

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

Allan Hsiao *

University of Chicago

January 4, 2022

([most recent version here](#))

Weak environmental regulation has global consequences. When domestic regulation of carbon-intensive industries fails, the international community can intervene by targeting these industries with import tariffs. I argue that import tariffs must possess two features – coordination and commitment – in order to be effective. Without coordination across importers, tariffs are undermined by leakage to unregulated markets. Without commitment to upholding tariffs over the long term, tariffs are reduced over time as importers give in to static incentives. I develop a dynamic empirical framework for quantifying these forces in settings with incomplete regulation and sunk investment, and I apply it to the market for palm oil, a major driver of deforestation and one of the largest sources of emissions globally.

*Email: ajhsiao@uchicago.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 David Atkin, Clare Balboni, Abhijit Banerjee, Dave Donaldson, Namrata Kala, Jing Li, Rachael Meager, Will Rafey, 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, domestic regulation is often constrained by low incentives from free riding and political constraints (Oates and Portney 2003), as well as implementation barriers from administrative limits and potential corruption (Burgess et al. 2012; Oliva 2015). The conventional approach attempts to address these challenges, such as by improving enforcement (Duflo et al. 2018), but doing so at scale can be infeasible. Trade policy offers an alternative for regulating the 60% of global CO₂ emissions embodied in traded goods (Davis et al. 2011). In particular, import tariffs circumvent domestic obstacles to regulation by directly targeting the prices emitters receive in world markets.

How effective are international import tariffs as a substitute for domestic regulation? This paper develops a dynamic empirical framework to answer this question quantitatively. I apply the framework to study the Indonesian and Malaysian palm oil industry, which accounts for a staggering 5% of global CO₂ emissions from 1990 to 2016 – more than the entire Indian economy (figure 1). I find that well designed import tariffs can be an effective substitute for a domestic palm oil tax, but face two significant challenges: a leakage problem under incomplete regulation, and a commitment problem from static incentives to reduce tariffs over time.

I begin by discussing the leakage and commitment problems. First, when importers do not coordinate, incomplete regulation leads to demand-side “leakage” (Fowlie 2009). That is, although tariffs lower consumption in regulated markets, in doing so they lower world prices and encourage consumption in unregulated markets. This offsetting effect constrains the size of tariffs, as large tariffs lead to large leakage and therefore low net benefits. Second, importers face a commitment problem. Most traded emissions are from industries in which sunk investments make up the bulk of production costs: fossil fuels, manufacturing, mining, transportation, and agriculture (Peters et al. 2011). The result is a static incentive to reduce tariffs over time: when investments are sunk, so too are emissions. For agriculture, emissions are sunk because they are released upon investment. Once land is cleared, the forest is gone. For other sectors, emissions are often sunk, even if released gradually, because investment ensures low marginal costs up to capacity. Once an oil well has been

Figure 1: CO₂ emissions from palm oil plantations over time




Figure 1a computes palm emissions using data on palm oil plantations (Xu et al. 2020; Song et al. 2018), tree biomass (Zarin et al. 2016), and peat deposits (Gumbricht et al. 2017). Figure 1b compares palm emissions to CO₂ emissions for the top seven emitters from 1990 to 2016. Palm emissions account for 4.95% of global emissions during this period. Global data come from the World Resources Institute and Global Carbon Atlas and include land-use change.

identified, explored, and drilled, extraction is cheap and thus proceeds to completion.

Palm oil and the resulting deforestation offer an ideal setting for studying environmental regulation by trade policy. I focus on palm oil from Indonesia and Malaysia, which together produce 84% of global supply. First, the industry is a major polluter. Land clearing for palm oil plantations threatens peatland forests that are particularly carbon-rich. Second, domestic incentives to regulate are limited. Despite its global consequences, palm oil is a major source of export revenue and has lifted millions out of poverty (Edwards 2019). 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). Third, foreign governments are actively discussing trade-policy interventions, with the EU passing recent legislation targeting palm oil imports (OJEU 2018). Fourth, satellite imagery provides a rich source of spatial data capturing the evolution of the industry over time and at a granular level.

I build a quantitative empirical model for evaluating palm oil import tariffs. I divide land into individual sites, which I treat as firms representing potential entrants. Firms deforest land for plantations, plantations produce fruit for mills, mills process fruit into palm oil for domestic and foreign consumers, and foreign consumers in

regulated markets pay import tariffs. The leakage problem depends on unregulated markets' elasticity of palm oil demand, which in turn depends on consumers' substitution between palm and other vegetable oils. The commitment problem depends on short-versus long-run elasticities of palm oil supply, which in turn depend on producers' expectations over future prices. The value of the structural model is that it accounts explicitly for cross-oil substitution on the demand side and price expectations on the supply side. A more reduced-form approach – that is, regressing palm oil demand and supply on prices (with instruments) – would account for neither, resulting in biased elasticity estimates in addition to ignoring equilibrium effects.

