm oil relative to the “like goods” of rapeseed and sunflower oils, which the EU produces domestically. Transfers act as compensation.

I evaluate import tariffs by imagining the EU as tariff coalition leader, and I calculate costs for the EU inclusive of the proposed transfers. I suppose transfers are to the point that all non-EU markets at least weakly prefer EU-led import tariffs to business as usual. For example, for long-run tariffs by the EU, China, and India in table 6, the EU itself incurs \$1.6B in welfare losses across EU consumer surplus and government revenue. I additionally consider EU transfers of \$32B to Indonesia and Malaysia as payment for ecosystem services, as well as \$6.9B to China and India for their participation as coalition members. There is no need for a transfer to other importers, who enjoy a welfare gain of \$27B through lower world prices. I then ask whether emission reductions are large enough to justify EU action. In doing so, I aim to assess import tariffs with distributional equity in mind, noting that I may overstate the feasibility of transfers, which are large and international, or conversely the need for transfers, which ignore that non-EU markets also desire emission reductions.²⁷

²⁷ Indeed, China and India bear 13% of the social costs of carbon. [Ricke et al. \(2018\)](#) construct country-level social costs of carbon with damage functions derived from historical climate-growth impacts. For each coalition group of interest, I sum over country-specific estimates, then I average across damage function specifications. I find that the EU, China and India, other importers, and

Figure 7: Import tariffs



I simulate EU-led import tariffs of increasing intensity. EU costs include losses to EU consumer surplus, net of tariff revenue and compensating transfers to other markets. Abatement is palm emission reductions. Costs are in inflation-adjusted, year-2000 USD per ton of CO₂, and abatement is in gigatons of CO₂. Both are totals from 1988 to 2016. Figure 7a shows long-run tariffs with coordination among all importers, an EU-China-India coalition, or the EU alone. Figure 7b shows coordinated tariffs with commitment from 1988 to 2016, 1988 to 2011, or 1988 to 2001.

Figure 7 presents the results. Even accounting for compensating transfers, I find that EU-led import tariffs reduce CO₂ emissions by up to 9.4 Gt from 1988 to 2016 at a cost to the EU of less than \$25 per ton of CO₂. Abatement at 9.4 Gt calls for coordinated, long-run import tariffs of 350%, noting that transfers serve as compensation for the welfare losses that these large tariffs impose on Indonesia and Malaysia, at least in principle. Import tariffs of 100%, as in table 6, reduce emissions by 5.4 Gt at a cost to the EU of \$15 per ton. Palm oil tariffs thus compare favorably to other means of abatement, including those receiving active EU investment.²⁸ However, the effectiveness of tariffs relies on coordination and commitment.

Figure 7a plots EU abatement costs for long-run tariffs across levels of coordination. For a target abatement cost of \$25 per ton of CO₂, an EU-China-India coalition achieves only 2.4 Gt of abatement with tariffs of 125%, while the EU itself achieves only 0.5 Gt with tariffs of 50%. Larger tariffs increase abatement, but at much higher cost because of leakage. I also note that EU-China-India tariffs are somewhat less costly at the lowest levels of abatement. At these levels, coordinated tariffs are costlier

Indonesia and Malaysia bear 3%, 13%, 82%, and 2% of the social costs of carbon, respectively.

²⁸ For now, direct air capture costs still far exceed the industry target of \$100 per ton (IEA 2022).

because they are more punitive for Indonesia and Malaysia and thus require larger compensating transfers. Unilateral tariffs are costlier because they lack the market power of a larger coalition.


Figure 7b plots EU abatement costs for coordinated tariffs across levels of commitment. For medium-run tariffs upheld from 1988 to 2011, abatement costs remain below \$25 per ton of CO₂ until 4.0 Gt of abatement with tariffs of 175%. Costs then rise convexly with a kink where the tariff coalition's initial welfare gains, which derive from market power, turn into welfare losses. At this point, the EU begins compensating transfers to China, India, and other importers to maintain the coalition. For short-run tariffs, abatement is much costlier. Tariffs upheld from 1988 to 2001 at 350% achieve 1.0 Gt of abatement at a cost of \$74 per ton. Forward-looking producers do not react strongly to short-run tariffs, and so these tariffs impose welfare costs with little abatement. The initial fall in costs occurs because small, short-run tariffs are particularly ineffective at inducing abatement.

Comparing figures 7a and 7b, I highlight the importance of commitment. In particular, I note that unilateral EU tariffs upheld from 1988 to 2016 dominate coordinated tariffs upheld from 1988 to 2001. The former achieves 1 Gt, 1.2 Gt, and 1.3 Gt of abatement at target costs of \$50, \$75, and \$100 per ton of CO₂, while the latter achieves 0 Gt, 1.0 Gt, and 1.3 Gt at the same costs. Even with full coordination across importers, abatement relies on commitment over the long run. The 14-year period from 1988 to 2001 is already not especially short, and shorter commitment would yield even less abatement. Commitment to long-run policy will be difficult, especially globally. Unilateral EU action may offer a more feasible path forward.

6.5 Export taxes

Export taxes also target traded consumption. They may appeal to Indonesia and Malaysia for several reasons. First, they allow these countries to exercise market power and raise tax revenue from foreign consumers. Second, they tax foreign but not domestic consumers, and so they raise domestic consumer surplus by shielding domestic consumers from foreign competition. Third, they are implementable with relatively limited administrative burden. Directly taxing production requires monitoring individual mills and plantations, while taxing exports requires enforcement only at international ports. Fourth, they are better than import tariffs. Export taxes

Figure 8: Export taxes



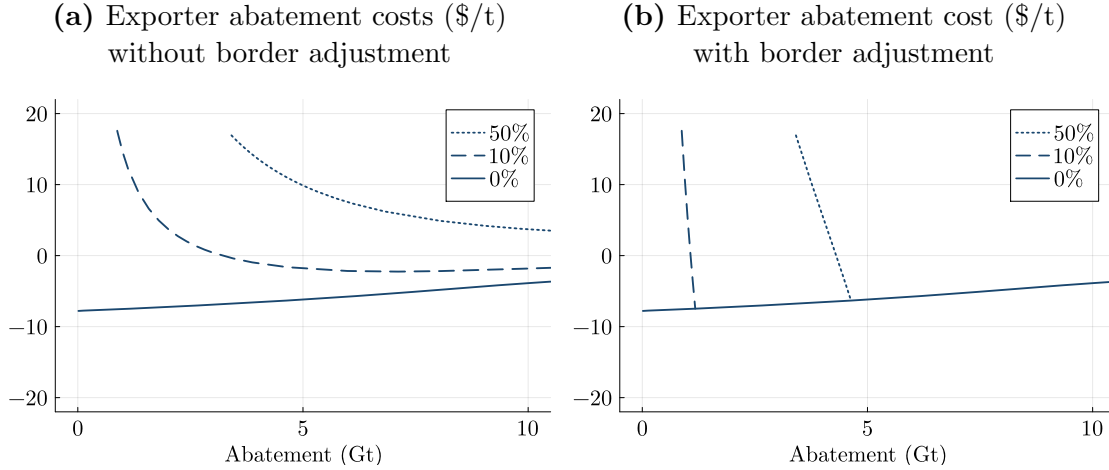
I simulate coordinated, long-run export taxes of increasing intensity by Indonesia and Malaysia. I compare these export taxes to coordinated, long-run production taxes by Indonesia and Malaysia and coordinated, long-run import tariffs lead by the EU. Figure 8a shows exporter costs, which include losses to Indonesian and Malaysian consumer and producer surplus, net of tax revenue and compensating transfers from the EU. Figure 8b shows EU costs, which include losses to EU consumer surplus, net of tariff revenue and compensating transfers to Indonesia and Malaysia. Abatement is palm emission reductions. Costs are in inflation-adjusted, year-2000 USD per ton of CO₂, and abatement is in gigatons of CO₂. Both are totals from 1988 to 2016.

by Indonesia and Malaysia and import tariffs by other markets are both aimed at the same set of goods: those that leave Indonesia and Malaysia for world markets.²⁹ Thus, both impose the same pressures to abate, but only the export taxes generate government revenue for Indonesia and Malaysia.

Figure 8a plots the costs of abatement for Indonesia and Malaysia. Negative costs imply that export taxes are welfare-enhancing, rather than welfare-reducing, for Indonesia and Malaysia. Export taxes reduce foreign consumption, which becomes increasingly inelastic as price-sensitive consumers exit. Export taxes also raise domestic consumption, which becomes increasingly elastic as domestic consumption reaches satiation. As a result, domestic losses are modest because the burden of export taxes falls primarily on foreign consumers – the inelastic party. Production taxes are also welfare-enhancing because of Indonesian and Malaysian market power, but export taxes avoid taxing domestic consumers and thus are more attractive. Import tariffs are the least attractive, even with compensating transfers from the EU that

²⁹ Note that symmetry in the sense of [Lerner \(1936\)](#) would instead compare export taxes by a given market to import tariffs by the same market.

Figure 9: Carbon border adjustment mechanism



I simulate coordinated, long-run production taxes of increasing intensity by Indonesia and Malaysia. Exporter costs include losses to Indonesian and Malaysian consumer and producer surplus, net of tax revenue but without compensating transfers from the EU. I study how these costs interact with coordinated, long-run import tariffs led by the EU, which I simulate at levels of 0%, 10%, and 50%. Abatement is palm emission reductions. Costs are in inflation-adjusted, year-2000 USD per ton of CO₂, and abatement is in gigatons of CO₂. Both are totals from 1988 to 2016. Figure 9a shows import tariffs that do not adjust with domestic regulation. Figure 9b shows a carbon border adjustment mechanism that combines import tariffs with credits for domestic regulation.

make import tariffs welfare-neutral. Without these compensating transfers, import tariffs would be welfare-reducing and even less attractive.

Figure 8b illustrates the European perspective. The EU prefers production taxes to export taxes because Indonesian and Malaysian consumers share in the tax burden when production is taxed domestically. Export taxes thus imply larger losses for EU consumer surplus at all levels of abatement. But the EU still prefers export taxes to import tariffs, which call for large compensating transfers to other markets. It is better for the EU to accept Indonesian and Malaysian export taxes in place of EU-led import tariffs, even if these export taxes are less effective than production taxes.

6.6 Carbon border adjustment mechanism

A concern is that EU-led import tariffs may crowd out domestic regulation in Indonesia and Malaysia. Figure 9a plots the costs of abatement for Indonesia and Malaysia when they impose production taxes against the backdrop of import tariffs. Focusing on coordinated, long-run production taxes and coordinated, long-run import

tariffs, I find that import tariffs raise the domestic costs of production taxes. Absent import tariffs, Indonesia and Malaysia can exercise market power to draw tax revenue from foreign consumers. Welfare rises at all levels of abatement, even as emissions fall. When importers impose tariffs of 10%, production taxes at modest levels lead to welfare losses for Indonesia and Malaysia. With import tariffs of 50%, production taxes at all levels lead to welfare losses. The reason is that import tariffs push the burden of production taxes onto domestic consumers. The intuition is clear at the extreme: if importers shut down imports by imposing infinite tariffs, then Indonesia and Malaysia lose their market power as exporters. Production taxes are then especially costly for Indonesia and Malaysia because these taxes fall solely on domestic consumers and producers.

A carbon border adjustment mechanism addresses this concern by combining import tariffs with credits for domestic regulation. Figure 9b shows that this mechanism restores Indonesian and Malaysian welfare gains from taxing production. When import tariffs fall as production taxes rise, Indonesia and Malaysia maintain market power and thus the incentive to tax production. When import tariffs fall to zero for production taxes at high levels, the three curves align on abatement at negative cost. The EU could also credit export taxes to similar effect, rather than crediting production taxes alone. Although the typical carbon border adjustment mechanism would not credit export taxes, export taxes and production taxes are similarly attractive for the EU in this setting (figure 8b).