I model palm oil demand by consumer market with an almost ideal demand system in which consumers choose between palm and other vegetable oils ([Deaton and Muellbauer 1980](#)). For estimation, I apply the iterated linear least squares approach of [Blundell and Robin \(1999\)](#) using annual panel data on vegetable oil prices and consumption by country. I address price endogeneity by instrumenting with weather shocks to oil production, which shift supply. I then obtain world palm oil demand curves, which shift right over time as demand expands.

I model palm oil supply with a dynamic model of sunk investment in palm oil production. Forward-looking firms invest along two margins. On the extensive margin, firms make a discrete choice over whether to build mills – a prerequisite for plantations. On the intensive margin, firms with mills make a continuous choice over how much land to deforest and develop into plantations. Once operational, plantations produce palm oil each period and thus yield a stream of revenues. These revenues depend on world prices, which in turn depend on aggregate production. Firms therefore play a dynamic competitive equilibrium as in the entry and investment game of [Hopenhayn \(1992\)](#). The dynamic model allows me to infer firms' hypothetical responses to long-run tariffs from their observed responses to short-run price variation, while accounting for price expectations in a disciplined way.

I take an Euler approach for estimating the supply model, combining standard continuous Euler methods for the intensive margin with more recent discrete Euler methods for the extensive margin ([Hall 1978](#); [Scott 2013](#)). In both cases, I compare investing today versus tomorrow: investing today brings forward plantation revenues, but it also brings forward investment costs. On the intensive margin, I form an Euler equation from the first order condition for investment. On the extensive margin, I use

discrete, short-term perturbations that hold long-term investments fixed. Continuation values difference out, greatly simplifying computation, and estimation reduces to linear regression with instruments. Identification comes from two sources: exogenous variation in world palm oil prices over time, given shifts in the world demand curve estimated above, and exogenous variation in palm oil yields over space, given differences in climate. Intuitively, high prices raise revenues most for high-yield plantations.

For counterfactuals, I solve by backward induction from the steady state. The model assumes no exit and thus reaches a steady state when land is exhausted. The challenge is that reaching this point takes many periods, and backward induction over long horizons suffers from a curse of dimensionality. I address this computational difficulty by iterating on two dimensions. In the outer loop, I solve over a manageable horizon treating the final period as the steady state. I then improve the solution by solving over a longer horizon, and I repeat until convergence. In the inner loop, I backward induct iteratively with a limited look-ahead window. I quantify emissions by combining the model’s spatial predictions with a map of carbon stocks.

I evaluate how coordination and commitment affect import tariffs, both individually and in combination. I consider tariffs that maximize social welfare and treat all palm oil uniformly as a baseline, then I relax each condition. When coordination and commitment hold, import tariffs are an effective substitute for domestic regulation by Indonesia and Malaysia. Tariffs reduce carbon emissions by 56% relative to business as usual, while the domestic tax reduces emissions by 64%. Import tariffs miss domestic consumption by Indonesia and Malaysia, but the loss is modest because inelastic domestic demand limits leakage on this margin.

Both coordination and commitment are necessary. When either fails, import tariffs are low and have little effect. Even under full commitment, emission reductions fall from 56% under full coordination, to 17% under an EU-China-India coalition, to 2% under unilateral EU action because of elastic importer demand. Emission reductions fall disproportionately more than tariff coverage – 80%, 35%, and 12% of world consumption, respectively – because leakage concerns lead to smaller tariffs. Even under full coordination, emission reductions fall from 56% under full commitment to 0% under no commitment because of time to build. Without commitment, imposing tariffs is never statically optimal: tariffs affect neither past development, which is sunk, nor new development, which does not yet generate taxable production.

Furthermore, coordination and commitment interact. I consider limited commitment in which importers commit to finite tariff regimes – e.g., “five-year plans” – and revise tariffs at the end of each regime. Achieving 95% of full-commitment emission reductions requires a commitment period of only five years when importers coordinate, but more than twenty years when the EU acts unilaterally. Thus, coordination lessens the burden of commitment. The mechanism is that producers can deforest and sell to unregulated markets, then shift sales to regulated markets as commitment failures lead to tariff reductions. The worse the leakage problem, the more unregulated markets can absorb, and the more easily producers can skirt tariffs.

Finally, the division of surplus among countries reveals why coordination and commitment are difficult to achieve in practice, particularly for countries that do not value emission reductions. Coordination is difficult because joining the coalition reduces consumer surplus, while defectors enjoy the lower world prices induced by the coalition. Commitment is similarly difficult, as longer commitment sacrifices more consumer surplus. Moreover, all tariff scenarios lead to lower surplus for Indonesia and Malaysia, which may raise equity concerns. I use the model to assess each of these challenges and quantify the transfers needed to address them.

This paper develops a new dynamic empirical framework for assessing emission-based trade policy. In doing so, I build on a rich literature studying environmental regulation in trade-exposed markets, where leakage motivates border adjustment taxes ([Markusen 1975](#); [Copeland and Taylor 1994, 1995](#); [Hoel 1996](#); [Rauscher 1997](#); [Elliott et al. 2010](#); [Fowlie et al. 2016](#); [Kortum and Weisbach 2017, 2021](#)) and carbon coalitions ([Nordhaus 2015](#); [Böhringer et al. 2016](#); [Farrokhi and Lashkaripour 2021](#)), or where trade policy distorts environmental incentives more broadly ([Shapiro 2020](#)). 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 2016, 2020](#); [Battaglini and Harstad 2016](#); [Acemoglu and Rafey 2019](#)). I provide novel analysis of how leakage and commitment interact, and I quantify these challenges empirically. By focusing on one important industry, I capture the rich dynamics and spatial heterogeneity often missed by the computable general equilibrium models most common in the trade policy literature.

Methodologically, I build on dynamic models of industry dynamics in the tradition of [Hopenhayn \(1992\)](#) and [Ericson and Pakes \(1995\)](#), with empirical applications

including [Ryan \(2012\)](#) and [Collard-Wexler \(2013\)](#). I draw on a growing literature, formalized by [Aguirregabiria and Magesan \(2013\)](#), [Scott \(2013\)](#), and [Kalouptsi et al. \(2021\)](#), that develops Euler conditional choice probability methods for estimating dynamic discrete choice models. Using standard dynamic discrete choice 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. My contribution is to show how to combine both continuous and discrete Euler techniques in a single framework, with a generalizable model containing discrete entry choices on the extensive margin and continuous investment choices on the intensive margin.

More broadly, trade policy enables regulation of industries operating 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 ([Burgess et al. 2019](#); [Souza-Rodrigues 2019](#); [Araujo et al. 2021](#); [Assunção et al. 2021](#)) or conservation contracting ([Harstad 2012, 2016](#); [Harstad and Mideksa 2017](#)), and it scales readily, unlike direct payments for ecosystem services ([Jayachandran et al. 2017](#); [Edwards et al. 2020](#)). [Domínguez-Iino \(2021\)](#) and [Harstad \(2021\)](#) also study trade policy and deforestation, but in static and theoretical settings, respectively. This paper shows that coordination and dynamic concerns both matter empirically for trade policy. And it quantifies these challenges in an industry that is pivotal in the fight against climate change.

2 Illustrative Model

This section studies optimal tariffs for an emission-intensive traded good in a setting with incomplete regulation and sunk investment.

2.1 Import tariffs, leakage, and commitment

Consider an unregulated “domestic” market u and a regulated “foreign” market r . I study an agricultural good produced in u and consumed in both u and r . Consumers have consumption utility described by inverse demand curves $P_t^{Dr}(q)$ and $P_t^{Du}(q)$. Price-taking farmers produce the good by establishing plantations, subject

to upfront development costs described by inverse supply curve $P_t^S(q)$. Investment in plantations is sunk and causes upfront emissions e via deforestation. Beginning one period after development, plantations produce goods every period at zero marginal cost. Social welfare is consumer and producer surplus net of emission damages over time. It depends on old development $Q_t^o = Q_{t-1}^o + Q_{t-1}^n$, the path of new development $\{Q_t^n, Q_{t+1}^n, \dots\}$ for $Q_t^n = Q_t^{rn} + Q_t^{un}$, and how production is allocated across markets.

$$W_t(Q_t^{rn}, Q_{t+1}^{rn}, \dots, Q_t^{un}, Q_{t+1}^{un}, \dots; Q_t^o) \\ = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t \left[\int_0^{Q_{t+s}^{ro}} P_{t+s}^{Dr}(q) dq + \int_0^{Q_{t+s}^{uo}} P_{t+s}^{Du}(q) dq - \int_{Q_{t+s}^o}^{Q_{t+s}^o + Q_{t+s}^n} \left(P_{t+s}^S(q) + e \right) dq \right].$$

Domestic regulation

The first best is a domestic Pigouvian tax that captures the full externality.

$$\tilde{\tau}_t^{\text{FB}} = e,$$

where the tilde denotes net present value. There is no leakage problem because direct domestic regulation of supply achieves complete regulation. There is no commitment problem because the regulator can impose the full tax upfront with a license fee.