6.7 General lessons

I discuss general lessons for green trade policy. First, a leakage problem arises from incomplete regulation. Coordinated trade policy can help, but only for traded goods. For palm oil, Indonesia and Malaysia export 80% of production, and so import tariffs have wide scope for impact. More broadly, global exports account for 68% of manufacturing GDP and 51% of agricultural GDP ([World Bank 2023](#), [WTO 2023b](#)). Both export shares are relatively large and indicate a role for trade policy. Among fossil fuels, global exports range from 54% of crude oil production to 28% for natural gas to 14% for coal ([EIA 2023](#)). Trade policy will be less effective at curbing coal emissions. The Amazon is another important frontier of deforestation, and Brazilian exports amount to 46% of soy production but only 14% of beef production ([USDA](#)

2025).³⁰ If soy expansion occurs on land previously deforested for cattle pasture, then trade policy can still play a role despite limited beef exports.

Second, a commitment problem arises from sunk emissions, which create static incentives to deregulate. Indeed, emissions are sunk in many sectors, including those accounting for the majority of traded emissions: agriculture, manufacturing, fossil fuels, mining, and transportation (Davis et al. 2011, Peters et al. 2011). For agriculture, including palm oil, emissions are sunk upon investment. Once land is cleared, the forest is gone. For other sectors, emissions are sunk – even if released gradually – if upfront investment yields low marginal costs. Once an oil well has been identified, explored, and drilled, extraction is cheap and proceeds to completion. Once older ships are built, they continue to operate and emit, even if new shipbuilding faces new regulation (Peters 2024). Committed trade policy helps by imposing long-run regulation and resisting the temptation to deregulate after emissions are sunk.

Third, trade policy need not be punitive. Compensating transfers can ensure equity, recognizing the global good that targeted markets provide by curbing emissions. These transfers act as payment for ecosystem services at global scale. Moreover, targeted markets may have their own fiscal incentives to regulate, even independent of emission concerns. Market power encourages domestic regulation, and export taxes avoid targeting domestic consumers. Trade policy can undercut these domestic incentives to regulate, but a carbon border adjustment mechanism restores the incentives.

Fourth, trade policy also faces challenges. Trade policy for palm oil should tax palm oil in all forms, but palm oil takes many forms indeed. Must a cookie importer be taxed for the 7 grams of palm oil in a 28-gram chocolate chip cookie? There is precedent for palm-based biofuels, which EU trade policy already covers: palm oil repackaged as biofuel remains subject to tariffs. But there is no such precedent in the cookie domain. Trade policy also faces political obstacles. Coordination and commitment must navigate complex, dynamic, multilateral bargaining environments. Palm oil tariffs may lead to trade disputes and escalation that I do not model, although I compute compensating transfers that acknowledge these frictions.

³⁰ I compute export-to-production ratios for the study period from 1988 to 2016. World Bank data give global agricultural and manufacturing GDP, WTO data give global total and agricultural export values, EIA data give global fossil fuel production and export volumes, and USDA data give Brazilian production and export volumes for cattle and soy. For global manufacturing, I define manufacturing as non-agriculture. For Brazilian soy, I pool oilseed, oil, and meal by weight.

7 Conclusion

Trade policy allows the international community to intervene when domestic policies fail. This paper develops a dynamic empirical framework to quantify the impacts of such policy. I use the framework to evaluate EU tariffs on imports of palm oil, a major driver of deforestation and global emissions. I document opportunities to achieve large emission reductions at low cost. Direct regulation with a production tax of 50% can reduce CO₂ emissions by 7.4 Gt. By comparison, EU import tariffs of similar magnitude can reduce emissions by 5.4 Gt if coordinated with other importers and upheld over the long run. The cost of these tariffs is only \$15 per ton of CO₂, inclusive of compensating transfers to Indonesia and Malaysia as payment for ecosystem services.

Green trade policy will be an important tool for protecting the vast forests that remain intact, at least for now. More broadly, international climate action will be crucial for meeting our global climate targets. But it relies on coordination and commitment, which are fundamentally difficult. And it imposes economic costs on lower-income countries that must also prioritize economic growth. How can we make progress in this increasingly fragmented and unequal world? Future work grounded in political realities will help to chart the path forward.

References

- Abrego, Lisandro and C Perroni. Investment subsidies and time-consistent environmental policy. *Oxford Economic Papers*, 54(4):617–635, 2002.
- Acemoglu, Daron and W Rafey. Mirage on the horizon: Geoengineering and carbon taxation without commitment. *Journal of Public Economics*, 219:104802, 2023.
- Aguirregabiria, Victor and A Magesan. Euler equations for the estimation of dynamic discrete choice structural models. *Advances in Econometrics*, 31:3–44, 2013.
- Almagro, Milena and T Domínguez-Iino. Location sorting and endogenous amenities: Evidence from Amsterdam. 2024.
- Araujo, Rafael, F Costa, and M Sant’Anna. Efficient forestation in the Brazilian Amazon: Evidence from a dynamic model. 2024.
- Arcidiacono, Peter and P Ellickson. Practical methods for estimation of dynamic

- discrete choice models. Annual Review of Economics, 3:363–394, 2011.
- Arcidiacono, Peter and R Miller. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. Econometrica, 79(6):1823–1867, 2011.
- Assunção, Juliano, R McMillan, J Murphy, and E Souza-Rodrigues. Optimal environmental targeting in the Amazon rainforest. Review of Economic Studies, 90(4):1608–1641, 2023.
- Badan Pusat Statistik. Direktori industri manufactur. 2017.
- Bergquist, Lauren and M Dinerstein. Competition and entry in agricultural markets: Experimental evidence from Kenya. American Economic Review, 110(12):3705–3747, 2020.
- Berry, Steven, J Levinsohn, and A Pakes. Automobile prices in market equilibrium. Econometrica, 63(4):841–890, 1995.
- Blevins, Jason. Nonparametric identification of dynamic decision processes with discrete and continuous choices. Quantitative Economics, 5(3):531–554, 2014.
- Blundell, Richard and JM Robin. Estimation in large and disaggregated demand systems: An estimator for conditionally linear systems. Journal of Applied Econometrics, 14(3):209–232, 1999.
- Böhringer, Christoph, J Carbone, and T Rutherford. The strategic value of carbon tariffs. American Economic Journal: Economic Policy, 8(1):28–51, 2016.
- Brunner, Steffen, C Flachsland, and R Marschinski. Credible commitment in carbon policy. Climate Policy, 12(2):255–271, 2012.
- Burgess, Robin, M Hansen, B Olken, et al. The political economy of deforestation in the tropics. Quarterly Journal of Economics, 127(4):1707–1754, 2012.
- Burgess, Robin, F Costa, and B Olken. National borders and the conservation of nature. 2024.
- Busch, Jonah, K Ferretti-Gallon, J Engelmann, et al. Reductions in emissions from deforestation from Indonesia’s moratorium on new oil palm, timber, and logging concessions. Proceedings of the National Academy of Sciences, 112(5):1328–1333, 2015.
- Chatterjee, Shoumitro. Market power and spatial competition in rural India. Quarterly Journal of Economics, 138(3):1649–1711, 2023.
- Climate Watch. Data explorer, 2020. URL www.climatewatchdata.org/data-explorer.

- Copeland, Brian and MS Taylor. North-South trade and the environment. Quarterly Journal of Economics, 109(3):755–787, 1994.
- Copeland, Brian and MS Taylor. Trade and transboundary pollution. American Economic Review, 85(4):716–737, 1995.
- Cramb, Rob and J McCarthy, editors. The Oil Palm Complex: Smallholders, Agribusiness and the State in Indonesia and Malaysia. NUS Press, 2016.
- Davis, Steven, G Peters, and K Caldeira. The supply chain of CO₂ emissions. Proceedings of the National Academy of Sciences, 108(45):18554–18559, 2011.
- De Groote, Olivier and F Verboven. Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. American Economic Review, 109(6):2137–2172, 2019.
- Deaton, Angus and J Muellbauer. An almost ideal demand system. American Economic Review, 70(3):312–326, 1980.
- Diamond, Rebecca, T McQuade, and F Qian. Who benefits from rent control? the equilibrium consequences of San Francisco’s rent control expansion. 2019.
- Direktorat Jenderal Perkebunan. Statistik perkebunan Indonesia 2017–2019: Kelapa sawit. 2018.
- Domínguez-Iino, Tomás. Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains. 2025.
- Duflo, Esther, M Greenstone, R Pande, and N Ryan. The value of regulatory discretion: Estimates from environmental inspections in India. Econometrica, 86(6): 2123–2160, 2018.
- Edwards, Ryan. Export agriculture and rural poverty: Evidence from Indonesian palm oil. 2019.
- Edwards, Ryan, W Falcon, G Hadiwidjaja, et al. Fight fire with finance: A randomized field experiment to curtail land-clearing fire in Indonesia. 2020.
- Elliott, Joshua, I Foster, S Kortum, et al. Trade and carbon taxes. American Economic Review: Papers & Proceedings, 100:465–469, 2010.
- Energy Information Administration. International energy data, 2023. URL www.eia.gov/international/data/world.
- Erdem, Tülin, S Imai, and M Keane. Brand and quantity choice dynamics under price uncertainty. Quantitative Marketing and Economics, 1:5–64, 2003.
- Ericson, Richard and A Pakes. Markov-perfect industry dynamics: A framework for empirical work. Review of Economic Studies, 62(1):53–82, 1995.

- Fairhurst, Thomas and D McLaughlin. Sustainable oil palm development on degraded land in Kalimantan. Technical report, World Wildlife Fund, 2009.
- Farrokhi, Farid and A Lashkaripour. Can trade policy mitigate climate change? 2025.
- Fick, Stephen and R Hijmans. WorldClim 2: New 1-km spatial resolution climate surfaces for global land areas. International Journal of Climatology, 37(12):4302–4315, 2017.
- Food and Agriculture Organization. ECOCROP database, 2025. URL [gaez.fao.org/pages/ecocrop](https://www.fao.org/pages/ecocrop).
- Fowlie, Meredith. Incomplete environmental regulation, imperfect competition, and emissions leakage. American Economic Journal: Economic Policy, 1(2):72–112, 2009.
- Fowlie, Meredith, M Reguant, and S Ryan. Market-based emissions regulation and industry dynamics. Journal of Political Economy, 124(1):249–302, 2016.
- GADM. Global administrative boundaries version 3.6, 2021. <https://www.gadm.org>.
- Gaveau, David, B Locatelli, M Salim, et al. Rise and fall of forest loss and industrial plantations in Borneo (2000–2017). Conservation Letters, 12(3):e12622, 2019.
- Gumbrecht, Thomas, RM Roman-Cuesta, L Verchot, et al. An expert system model for mapping tropical wetlands and peatlands reveals South America as the largest contributor. Global Change Biology, 23:3581–3599, 2017.
- Hall, Robert. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. Journal of Political Economy, 86(6):971–987, 1978.
- Hansen, Lars Peter and K Singleton. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica, 50(5):1269–1286, 1982.
- Harstad, Bård. Buy coal! A case for supply-side environmental policy. Journal of Political Economy, 120(1), 2012.
- Harstad, Bård. Technology and time inconsistency. Journal of Political Economy, 128(7):2653–2689, 2020.
- Harstad, Bård. Trade and trees. American Economic Review: Insights, 6(2):155–175, 2024.
- Harstad, Bård and T Mideksa. Conservation contracts and political regimes. Review of Economic Studies, 84:1708–1734, 2017.
- Helm, Dieter, C Hepburn, and R Mash. Credible carbon policy. Oxford Review of Economic Policy, 19(3):438–450, 2003.
- Hendel, Igal and A Nevo. Measuring the implications of sales and consumer inventory

- behavior. Econometrica, 74(6):1637–1673, 2006.
- Hoel, Michael. Should a carbon tax be differentiated across sectors? Journal of Public Economics, 59:17–32, 1996.
- Hoffmann, Munir, A Castañeda-Vera, M van Wijk, et al. Simulating potential growth and yield of oil palm (*Elaeis guineensis*) with PALMSIM: Model description, evaluation and application. Agricultural Systems, 131:1–10, 2014.
- Hopenhayn, Hugo. Entry, exit, and firm dynamics in long run equilibrium. Econometrica, 60(5):1127–1150, 1992.
- Hotz, V. Joseph and R Miller. Conditional choice probabilities and the estimation of dynamic models. Review of Economic Studies, 60(3):497–529, 1993.
- Hsiao, Allan. Sea level rise and urban adaptation in Jakarta. 2025.
- Hsiao, Allan, J Moscona, and K Sastry. Food policy in a warming world. 2025.
- IndexMundi. Commodity prices, 2019. URL www.indexmundi.com/commodities.
- International Energy Agency. Direct air capture: A key technology for net zero. 2022.
- International Monetary Fund. Primary commodity prices, 2023. URL www.imf.org/en/Research/commodity-prices.
- Iskhakov, Fedor, T Jørgensen, J Rust, and B Schjerning. The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. Quantitative Economics, 8(2):317–365, 2017.
- Jayachandran, Seema, J de Laat, E Lambin, et al. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. Science, 357:267–273, 2017.
- Kalouptsi, Myrto, P Scott, and E Souza-Rodrigues. Linear IV regression estimators for structural dynamic discrete choice models. Journal of Econometrics, 222(1):778–804, 2021.
- Kortum, Samuel and D Weisbach. The design of border adjustments for carbon prices. National Tax Journal, 70(2):421–446, 2017.
- Kortum, Samuel and D Weisbach. Optimal unilateral carbon policy. 2024.
- Lerner, Abba. The symmetry between import and export taxes. Economica, 3(11):306–313, 1936.
- Magnac, Thierry and D Thesmar. Identifying dynamic discrete decision processes. Econometrica, 70(2):801–816, 2002.
- Malaysian Palm Oil Board. Number and capacities of palm oil sectors. 2016.
- Malaysian Palm Oil Board. Malaysian oil palm statistics. 2018.

- Man, Elaine and A Baharum. A qualitative approach of identifying major cost influencing factors in palm oil mills and the relations towards production cost of crude palm oil. American Journal of Applied Sciences, 8(5):441–446, 2011.
- Markusen, James. International externalities and optimal tax structures. Journal of International Economics, 5:15–29, 1975.
- Marsiliani, Laura and T Renström. Time inconsistency in environmental policy: Tax earmarking as a commitment solution. Economic Journal, 110(462):C123–C138, 2000.
- Meijer, Johan, M Huijbregts, K Schotten, and A Schipper. Global patterns of current and future road infrastructure. Environmental Research Letters, 13(6):064006, 2018.
- Murphy, Alvin. A dynamic model of housing supply. American Economic Journal: Economic Policy, 10(4):243–267, 2018.
- National Geospatial-Intelligence Agency. World port index, 2019. URL msi.nga.mil/Publications/WPI.
- Nordhaus, William. Climate clubs: Overcoming free-riding in international climate policy. American Economic Review, 105(4):1339–1370, 2015.
- Oates, Wallace and P Portney. The political economy of environmental policy. Handbook of Environmental Economics, 1:325–354, 2003.
- Official Journal of the EU. Council directive 2011/64/EU of 21 June 2011. 2011.
- Official Journal of the EU. Regulation (EU) 2023/1115 of the European Parliament and of the Council of 31 May 2023. 2023.
- Okarda, Beni and P Manalu. Oil palm mills database, 2017. URL doi.org/10.17528/CIFOR/DATA.00098.
- Oliva, Paulina. Environmental regulations and corruption: Automobile emissions in Mexico City. Journal of Political Economy, 123(3):686–724, 2015.
- Palm Oil Analytics. Essential palm oil statistics. 2017.
- Peters, Allen. Beached assets? Capital turnover and emissions in shipping. 2024.
- Peters, Glen, J Minx, C Weber, and O Edenhofer. Growth in emission transfers via international trade from 1990 to 2008. Proceedings of the National Academy of Sciences, 108(21):8903–8908, 2011.
- Rauscher, Michael. International Trade, Factor Movements, and the Environment. Oxford University Press, 1997.
- Ricke, Katharine, L Drouet, K Caldeira, and M Tavoni. Country-level social cost of

- carbon. Nature Climate Change, 8:895–900, 2018.
- Rubens, Michael. Market structure, oligopsony power, and productivity. American Economic Review, 113(9):2382–2410, 2023.
- Sant’Anna, Marcelo. How green is sugarcane ethanol? Review of Economics and Statistics, 106(1):202–216, 2024.
- Scott, Paul. Dynamic discrete choice estimation of agricultural land use. 2013.
- Shapiro, Joseph. The environmental bias of trade policy. Quarterly Journal of Economics, 136(2):831–886, 2021.
- Sheffield, Justin, G Goteti, and E Wood. Development of a 50-year high-resolution global dataset of meteorological forcings for land surface modeling. Journal of Climate, 19(13):3088–3111, 2006.
- Sofiyuddin, Muhammad, A Rahmanulloh, and S Suyanto. Assessment of profitability of land use systems in Tanjung Jabung Barat district, Jambi province, Indonesia. Open Journal of Forestry, 2(4):252–256, 2012.
- Song, Xiao-Peng, M Hansen, S Stehman, et al. Global land change from 1982 to 2016. Nature, 560:639–643, 2018.
- Souza-Rodrigues, Eduardo. Deforestation in the Amazon: A unified framework for estimation and policy analysis. Review of Economic Studies, 86:2713–2744, 2019.
- Traiberman, Sharon. Occupations and import competition: Evidence from Denmark. American Economic Review, 109(12):4260–4301, 2019.
- USDA Foreign Agricultural Service. Indonesia biofuels annual report. 2019a.
- USDA Foreign Agricultural Service. Malaysia biofuels annual report. 2019b.
- USDA Foreign Agricultural Service. International production assessment division crop calendars, 2021. URL ipad.fas.usda.gov.
- USDA Foreign Agricultural Service. Production, supply and distribution online, 2025. URL apps.fas.usda.gov/psdonline.
- Wagstaff, Kiri, C Cardie, S Rogers, and S Schroedl. Constrained k-means clustering with background knowledge. Proceedings of the Eighteenth International Conference on Machine Learning, 1:577–584, 2001.
- Warren, Matthew, K Hergoualc’h, JB Kauffman, et al. An appraisal of Indonesia’s immense peat carbon stock using national peatland maps: Uncertainties and potential losses from conversion. Carbon Balance and Management, 12:12, 2017.
- World Bank. Indonesia database for policy and economic research (INDO-DAPOER), 2022. URL databank.worldbank.org/source/

- [indonesia-database-for-policy-and-economic-research](#).
- World Bank. World Bank open data, 2023. URL [data.worldbank.org](#).
- World Port Source. Ports by country, 2020. URL [www.worldportsource.com/countries.php](#).
- World Resources Institute. Universal mill list, 2019. URL [data.globalforestwatch.org/documents/gfw:universal-mill-list](#).
- World Trade Organization. Tariff analysis online, 2023a. URL [tao.wto.org](#).
- World Trade Organization. WTO stats, 2023b. URL [stats.wto.org](#).
- World Wildlife Fund. Palm oil buyers scorecard: Measuring the progress of palm oil buyers. 2016.
- Xu, Yidi, L Yu, W Li, et al. Annual oil palm plantation maps in Malaysia and Indonesia from 2001 to 2016. Earth System Science Data, 12:847–867, 2020.
- Zarin, Daniel, N Harris, A Baccini, et al. Can carbon emissions from tropical deforestation drop by 50% in 5 years? Global Change Biology, 22:1336–1347, 2016.
- Zavala, Lucas. Unfair trade? Monopsony power in agricultural value chains. 2024.