The leakage problem

Import tariffs target regulated consumption but miss unregulated consumption. Such regulation is incomplete. Under full commitment, the optimal tariff is

$$\tilde{\tau}_t^{\text{C}} = \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du}} \right) e < \tilde{\tau}_t^{\text{FB}},$$

where $\varepsilon_t^S > 0$ and $\varepsilon_{t+1}^{Du} < 0$ are elasticities of supply and unregulated demand. The tariff is smaller than the first-best tax. First, leakage lowers the benefits of the tariff. Although tariffs decrease regulated consumption, these benefits are partially offset by unregulated consumption that expands as tariffs lower world prices. Second, leakage raises the costs of the tariff. Tariffs produce allocative inefficiency as they shift consumption from regulated consumers with higher willingness to pay to unregulated consumers with lower willingness to pay.

The commitment problem

Import tariffs tax consumption – not development directly – and thus are applied over time. But sunk investment, time to build, and leakage together induce a commitment problem. Tariffs have no benefit today: they cannot prevent prior development, which is sunk, and they cannot prevent new development, which under time to build does not generate taxable production until a future period. Furthermore, tariffs are costly: leakage leads to allocative inefficiency as tariffs distort consumption between markets. In combination, these forces make it statically optimal to set tariffs to zero. Without commitment, importers never levy tariffs at all.

Under limited commitment, importers uphold tariffs for L periods at a time before resetting them. Importers cycle between removing and imposing tariffs as they give in to static incentives when tariffs are reset. Tariffs have net present value

$$\tilde{\tau}_t^{\text{LC}}(L) = \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du} [1 + \Lambda(L, \varepsilon)]} \right) e,$$

for $\Lambda(L, \varepsilon) = \frac{(1-\beta)\beta^L}{\beta-\beta^L} \frac{P_{t+L}^{Du}}{P_{t+1}^{Du}} \left(1 - \frac{Q_{t+1}^o \varepsilon_t^S}{Q_{t+L}^{ro} \varepsilon_{t+1}^{Dr} + Q_{t+L}^{uo} \varepsilon_{t+1}^{Du}} \right) > 0$. Tariffs are increasing in L and approach full commitment as $L \rightarrow \infty$.

$$0 = \tilde{\tau}_t^{\text{NC}} < \tilde{\tau}_t^{\text{LC}}(L) < \tilde{\tau}_t^{\text{C}} = \lim_{L \rightarrow \infty} \tilde{\tau}_t^{\text{LC}}(L)$$

Moreover, tariffs $\tilde{\tau}_t^{\text{LC}}$ fall over time because they do less to reduce emissions as the stock of sunk investment grows. At the extreme, tariffs are set to zero once all forest is cleared, at which point tariffs no longer reduce emissions (and $\varepsilon_t^S = 0$). Also note that while emissions from deforestation are released upfront, the same framework applies to emissions released over time if sunk investment in brown technology leads to permanently low marginal costs of production. In this case, investment locks in future production and emissions, and externality e captures the net present value.

How leakage and commitment interact

Incomplete regulation allows producers to skirt high tariffs in any given period by directing sales to the unregulated market until tariffs fall. Leakage therefore exacerbates the commitment problem. The greater the leakage problem, the more

Table 1: Palm oil statistics by country (1988-2016)

	Production	Consumption	Exports	Imports
Indonesia	0.44	0.14	0.41	0.00
Malaysia	0.40	0.06	0.48	0.02
European Union	0.00	0.12	0.00	0.17
China	0.00	0.11	0.00	0.15
India	0.00	0.12	0.00	0.16
Rest of world	0.16	0.45	0.10	0.50

Data are from the USDA Foreign Agricultural Service. Columns show ratios of global totals.

the unregulated market can absorb, and thus the greater the loss from commitment failures. At the extreme, eliminating leakage also solves the commitment problem. Commitment failures lead to low future tariffs, but the regulator can compensate by frontloading high tariffs if, absent leakage, producers are unable to skirt them.

3 Empirical Setting and Data

This section provides institutional details and describes the data. Both make the palm oil market ideal for studying environmental regulation by trade policy.