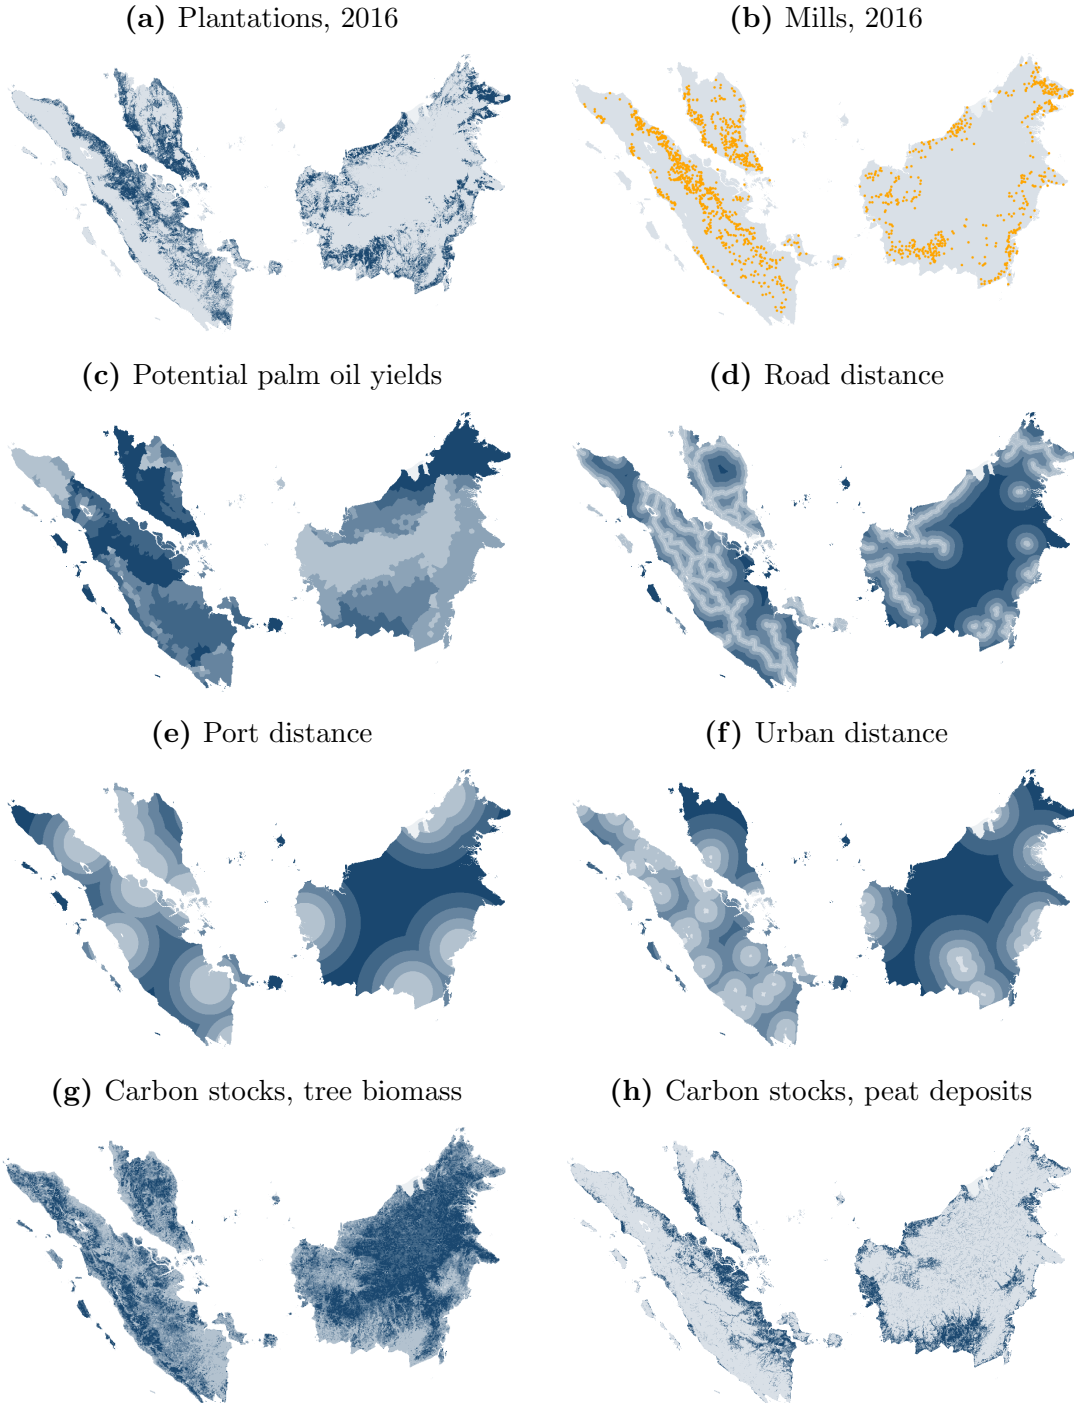
ONLINE APPENDIX

A Data

Table A1: Data sources

Source	Period	Coverage	Description
USDA (2025)	1960-2019	World	Annual consumption, production, area harvested, imports, and exports by country and oilcrop
IMF (2023), IndexMundi (2019)	1980-2019	World	Monthly prices by oilcrop
Sheffield et al. (2006)	1980-2016	World	Daily precipitation and temperature, 28km resolution
USDA (2021), FAO (2025)	—	World	Crop calendars and growing conditions
Xu et al. (2020)	2001-2016	Indonesia, Malaysia	Palm oil plantations, 100m resolution
Song et al. (2018)	1982-2016	World	Land cover change, 5.6km resolution
WRI (2019)	2018	Indonesia, Malaysia	List of mill coordinates
Okarda and Manalu (2017)	2017	Indonesia	List of mill coordinates
GADM (2021)	2018	World	GIS maps of administrative boundaries
Fick and Hijmans (2017)	1970-2000	World	Average monthly solar radiation and precipitation
DJP (2018), World Bank (2022)	1996-2017	Indonesia	Annual yields by province
MPOB (2018)	1990-2018	Malaysia	Annual yields by state
Meijer et al. (2018)	2018	World	Road networks
NGA (2019), WPS (2020)	2019-2020	World	Port coordinates
Zarin et al. (2016)	2000	World	Aboveground biomass, 30m resolution
Gumbricht et al. (2017)	2011	World	Peatland deposits, 231m resolution
Climate Watch (2020)	1990-2016	World	CO ₂ emissions by country

Figure A1: Plantations, mills, yields, and cost factors



I study Sumatra, Kalimantan, Peninsular Malaysia, and East Malaysia, which together total 134 Mha in surface area. Darker blue indicates higher yields, farther distances, and larger carbon stocks. Urban areas include administrative cities (Indonesia) and federal territories (Malaysia).

Plantations and mills

Spatial panel data on palm oil plantations from 2001 to 2016 come from [Xu et al. \(2020\)](#), who construct the data at 100m resolution from Phased Array type L-band Synthetic Aperture Radar (PALSAR), PALSAR-2, and Moderate Resolution Imaging Spectroradiometer (MODIS) satellite imagery. The data measure how much of each tile is covered by palm oil plantations, inclusive of both young and mature as well as industrial and smallholder plantations. I use midpoints of the upper and lower bounds in years where bounds are provided and point estimates otherwise. I aggregate the data to 1km resolution by averaging. As discussed in [Xu et al. \(2020\)](#), I impose that development is uni-directional, such that plantation area for each tile is non-decreasing over time. [Xu et al. \(2020\)](#) restrict their attention to Sumatra, Kalimantan, Peninsular Malaysia, and East Malaysia, and I do the same in my analysis. These regions cover virtually all palm oil production in Indonesia and Malaysia during the study period.

I extend the plantations data back to 1988 using data on tree canopy cover from [Song et al. \(2018\)](#), who analyze satellite imagery from the Advanced Very High Resolution Radiometer (AVHRR), MODIS, and Landsat Enhanced Thematic Mapper Plus (ETM+). These data extend from 1982 to 2016, overlapping the [Xu et al. \(2020\)](#) data from 2001 to 2016. Focusing on tiles that the [Xu et al. \(2020\)](#) data identify as having plantations, I estimate the empirical relationship between plantation development and tree cover loss during the period of overlap, and I use these estimates to impute plantations prior to 2001. For tiles i in years t ,

$$\Delta\text{Plantation}_{it} = \sum_{s=0}^3 \beta_s \Delta\text{Tree cover}_{it-s} + \varepsilon_{it},$$

where $\Delta\text{Plantation}_{it}$ is new plantation development and $\Delta\text{Tree cover}_{it-s}$ terms are tree cover loss in preceding years. The [Song et al. \(2018\)](#) data are at 5.6km resolution, so I downscale them to match the 1km resolution of the aggregated [Xu et al. \(2020\)](#) data. Table [A2](#) shows the resulting estimates. Significant, negative coefficients indicate that tree cover loss, especially in the preceding two years, is predictive of plantation development. I take third column, which includes tile fixed effects, as my preferred specification. For each tile, I combine predicted plantation development

Table A2: Plantations vs. tree cover (2001-2016)

	$\Delta\text{Plantation}_t$	$\Delta\text{Plantation}_t$	$\Delta\text{Plantation}_t$
$\Delta\text{Tree cover}_t$	-0.00314*** (0.000156)	-0.00253*** (0.000155)	-0.00261*** (0.000153)
$\Delta\text{Tree cover}_{t-1}$	-0.00524*** (0.000192)	-0.00441*** (0.000191)	-0.00435*** (0.000190)
$\Delta\text{Tree cover}_{t-2}$	-0.00102*** (0.000194)	0.000203 (0.000193)	0.000414** (0.000193)
$\Delta\text{Tree cover}_{t-3}$	-0.000672*** (0.000162)	6.42e-05 (0.000161)	7.27e-05 (0.000160)
Year FE	x	x	x
District FE		x	
Tile FE			x
Observations	9,098,040	9,098,040	9,098,040


Each column is a regression, and each observation is a tile-year. I regress plantation development on tree canopy cover. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

with observed levels in 2001 to impute pre-2001 plantations, imposing a minimum of zero for plantation development. The downscaling of the coarser [Song et al. \(2018\)](#) implies that the imputed data should not be analyzed below 5.6km resolution, and indeed my core analysis analyzes aggregated sites and not individual tiles.

Figure [A2](#) plots the resulting data. First, imposing uni-directional development rules out exit. Limited exit makes this assumption reasonable: [Xu et al. \(2020\)](#) measure cumulative exit of only 4.6% between 2007 and 2016, noting that exit can include misclassified oil palm. Second, the tree cover data imply a reasonable pattern of plantation development pre-2001. Third, I verify the quality of the satellite data, both observed and imputed, by comparing them to aggregate government estimates compiled by [USDA \(2025\)](#). The data match well, although the satellite data reveals modestly higher levels of plantation development in later years.

Spatial data on palm oil mills come from the 2018 Universal Mill List, a joint effort led by the World Resources Institute and Rainforest Alliance that collects data from palm oil processors, traders, corporate consumers, and NGOs ([WRI 2019](#)). Mills are geocoded and manually verified by satellite. I combine these data with the 2017 Center for International Forestry Research database, an independent effort that combs traceability reports for major palm oil processors and also verifies coordinates manually by satellite ([Okarda and Manalu 2017](#)). I merge the datasets spatially, match-

Figure A2: Total plantations (Mha)



The left figure imposes no exit, the middle figure extends the plantation data using tree cover data, and the right figure shows government data. Plantation area is in megahectares.

ing mills within one kilometer of each other, and I validate mills with Landsat and DigitalGlobe satellite images from Google Earth by identifying nearby plantations, storage tanks, and effluent ponds. I omit mills in Java, which houses refineries and administrative offices but few plantations. I correct coordinates where necessary.

I identify 1,526 mills, of which 1,497 lie within the study area. I obtain 1,482 mills that are consistent with the plantation data, as I will describe in the next section. Table A3 validates the data against official government data from Statistics Indonesia and the Malaysian Palm Oil Board. Within the study area, I identify 805 establishments that produce palm oil as their main product in the 2017 Indonesian manufacturing directory (BPS 2017). The Malaysian government data record 453 mills in 2016 (MPOB 2016). My baseline mill data capture a larger number of mills in Indonesia, where the official directory is constrained to medium and large firms with complete listing information, but the spatial distribution remains similar. I match the Malaysian data closely, while also capturing a slightly larger number of mills in Peninsular Malaysia.

Sites

To divide land into sites, I first compute the maximum number of sites \bar{k} for each province: $\bar{k} = \max\{\text{floor}(\text{area}/535), \text{number of observed mills}\}$. I use a benchmark site size of 535 km², which I obtain as the average of two calculations. The first

Table A3: Mill counts by region (2016)

	Mill data	Official data
Indonesia	1,010	805
Sumatra	681	567
Kalimantan	329	238
Malaysia	472	453
Peninsular Malaysia	266	247
East Malaysia	206	206
Total	1,482	1,258

Government data come from Statistics Indonesia and the Malaysian Palm Oil Board.


calculation considers provinces with the highest mill densities observed in the data. For each province, I compute the mill density at the end of the study period in 2016. The 75th-percentile province has a density of one mill per 453 km². The second calculation considers circular sites with radii given by the largest plantation-mill distances observed in the data. For each province, I compute the average distance between plantations and their closest mills in 2016. The 75th-percentile province has a distance of 14 km, and a radius of 14 km gives a circular site size of 617 km².

Second, I define sites by k -means clustering on geographic coordinates. I ensure consistency with the plantations and mills observed in 2016 by imposing (1) that observed mills be assigned to unique sites and (2) that observed plantations be clustered with observed mills. I adapt the constrained k -means clustering algorithm of [Wagstaff et al. \(2001\)](#). I apply multiple starts because convergence is to local optima.

1. Choose initial cluster centers C_1, C_2, \dots, C_k .
2. For the m mills observed in the data, move the m closest centers to the mill coordinates.
3. Assign points to the nearest cluster centers.
4. Update each cluster center by averaging over the points assigned to it.
5. Repeat (2) to (4) until convergence.
6. For clusters without mills but with at least 10 km² of plantations, reassign all points to the nearest clusters with mills.

Step (2) ensures consistency with observed mills, and step (6) with observed plantations. Figure [A3](#) maps the resulting 2,050 potential sites.

Figure A3: Potential sites



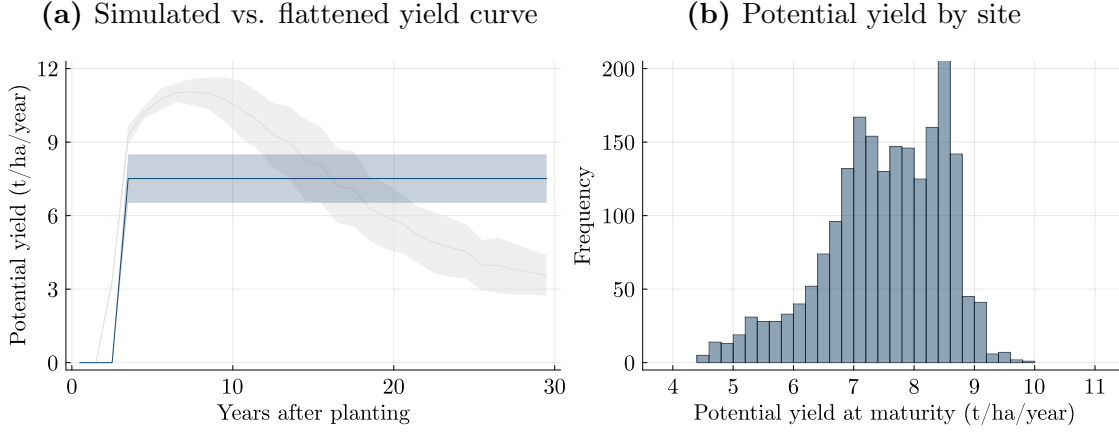
Blue shading indicates different potential sites. Orange dots are palm oil mills observed by 2016.

I overlay plantations and mills with the site boundaries to construct a panel of palm oil investment by site and year. From the panel data on plantations, I compute plantation area for each site, and I define plantation development as the change in area each year. From the cross-sectional data on mills, I define mill construction as the first year in which I observe meaningful plantations. I define the cutoff for “meaningful” to be 413 hectares, which I choose so that average plantation development in the year of mill construction is equal to average plantation development in the years after mill construction. A cutoff of zero hectares would yield little variation on the extensive margin, as measurement error in the satellite data and classification error in clustering sites each place small patches of plantations on most sites, even in early years. Finally, I harmonize the data by dropping plantations without mills and mills without plantations. I drop 0.3% of plantations and 1% of mills observed in 2016.

Yields

I construct site-level data on palm oil yields over time. First, I compute potential yields by site using the agronomic PALMSIM model of [Hoffmann et al. \(2014\)](#), which predicts yields under optimal growing conditions as a function of climate. As inputs, I use average monthly solar radiation and precipitation data at 1km resolution from WorldClim. To facilitate computation, I run the PALMSIM model on average climate measures by site. Figure [A4a](#) shows the resulting 30-year yield curve, which starts at

Figure A4: Potential palm oil yields



Yield curves are computed from the PALMSIM model (Hoffmann et al. 2014) using site-level average monthly solar radiation and precipitation from WorldClim. On the left, the gray curve shows the average output of the PALMSIM model, and the navy blue line flattens the curve to two levels – “immature” (zero-yield) and “mature” – while maintaining the same average over time. Shaded areas show standard deviations. On the right, I show the dispersion of (flattened) mature yields across sites. Yields are in tons per hectare per year.

zero before increasing steeply then declining gradually. Because the data on actual yields distinguish only between “immature” and “mature” crops, I flatten the curve to these levels while fixing average yields over time. Figure A4b plots these flattened yields at maturity. The yields are time-invariant because yields under optimal conditions are an inherent characteristic of the oil palm plant.

Second, I compile data on actual yields by province and year, drawing on government statistics from the Indonesian Ministry of Agriculture, the World Bank INDO-DAPOER database, and the Malaysian Palm Oil Board (DJP 2018, WB 2022, MPOB 2018). Each reports yields for mature crops, omitting immature crops that do not yet produce fruit. Figure A5a shows that, on average, actual yields are increasing over time as technology improves, although they fall far short of potential yields in all provinces and years. Crop age mix also affects yields over time, but two effects are potentially offsetting: young crops approaching their peak have increasing yields, while aging crops past their peak have decreasing yields. Across provinces and years, the average observed annual yield per hectare is 3.3 tons.

Third, I combine these data to compute actual yields y_{it} by site i and year t . I assume that y_{it} combines yield gaps γ_{mt} and potential yields y_i^p , where yield gaps are

Figure A5: Actual palm oil yields (t/ha/year)

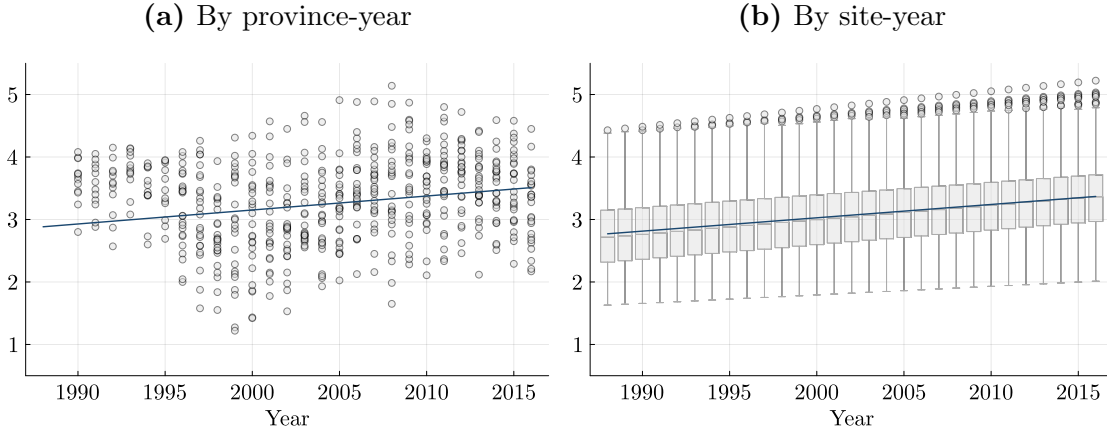


Figure A5a is annual yields by province (Indonesia) or state (Malaysia) as recorded in government statistics. Figure A5b is annual yields by site as computed from potential yields by site (PALMSIM) and actual yields by province (government statistics). Yields are in tons per hectare per year.

fixed across sites within a given province m and year t .

$$y_{it} = (1 - \gamma_{mt})y_i^p \quad (11)$$

I observe actual yields y_{mt} by province-year, which aggregate over site-years. That is, $y_{mt} = (\sum_{i \in \mathcal{I}_m} y_{it} N_{it}) / (\sum_{i \in \mathcal{I}_m} N_{it})$ for plantation acreage N_{it} . I substitute equation 11 to obtain yield gaps γ_{mt} .


$$\gamma_{mt} = 1 - y_{mt} \left(\frac{\sum_{i \in \mathcal{I}_m} y_i^p N_{it}}{\sum_{i \in \mathcal{I}_m} N_{it}} \right)^{-1}$$

Thus, I can compute yield gaps as a function of observed quantities: actual yields y_{mt} , potential yields y_i^p , and plantation acreage N_{it} . I then isolate the underlying levels and trends of these yield gaps with the specification

$$\gamma_{mt} = \alpha_m + \beta t + \varepsilon_{mt},$$

which also allows me to extrapolate back to 1988 and beyond 2016. I compute fitted values and substitute into equation 11 to obtain actual yields y_{it} by site-year. Figure A5b shows the resulting estimates, which combine the uptrend of figure A5a with the site-level dispersion of figure A4b.

Figure A6: Plantation development vs. emissions



Emissions are in gigatons of CO₂, and plantation development area is in megahectares.

Carbon stocks

I observe carbon stocks over space with data on aboveground tree biomass from [Zarin et al. \(2016\)](#) and on belowground peat biomass from [Gumbricht et al. \(2017\)](#). I convert aboveground biomass to carbon with a typical biomass-to-carbon conversion factor of 0.5. I convert belowground biomass with a conversion factor of 65.1 kg C/m³ peat, as in [Warren et al. \(2017\)](#). I convert carbon to CO₂ emissions with a molecular-weight conversion factor of 3.67.³¹ I treat carbon stocks as predetermined, but they are not measured before the study period. Tree biomass is measured for the year 2000, and peat deposits for 2011. The data may therefore miss carbon stocks destroyed before these years. For tree biomass, I impute 1988 values by combining the 2000 values with the proportion of tree cover loss between 1988 and 2000, as measured by [Song et al. \(2018\)](#). For peat deposits, bias is limited because [Gumbricht et al. \(2017\)](#) rely primarily on precipitation and topography – predetermined features – to identify wetlands as areas where water is likely to pool because precipitation exceeds evapotranspiration. MODIS satellite imagery from 2011 then allow the authors to distinguish between different kinds of wetlands. Indeed, figure A6 shows that the relationship between plantation development and the resulting emissions is consistent over time. If the data missed peatlands destroyed before 2011, then peatland emissions would be much smaller for plantation development before 2011.

³¹ I focus on CO₂ emissions because the carbon content of peatlands is well documented and because they account for 73% of total greenhouse gas emissions during the study period ([Climate Watch 2020](#)). Palm oil production also involves the release of methane and nitrous oxide, but precise estimates of these emissions are not yet well established.

Table A4: Vegetable oil producers

Oil	Producers
Coconut	Philippines 52%, Indonesia 33%, India 15%
Olive	EU 86%, Tunisia 8%, Turkey 6%
Palm	Indonesia 49%, Malaysia 45%, Nigeria 6%
Rapeseed	EU 36%, China 27%, Canada 23%, India 14%
Soybean	US 44%, Brazil 29%, Argentina 18%, China 8%
Sunflower	EU 29%, Russia 23%, Ukraine 23%, Argentina 17%, China 8%

Each row sums to 100% and covers 1988 to 2016. I omit producers below 5% of world production.

Weather shocks to vegetable oil production

Weather data come from the Global Meteorological Forcing Dataset, which records daily rainfall and average surface temperature from 1988 to 2016 at 0.25° resolution. I use these data to construct annual measures of weather shocks to the production of coconut, olive, palm, rapeseed, soybean, and sunflower oils over the study period. I omit cottonseed and peanut oils given a lack of price data and relatively small volumes at 5% of vegetable oil consumption volume in 2016.

First, I isolate day-tile observations within oil-producing regions and during the growing season. I define oil-producing regions as countries that account for at least 5% of world production for any of the aforementioned oils during the study period, as measured by data from the USDA Foreign Agricultural Service. Table A4 lists these countries for each oil (aggregating EU countries). For Argentina, Brazil, Canada, China, India, Indonesia, Malaysia, Russia, and the United States, I further consider subnational regions – namely states and provinces – using data from the USDA and national government sources. I define country-specific growing seasons for rapeseed, soybean, and sunflower oils to be as specified in USDA crop calendars (USDA 2021), and I take the growing season for coconut, olive, and palm oils to be year-round.

Second, I compute crop-specific weather shocks at the year-tile level. For rainfall, I first aggregate from daily to monthly values for each tile, as daily variation in rainfall is less detrimental to crop growth than daily variation in temperature. I then compute shocks as absolute deviations from optimal levels for each crop. The FAO ECOCROP database records optimal windows by crop for both rainfall and temperature (FAO 2025), and I take the midpoint of these windows as optimal levels. The FAO database

Table B1: Weather shocks as price instruments

	All	All	Palm	Other
Rainfall shocks (100 mm)	0.208*** (0.0318)	0.212*** (0.0279)	0.139*** (0.0326)	0.224*** (0.0320)
Temperature shocks (°C)	0.297*** (0.0335)	0.308*** (0.0316)	0.687 (0.807)	0.315*** (0.0335)
Oil FE	x	x		
Oil-year trend		x		
Year trend			x	x
Observations	174	174	29	145
F-statistic	40.78	49.03	10.46	48.68

Each column is a regression, and the outcome variable is log prices. Annual data cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. Weather shocks are absolute deviations from optimal conditions during the growing season, aggregated over producing provinces and states. Oil fixed effects and oil-specific time trends differentiate between palm and other oils. Newey-West standard errors account for serial correlation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

specifies optimal annual rainfall, which I divide by twelve to obtain optimal monthly rainfall. Having computed monthly deviations from optimal levels for rainfall, as well as daily deviations for temperature, I aggregate over time to obtain average deviations by year for each tile.

Third, I aggregate to obtain annual weather shocks by vegetable oil. I average over tiles for each oil-producing region, and then I average across oil-producing regions for each oil in proportion to production volumes. I weight by total production over the study period rather than annual production, which depends on annual weather.

B Demand

Table B1 shows that weather shocks increase world vegetable oil prices in the first stage. The first two columns pool across oil products, and the last two consider palm and other oils separately. The instruments remain strong for palm oil despite a smaller sample size, although imprecise temperature estimates reflect limited temperature variation in the palm-producing tropics. Table B2 shows that weather shocks do not have domestic income or expenditure effects, which would violate the exclusion restriction by influencing demand beyond the price channel. These results also provide reassurance that the instruments do not simply capture idiosyncratic fluctuations in

Table B2: Weather shocks vs. incomes and expenditures

Market	Outcome	Rainfall		Temperature		Obs
		Estimate	SE	Estimate	SE	
European Union	CPI	0.00357	(0.00277)	0.00259	(0.00246)	174
	GDP	0.00517	(0.00765)	0.00395	(0.00737)	174
	GDE	0.00574	(0.00785)	0.00424	(0.00749)	174
China/India	CPI	0.00644	(0.0110)	0.00353	(0.0113)	174
	GDP	1.10e-05	(0.0103)	-0.00350	(0.00987)	174
	GDE	-0.00166	(0.00971)	-0.00436	(0.00923)	174
Other importers	CPI	0.00573	(0.00779)	0.000981	(0.00787)	174
	GDP	0.00352	(0.00450)	0.00172	(0.00411)	174
	GDE	0.00422	(0.00417)	0.00228	(0.00373)	174
Indonesia/Malaysia	CPI	-0.0231	(0.0247)	-0.0220	(0.0243)	174
	GDP	0.0113	(0.0154)	0.00534	(0.0157)	174
	GDE	0.00921	(0.0147)	0.00420	(0.0152)	174

Each row is a regression. Annual data cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. For outcome variables, GDPs and GDEs are in logs, and CPIs aggregate national data weighted by household GDE. Weather shocks are absolute deviations from optimal conditions during the growing season, aggregated over producing regions. Oil fixed effects and oil-specific time trends differentiate between palm and other oils. Newey-West standard errors account for serial correlation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B3: Demand elasticities for other oils

	Palm		Other	
	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.943***	(0.024)
China/India	-0.692***	(0.168)	-0.867***	(0.045)
Other importers	-0.876***	(0.128)	-0.892***	(0.036)
Indonesia/Malaysia	-0.925***	(0.046)	-0.531*	(0.309)

Each pair of columns shows own-price elasticities by consumer market. I report elasticities of total consumption with respect to a 1% increase in prices from 1988 to 2016. Other oils aggregate over non-palm vegetable oils. Demand estimation draws on annual data covering coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

macroeconomic conditions. Table B3 shows that own-price elasticities for other oils are similar to those for palm oil. Elasticities for other oils are more precisely estimated for the EU, China, India, and other importers, where observed consumption of other oils is high. They are much less precisely estimated for Indonesia and Malaysia, where observed consumption of other oils is particularly low.

Table B4: Demand elasticities with dynamics and longer-run variation

Price lags	None		One-year		Two-year		Three-year	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.572**	(0.248)	-0.547**	(0.236)	-0.505*	(0.262)
China/India	-0.692***	(0.168)	-0.653***	(0.160)	-0.786***	(0.110)	-0.733***	(0.100)
Other importers	-0.876***	(0.128)	-0.786***	(0.134)	-0.761***	(0.132)	-0.761***	(0.150)
Indonesia/Malaysia	-0.925***	(0.046)	-0.949***	(0.039)	-0.949***	(0.033)	-0.965***	(0.032)
Price leads	None		One-year		Two-year		Three-year	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.844***	(0.169)	-0.872***	(0.188)	-0.851***	(0.186)
China/India	-0.692***	(0.168)	-0.646***	(0.171)	-0.727***	(0.178)	-0.733***	(0.221)
Other importers	-0.876***	(0.128)	-0.876***	(0.119)	-0.860***	(0.137)	-0.810***	(0.147)
Indonesia/Malaysia	-0.925***	(0.046)	-0.940***	(0.052)	-0.936***	(0.047)	-0.988***	(0.034)
Rolling maximum	None		One-year		Two-year		Five-year	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.554***	(0.166)	-0.488*	(0.265)	-0.624***	(0.212)
China/India	-0.692***	(0.168)	-0.437***	(0.097)	-0.662***	(0.206)	-0.885***	(0.130)
Other importers	-0.876***	(0.128)	-0.700***	(0.143)	-0.731***	(0.277)	-0.295	(0.845)
Indonesia/Malaysia	-0.925***	(0.046)	-1.010***	(0.040)	-1.046***	(0.051)	-1.045***	(0.086)
Rolling minimum	None		One-year		Two-year		Five-year	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
European Union	-0.723***	(0.210)	-0.393	(0.518)	-0.271	(0.712)	-1.142***	(0.399)
China/India	-0.692***	(0.168)	-0.668*	(0.404)	-0.839**	(0.416)	-0.971***	(0.285)
Other importers	-0.876***	(0.128)	-0.927***	(0.222)	-1.119***	(0.204)	-1.116***	(0.178)
Indonesia/Malaysia	-0.925***	(0.046)	-0.810***	(0.103)	-0.871***	(0.099)	-1.074***	(0.084)

Each pair of columns shows own-price elasticities for palm oil by consumer market. Elasticities and estimation are as described in table B3. I explore dynamics with price lags and leads of up to three years. I exploit longer-run variation in the maximum and minimum shocks within rolling windows of up to five years before and after each year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I take a static, short-run approach to estimating demand. Estimation is static because the demand model is static, and estimation is short-run because it relies on annual variation in prices. I underestimate price elasticities if substitution is slow or if switching is costly. If substitution is slow, then consumers respond to price changes with a delay or in anticipation. If switching is costly, then consumers respond only to longer-run price changes. In both cases, demand responds weakly to contemporaneous prices. However, table B4 finds limited evidence for such effects. I evaluate delayed and anticipatory responses with price lags and leads of up to three years, and I draw

on longer-run variation over rolling windows that extend up to five years before and after each year. I obtain similar estimates across specifications.

I may also overestimate price elasticities. If consumers stockpile or if permanent switching is more costly than temporary switching, then demand responds especially strongly to contemporaneous prices. Table B4 does not reveal strong upward bias in the baseline estimates. Indeed, the USDA data record limited stockpiling that averages only 12% of consumption by volume (USDA 2025), relative to 342% and 188% in other contexts (Erdem et al. 2003, Hendel and Nevo 2006). These studies document weekly, individual-level stockpiling, which aggregates out in my annual, national-level measures of consumption.

C Supply

Model

I define choice-specific conditional value functions

$$\bar{v}(0, s_{it}) = \beta \mathbb{E}[\bar{V}(s_{it+1}) | s_{it}], \quad (12a)$$

$$\bar{v}(1, s_{it}) = -\bar{c}(s_{it}) + V(0, s_{it}). \quad (12b)$$

To lighten notation, I denote arguments with subscripts for the rest of this section. By the logit shocks of equation 5, mill construction probabilities are

$$\pi_{it} = \frac{\exp[\bar{v}_{it}(1)]}{\exp[\bar{v}_{it}(0)] + \exp[\bar{v}_{it}(1)]}, \quad (13)$$

and expected utility is given by the log-sum formula

$$\bar{V}_{it} = \ln\{\exp[\bar{v}_{it}(0)] + \exp[\bar{v}_{it}(1)]\} = \bar{v}_{it}(1) - \ln \pi_{it}. \quad (14)$$

Arcidiacono and Ellickson (2011) document this expression as the logit special case of Arcidiacono and Miller (2011) Lemma 1. For linear revenue function r_{it} , I can rewrite equation 6 as

$$V_{it}(N_{it}) = \alpha p_t y_{it} N_{it} + \mathbb{E}_{it} \left[\max_{n_{it}} \left\{ -c_{it}(n_{it}) + \sum_{s=3}^{T-t} \alpha \beta^s p_{t+s} y_{it+s} n_{it} \right\} + \beta V_{it+1}(N_{it}) \right]. \quad (15)$$

I express flow revenues for new development n_{it} as a net present value, while keeping revenues for mature plantations N_{it} in recursive form. For quadratic cost function c_{it} , the first order condition gives plantation development

$$n_{it} = \frac{1}{\psi} \left(-\gamma_g^0 - \gamma_g^1 t - x_i \delta - \varepsilon_{it} + \sum_{s=3}^{T-t} \alpha \beta^s \mathbb{E}_{it}[p_{t+s} y_{it+s}] \right). \quad (16)$$

Substituting back into equation 15, I obtain

$$V_{it}(0) = \frac{1}{2} \psi n_{it}^2 + \beta \mathbb{E}_{it}[V_{it+1}(0)] = \sum_{s=0}^{T-t} \beta^s \mathbb{E}_{it} \left[\frac{1}{2} \psi n_{it+s}^2 \right]. \quad (17)$$

Production

I can compute production by site from supply parameters $\{\alpha, \gamma_g^0, \gamma_g^1, \delta\}$ and states $\{s_{i1}, \dots, s_{iT}\}$. First, I calculate mill construction probabilities and plantation development $\{\pi_{it}, n_{it}\}$.

1. Compute $\{n_{i1}, \dots, n_{iT}\}$ with equation 16.
2. Compute $\{V_{i1}(0), \dots, V_{iT}(0)\}$ with equation 17.
3. Compute $\{\bar{v}_{i1}(1), \dots, \bar{v}_{iT}(1)\}$ with equation 12b.
4. Compute $\{\pi_{i1}, \dots, \pi_{iT}\}$ backward from terminal year T . In year T , $\bar{v}_{iT}(0) = 0$ by normalization; π_{iT} follows from equation 13. In prior years, $\bar{v}_{it}(0) = \beta \mathbb{E}_{it}[\bar{v}_{it+1}(1) - \ln \pi_{it+1}]$ by equations 12a and 14; π_{it} follows from equation 13.

Second, I calculate expected mill stocks and plantation acreage $\{\hat{M}_{it}, \hat{N}_{it}\}$.

$$\hat{M}_{it+1} = \hat{M}_{it} + (1 - \hat{M}_{it})\pi_{it}, \quad \hat{N}_{it+3} = \hat{N}_{it} + \hat{M}_{it+1}n_{it}, \quad (18)$$

where initial conditions $\{M_{i1}, N_{i1}\}$ are as observed in the data. For sites that contain a mill initially, $M_{i1} = 1$ implies $\pi_{it} = 0$ and thus $\hat{M}_{it} = 1$. For other sites, I obtain expected mill investment $\hat{M}_{it} \in [0, 1]$ as a function of investment probabilities π_{it} . These expected values are continuous, unlike the binary observed values of mill stocks $M_{it} \in \{0, 1\}$, and I sum over sites to compute aggregate production in each year. As in [Hopenhayn \(1992\)](#), atomistic firms and the law of large numbers imply that realized production is given simply by expected production, which equations 18 deliver in closed form. Without atomistic firms, calculating realized production is

more computationally intensive. It requires simulating discrete realizations of mill construction choices m_{it} with choice probabilities π_{it} , then integrating over the distribution of potential realizations.

Estimation

I derive regression equations for estimating the supply model. On the intensive margin, differentiating equation 6 gives first order conditions for plantation development (n_{it}, n_{it+1}) .

$$[n_{it}] \quad \frac{\partial c_{it}}{\partial n_{it}} = \beta^3 \mathbb{E}_{it} \left[\frac{\partial r_{it+3}}{\partial n_{it}} \right] + \beta^4 \mathbb{E}_{it} \left[\frac{\partial V_{it+4}}{\partial n_{it}} \right], \quad [n_{it+1}] \quad \beta \mathbb{E}_{it} \left[\frac{\partial c_{it+1}}{\partial n_{it+1}} \right] = \beta^4 \mathbb{E}_{it} \left[\frac{\partial V_{it+4}}{\partial n_{it+1}} \right].$$

By the envelope theorem, impacts on future actions are negligible. Differencing gives Euler equation

$$\frac{\partial c_{it}}{\partial n_{it}} - \beta \mathbb{E}_{it} \left[\frac{\partial c_{it+1}}{\partial n_{it+1}} \right] = \beta^3 \mathbb{E}_{it} \left[\frac{\partial r_{it+3}}{\partial n_{it}} \right],$$

which compares payoffs of plantation development in year t relative to year $t + 1$. Specializing with r_{it} and c_{it} , the difference in revenues is $\mathbb{E}_{it} \left[\frac{\partial r_{it+3}}{\partial n_{it}} \right] = \alpha \mathbb{E}_{it} [p_{t+3} y_{it+3}]$. The difference in costs is

$$\frac{\partial c_{it}}{\partial n_{it}} - \beta \mathbb{E}_{it} \left[\frac{\partial c_{it+1}}{\partial n_{it+1}} \right] = (1 - \beta)(\gamma_g^0 + \gamma_g^1 t + x_i \delta) - \beta \gamma_g^1 + \psi n_{it} + \varepsilon_{it} - \beta \mathbb{E}_{it} [\psi n_{it+1} - \varepsilon_{it+1}].$$

Defining structural error $\mu_{it} = -\frac{1}{\psi} \varepsilon_{it} + \frac{\beta}{\psi} \varepsilon_{it+1}$ and expectational error

$$\eta_{it} = \mathbb{E}_{it} \left[\beta n_{it+1} + \frac{\alpha \beta^3}{\psi} p_{t+3} y_{it+3} + \frac{\beta}{\psi} \varepsilon_{it+1} \right] - \beta n_{it+1} - \frac{\alpha \beta^3}{\psi} p_{t+3} y_{it+3} - \frac{\beta}{\psi} \varepsilon_{it+1},$$

I substitute and rewrite to obtain regression equation 9.

On the extensive margin, $\bar{v}_{it}(1) = -\bar{c}_{it} + V_{it}(0)$ by definition, and $\bar{v}_{it}(0) = \beta \mathbb{E}_{it} [\bar{v}_{it+1}(1) - \ln \pi_{it+1}]$ by equations 12a and 14. Inverting equation 13 and substituting gives

$$\begin{aligned} \ln \pi_{it} - \ln(1 - \pi_{it}) &= \bar{v}_{it}(1) - \bar{v}_{it}(0) \\ &= -\bar{c}_{it} + V_{it}(0) + \beta \mathbb{E}_{it} [\bar{c}_{it+1} - V_{it+1}(0) + \ln \pi_{it+1}], \end{aligned}$$

which compares payoffs of mill construction in year t relative to year $t + 1$. Special-

izing with r_{it} , c_{it} , and \bar{c}_{it} , the difference in revenues is $V_{it}(0) - \beta \mathbb{E}_{it}[V_{it+1}(0)] = \frac{1}{2}\psi n_{it}^2$ by equation 17. The linear separability of equation 17 admits a finite-dependence argument, as the only gain from mill construction in year t relative to $t+1$ is the flow revenue from plantation development n_{it} , net of its upfront costs. Future development choices and payoffs remain unaffected, such that continuation values align and difference out. The difference in costs is

$$-\bar{c}_{it} + \beta \mathbb{E}_{it}[\bar{c}_{it+1}] = \beta \bar{\gamma}_g^1 - (1 - \beta)(\bar{\gamma}_g^0 + \bar{\gamma}_g^1 t + x_i \bar{\delta}) - \bar{\varepsilon}_{it} + \beta \mathbb{E}_{it}[\bar{\varepsilon}_{it+1}].$$

Defining structural error $\bar{\mu}_{it} = -\bar{\varepsilon}_{it} + \beta \bar{\varepsilon}_{it+1}$ and expectational error

$$\bar{\eta}_{it} = \mathbb{E}_{it}[\beta \ln \pi_{it+1} + \beta \bar{\varepsilon}_{it+1}] - \beta \ln \pi_{it+1} - \beta \bar{\varepsilon}_{it+1},$$

I substitute and rewrite to obtain regression equation 10.

Finally, a static model would ignore the durability of mills and plantations, as well as time to build. To this end, I consider laws of motion $M_{it} = m_{it}$ and $N_{it} = n_{it}$ that eliminate long-run continuation values. Equations 6 and 5 yield intensive- and extensive-margin conditions

$$\frac{\partial c_{it}}{\partial n_{it}} = \frac{\partial r_{it}}{\partial n_{it}}, \quad \ln \pi_{it} - \ln(1 - \pi_{it}) = \bar{v}_{it}(1) - \bar{v}_{it}(0) = -\bar{c}_{it} + V_{it}(0).$$

I substitute and rewrite to obtain regression equations

$$n_{it} = \frac{\alpha}{\psi} p_t y_{it} - \frac{1}{\psi} (\gamma_g^0 + \gamma_g^1 t + x_i \delta) - \frac{1}{\psi} \varepsilon_{it},$$

$$\ln \pi_{it} - \ln(1 - \pi_{it}) = \frac{1}{2} \psi n_{it}^2 - \bar{\gamma}_g^0 - \bar{\gamma}_g^1 t - x_i \bar{\delta} - \bar{\varepsilon}_{it}.$$

Elasticities

For small price changes within the study period, I can compute production by site from supply parameters $\{\alpha, \psi\}$. I do not need to specify states beyond the study period. These price changes have no direct effect on forward-looking choices $\{\hat{n}_{iS}, \hat{\pi}_{iS}\}$, where S is the final year of the study period, and so I can read these choices from data $\{n_{iS}, \pi_{iS}\}$. I then compute counterfactual choices $\{\hat{n}_{it}, \hat{\pi}_{it}\}$ for all

Table C1: Lagged instruments

	Future revenues
	$\beta^3 p_{t+3} y_{it+3}$
Current total oil consumption \times yields	112.8*** (1.811)
Current other oil rainfall shocks \times yields	-31.19*** (1.143)
Current other oil temperature shocks \times yields	-81.66*** (1.405)
Region FE	x
Region-year trend	x
Observations	37,754
F-statistic	6,760

The table shows the first stage for current revenues as an instrument for future revenues. Other oils aggregate over coconut, olive, rapeseed, soybean, and sunflower oils. I control for cost factors and cluster standard errors by district. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

preceding years directly from estimating equations 9 and 10. For differences

$$\Delta n_{it} = n_{it} - \beta n_{it+1}, \quad (19a)$$

$$\Delta \pi_{it} = \ln \pi_{it} - \ln(1 - \pi_{it}) - \beta \ln \pi_{it+1}, \quad (19b)$$

I proceed as follows for price changes from observed p_t to counterfactual \hat{p}_t .

1. Compute Δn_{it} by equation 19a.
2. Compute $\Delta \hat{n}_{it}$. By equations 9 and 19a, $\Delta \hat{n}_{it} = \Delta n_{it} + \frac{\alpha \beta^3}{\psi} (\hat{p}_{t+3} - p_{t+3}) y_{it+3}$.
3. Compute \hat{n}_{it} . By equation 19a, $\hat{n}_{it} = \Delta \hat{n}_{it} + \beta n_{it+1}$, where \hat{n}_{iS} is data.
4. Compute $\Delta \pi_{it}$ by equation 19b.
5. Compute $\Delta \hat{\pi}_{it}$. By equations 10 and 19b, $\Delta \hat{\pi}_{it} = \Delta \pi_{it} + \frac{1}{2} \psi (\hat{n}_{it} - n_{it})^2$.
6. Compute $\hat{\pi}_{it}$. By equation 19b, $\ln \hat{\pi}_{it} - \ln(1 - \hat{\pi}_{it}) = \Delta \hat{\pi}_{it} + \beta \ln \hat{\pi}_{it+1}$, where $\hat{\pi}_{iS}$ is data.
7. Compute $\{\hat{M}_{it}, \hat{N}_{it}\}$ by equations 18.

Estimates

Table C1 presents the first stage regression and shows that lagged demand shifters are strong instruments for revenues. Total vegetable oil consumption raises demand and thus prices for palm oil. Rainfall and temperature shocks to other vegetable oils instead reduce prices for palm oil. The substitution effect suggests that these shocks reduce supply and raise prices for other oils, thereby raising demand and thus prices

Table C2: Supply parameters by region

	γ_g^0/α (\$1K)		γ_g^1/α (\$1K)		$\bar{\gamma}_g^0/\alpha$ (\$1M)		$\bar{\gamma}_g^1/\alpha$ (\$1M)	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Indonesia, Sumatra	6.647***	(0.218)	-0.471***	(0.021)	99.77**	(45.30)	1.713**	(0.677)
Indonesia, Kalimantan	8.146***	(1.038)	-0.166	(0.104)	81.89**	(35.47)	1.641***	(0.412)
Malaysia, Peninsular	10.44***	(1.015)	-0.284***	(0.095)	35.83**	(16.26)	-0.329	(0.654)
Malaysia, East	6.911***	(0.761)	-0.682***	(0.030)	144.2**	(70.82)	2.338***	(0.707)

Each row shows cost parameters in dollar terms for a producing region. The first pair of columns is the fixed cost of plantation development in thousands of inflation-adjusted, year-2000 USD. The second pair is the annual trend in these costs. Similarly, the third pair of columns is the fixed cost of mill construction in millions of dollars. The fourth pair is the annual trend in these costs. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


for palm oil. But the income effect suggests that higher prices for other oils reduce price-adjusted incomes, thereby reducing demand and thus prices for palm oil. The income effect seems to dominate.³²

Table C2 presents average costs and time trends by region. Peninsular Malaysia has the highest costs of plantation development, but the lowest costs of mill construction. This combination rationalizes widespread but slow growth in palm oil production, which originates in this region but remains limited by land constraints during the study period. Other regions have lower plantation costs and higher mill costs, which rationalize spatially concentrated but rapid growth in production. Low plantation costs encourage intensive-margin expansion, while high mill costs discourage extensive-margin expansion. These effects strengthen over time with falling plantation costs and rising mill costs.

Figure C1 shows that estimation without instruments leads to attenuated estimates, as expected. For intensive-margin equation 9, larger revenue $p_{t+3}y_{it+3}$ implies larger unobserved costs ε_{it} and smaller expectational error η_{it} , which attenuate the effect of an increase in $p_{t+3}y_{it+3}$ on the dependent variable. Static estimation also leads to biased estimates, which I find are negative. Static estimation regresses on current prices, even though forward-looking investment depends on future prices. Current prices are noisy measures of future prices, and this noise biases estimates toward zero. Furthermore, investment may even slow when prices are high, as mean

³² Table 3 suggests a related force. For other importers (44% of consumption), I estimate δ to be large and positive. Palm oil expenditure shares fall as incomes fall, adding to the income effect.

Figure C1: Supply elasticities, alternative specifications



Each point is an elasticity of total production during the study period with respect to a 1% increase in prices throughout the study period. The first three points show dynamic IV, dynamic OLS, and static IV estimation. The fourth point uses disaggregated road, port, and urban distances and above- and belowground carbon stocks for estimation. The “smoothing” points show alternative basis functions for smoothing in extensive-margin estimation. The “selection” points show alternative samples for smoothing, again for extensive-margin estimation, where I compute plantation development for sites without mills from varying subsamples of sites with a mill. I plot 95% confidence intervals.

reversion implies that high prices today portend lower prices tomorrow.

Figure C1 also shows robustness to alternative specifications. First, I treat cost factors parsimoniously, but I obtain similar estimates with disaggregated cost factors that include road, port, and urban distances and above- and belowground carbon stocks. Second, extensive-margin estimation involves smoothing over observed choices to compute mill construction probabilities π_{it} and plantation development n_{it} for sites without mills. I smooth spatially with one-knot cubic basis splines, but I obtain similar estimates when smoothing non-spatially (omitting latitude and longitude as basis variables), with three knots, and with three-knot restricted cubic splines.

Third, I assume that unobserved mill and plantation costs $\{\bar{\varepsilon}_{it}, \varepsilon_{it}\}$ are uncorrelated with each other. Correlation threatens estimates of parameter ψ in equation 10. Consider a positive correlation, as is natural. On one hand, sites with a mill are positively selected on plantation development n_{it} , and so I overestimate n_{it} for sites without mills. This bias leads me to underestimate ψ . On the other hand, plantation development n_{it} is positively correlated with structural error $\bar{\mu}_{it}$, as n_{it} contains ε_{it} and $\bar{\mu}_{it}$ contains $\bar{\varepsilon}_{it}$. This correlation leads me to overestimate ψ . Figure C1 shows

that these forces seem to be inconsequential in my sample. In particular, I identify sites that are likely to have extreme values for unobserved costs. Sites with a mill despite low probabilities π_{it} of mill construction must have low costs $\bar{\varepsilon}_{it}$. Correlated costs then imply low costs ε_{it} that encourage high development n_{it} . I calculate π_{it} for all sites, and I drop those with a mill despite low π_{it} . These sites are most selected. I smooth over the remaining sites with a mill, and I obtain \hat{n}_{it} for sites without mills. I then estimate elasticities as in baseline. I obtain similar estimates when dropping the 10, 25, and 50% of sites with lowest π_{it} . More generally, however, the potential for correlation across margins remains a challenge for discrete-continuous models.

D Counterfactuals

Welfare

I compute undiscounted total consumer surplus, producer surplus, government revenue, and emissions over the study period. For a given year t ,

$$\begin{aligned} CS_t &= \sum_k \left(X(p_{1t}, p_{2t}; u_{kt}^0) - X(p_{1t}^0, p_{2t}^0; u_{kt}^0) \right), \\ PS_t &= \sum_i \left(r(N_{it}, s_{it}) - \frac{1}{\alpha} [c(n_{it}, s_{it})M_{it+1} + \bar{c}(s_{it})\pi_{it}(1 - M_{it})] \right), \\ G_t &= \sum_k p_{1t}q_{1kt}^D\tau_{kt}^D + \sum_i p_{1t}q_{1it}^S\tau_{it}^S, \quad E_t = \sum_i e_i n_{it} M_{it+1}. \end{aligned}$$

Consumer surplus is the increase in expenditures X_{kt} needed to maintain utility $u_{kt} = (\ln X_{kt} - \ln P_{kt})(\prod_o p_{ot}^{\delta_{ok}})^{-1}$, as derived in [Deaton and Muellbauer \(1980\)](#). It is relative to baseline prices $\{p_{1t}^0, p_{2t}^0\}$ and utility u_{kt}^0 . Producer surplus is relative to the outside option. In net-present-value terms, producer surplus is simply $\bar{V}(s_{it})$. I compute revenue net of costs to obtain producer surplus in undiscounted terms. Government revenue is from ad valorem taxes on world prices p_{1t} . Emissions depend on carbon stock density e_i , as measured in tons per hectare. I observe carbon stocks, and so I can read counterfactual emissions directly from data.

Emissions

On the demand side, I ignore emissions from substitution to other vegetable oils. The primary threat is South American soybean oil, which contributes to Amazonian

Table D1: Acacia vs. palm development

	All sites		1990 sites		All districts		1990 districts	
	Acacia	Acacia	Acacia	Acacia	Acacia	Acacia	Acacia	Acacia
Palm	0.0235*** (0.00888)	0.0134 (0.00859)	0.0233* (0.0116)	0.0125 (0.0124)	0.0350** (0.0159)	-0.0160 (0.0345)	0.0293 (0.0191)	-0.0189 (0.0422)
Site FE		x		x		x		x
Year FE		x		x		x		x
Observations	5,700	5,700	1,254	1,254	528	528	270	270

Each column is one regression. I measure palm and acacia development in hectares of new plantation. In the first four columns, each observation is a site-year. I consider the full sample and the sample with nonzero palm development in 1990. In the last four columns, each observations is a district-year. I again consider the full and 1990 samples. Standard errors are clustered by district. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

deforestation. The resulting bias is small because Amazonian deforestation is driven primarily by cattle, not soy ([Souza-Rodrigues 2019](#)), and it does not destroy peatlands, which are located away from the deforested outskirts of the forest ([Gumbrecht et al. 2017](#), [Song et al. 2018](#)). Furthermore, South American soybean oil is only 13% of total oil consumption. To capture these emissions, I would need to model the supply of soybean oil. I would then impose joint tariffs on palm and soybean oils.

On the supply side, I ignore emissions from substitution to other drivers of deforestation. I leave aside mining, which is governed by the exogenous distribution of deposits, and selective logging, which does not destroy peatlands. The primary threat that remains is substitution to acacia (paper pulp) plantations, which do destroy peatlands. The resulting bias is small because palm oil is seven times more profitable than acacia, which requires replanting upon harvest, such that switching to acacia is unappealing for many palm oil producers ([Sofiyuddin et al. 2012](#)). To capture these emissions, I would need to model the choice between palm and acacia development. I would then impose joint tariffs on palm oil and acacia.

Indeed, palm development greatly exceeds acacia development, and I find limited substitution between the two historically. [Gaveau et al. \(2019\)](#) measure palm and acacia plantations for the island of Borneo in five-year intervals from 1990 to 2015. For the average site, I observe 257 ha of palm development per year, relative to only 42 ha of acacia development per year. These measures align with the baseline data, which capture 288 ha of palm development per year. I test for substitution between

Table D2: Domestic regulation

	20% tax		33% tax		50% tax		67% tax	
	2016	2001	2016	2001	2016	2001	2016	2001
Production taxes (I + M)								
Emissions	-3	-0	-5	-0	-7	-1	-11	-1
Welfare: European Union	-7	-2	-12	-3	-22	-6	-33	-8
Welfare: China, India	-12	-3	-23	-5	-41	-9	-64	-13
Welfare: Other importers	-23	-8	-42	-14	-74	-22	-111	-31
Welfare: Indonesia	5	-1	5	-3	-0	-7	-11	-12
Welfare: Malaysia	14	5	24	9	38	14	49	17
Production taxes (I)								
Emissions	-2	-0	-3	-1	-5	-1	-6	-1
Welfare: European Union	-3	-1	-5	-1	-7	-2	-8	-2
Welfare: China, India	-6	-2	-10	-2	-13	-3	-15	-3
Welfare: Other importers	-12	-4	-18	-6	-22	-7	-26	-7
Welfare: Indonesia	-24	-16	-37	-27	-43	-32	-49	-32
Welfare: Malaysia	33	18	49	28	57	32	64	33
Production taxes (M)								
Emissions	-1	0	-1	-0	-2	-0	-2	-0
Welfare: European Union	-2	-1	-3	-1	-4	-1	-4	-1
Welfare: China, India	-4	-1	-6	-2	-7	-2	-8	-2
Welfare: Other importers	-7	-3	-10	-4	-13	-5	-14	-5
Welfare: Indonesia	25	15	33	22	37	24	39	25
Welfare: Malaysia	-20	-13	-27	-21	-30	-23	-32	-24

I compute total changes in global emissions and market-specific welfare from 1988 to 2016. Emissions are in gigatons of CO₂, and welfare is in billions of inflation-adjusted, year-2000 USD. Production taxes are levied in Indonesia (I), Malaysia (M), or both (I + M). Welfare for Indonesia and Malaysia includes consumer surplus, producer surplus, and government revenue. Welfare elsewhere includes consumer surplus. Columns are taxes of different levels, upheld from 1988 to 2016 or 2001.

palm and acacia as follows. For sites i and years t , I compare palm and acacia development with the specification

$$\Delta \text{acacia}_{it} = \beta \Delta \text{palm}_{it} + \alpha_i + \alpha_t + \varepsilon_{it}.$$

Acacia_{it} is total acacia plantation, and so $\Delta \text{acacia}_{it}$ is new acacia development. I include site and year fixed effects, and I cluster standard errors by district. Table D1 shows that palm development has very small effects on acacia development. Palm development does not displace acacia development and if anything slightly increases it, perhaps by opening up new lands. That is, palm and acacia are not substitutes and may even be weak complements. I isolate intensive-margin investments by focusing

Table D3: Trade policy

	25% tax		50% tax		100% tax		200% tax	
	2016	2001	2016	2001	2016	2001	2016	2001
Import tariffs (all imports)								
Emissions	-2	-0	-3	-0	-5	-0	-8	-1
Welfare: European Union	5	1	8	2	9	2	4	2
Welfare: China, India	7	1	10	1	9	1	-3	-0
Welfare: Other importers	13	4	17	6	13	6	-12	1
Welfare: Indonesia, Malaysia	-34	-15	-58	-25	-88	-39	-120	-54
Import tariffs (EU + CI)								
Emissions	-1	-0	-1	-0	-2	-0	-3	-0
Welfare: European Union	2	1	2	1	-2	-0	-14	-3
Welfare: China, India	3	0	1	0	-7	-1	-33	-4
Welfare: Other importers	9	2	16	5	27	8	41	12
Welfare: Indonesia, Malaysia	-12	-4	-20	-7	-32	-11	-45	-16
Import tariffs (EU)								
Emissions	-0	-0	-0	-0	-1	-0	-1	-0
Welfare: European Union	-0	0	-2	-0	-10	-2	-28	-5
Welfare: China, India	2	0	3	1	5	1	8	1
Welfare: Other importers	3	1	6	2	10	4	15	6
Welfare: Indonesia, Malaysia	-5	-2	-9	-4	-14	-6	-19	-9
Export taxes (all exports)								
Emissions	-2	-0	-3	-0	-5	-0	-8	-1
Welfare: European Union	-8	-2	-15	-4	-28	-7	-48	-11
Welfare: China, India	-14	-3	-28	-6	-51	-10	-90	-17
Welfare: Other importers	-28	-9	-52	-16	-94	-29	-161	-48
Welfare: Indonesia, Malaysia	35	12	61	20	98	33	145	50

I compute total changes in global emissions and market-specific welfare from 1988 to 2016. Emissions are in gigatons of CO₂, and welfare is in billions of inflation-adjusted, year-2000 USD. Import tariffs are levied on all imports, imports to the EU, China, and India (EU + CI), or imports to the EU alone (EU). Export taxes are levied on all exports from Indonesia and Malaysia. Welfare for Indonesia and Malaysia includes consumer surplus and producer surplus, as well as government revenue from export taxes. Welfare elsewhere includes consumer surplus, as well as government revenue from import tariffs. Columns are taxes of different levels, upheld from 1988 to 2016 or 2001.

on sites with nonzero initial palm development, and I allow for cross-site effects by aggregating to the district level. Each gives similar estimates.

Policy

Tables D2 and D3 summarize the impacts of domestic regulation and trade policy on global emissions and welfare by market. They correspond to table 6 and figures 6, 7, and 8.