3.1 Empirical setting

Palm oil is among the most widely used plant products in the world, with a low price point driven by high yields relative to other oilseeds. Palm oil is used as a cooking oil, as a replacement for trans fats in processed foods, and as an ingredient in non-food products ranging from soaps to cosmetics to biofuels. Table 1 shows that Indonesia and Malaysia account for 84% of global production, 90% of exports, and 20% of consumption; the European Union, China, and India account for another 35% of global consumption. The market is unconcentrated: the largest producer (FGV Holdings Berhad) is 4% of global production (POA 2017), and the largest consumer (Unilever) is 2% of global consumption (WWF 2016).

This empirical setting is appealing for several reasons. First, palm oil is among the largest sources of global carbon emissions. Deforestation for palm oil plantations has severe consequences because Indonesia and Malaysia are rich in peatland forests,

which contain deep layers of carbon-rich peat. I compute palm-related emissions in figure 1a and find that emissions from peat deposits exceed those from tree biomass by five to ten times.¹ Figure 1b shows that palm emissions account for more CO₂ from 1990 to 2016 than the entire Indian economy.

Second, there are significant challenges in implementing regulation domestically. Free-riding limits incentives to pass regulation, and weak enforcement hampers regulation that does pass. In 2010, Norway pledged US \$1 billion to Indonesia in cash incentives for domestic forest regulation, prompting a 2011 moratorium on new forest concessions. But the moratorium had little effect, failing to curb both (legal) deforestation within existing concessions and (illegal) deforestation outside of concessions, including in protected areas (Busch et al. 2015).

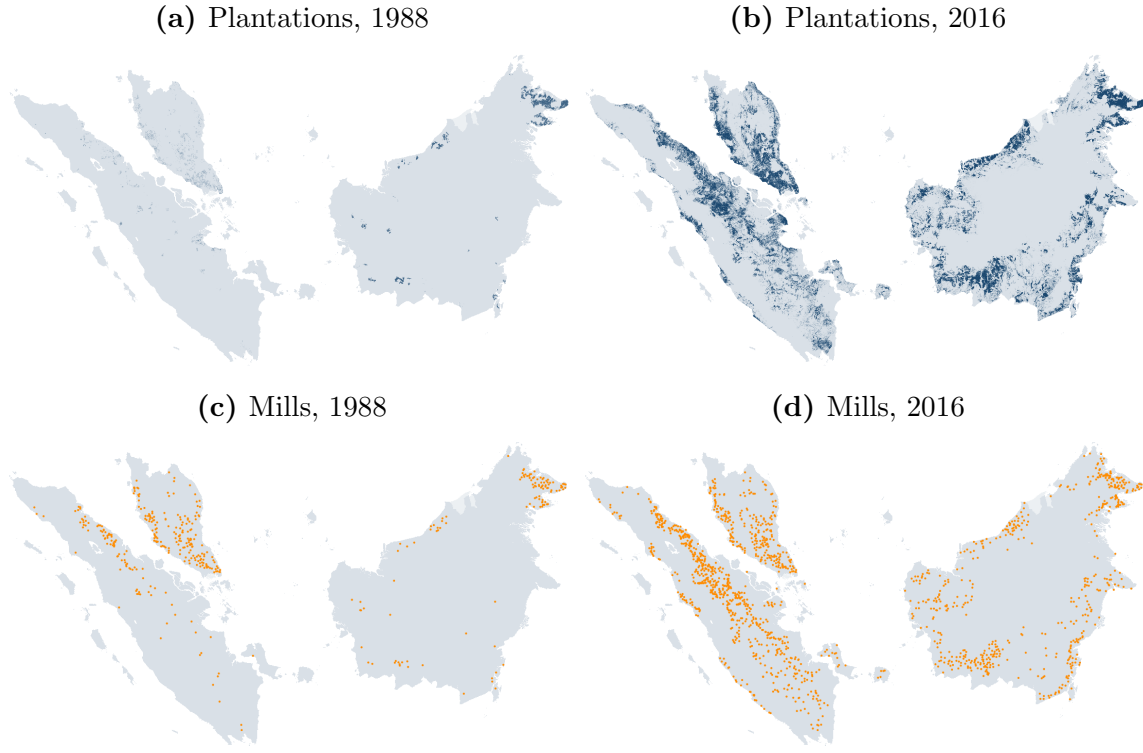
Third, European policymakers are actively discussing trade-policy interventions. French parliament debated a “Nutella” tax on palm oil in food products in 2016. The European Union is set to eliminate green subsidies for palm-based biofuels, cap production, and achieve a complete phase-out by 2030. Palm-based biofuels face the further loss of green tax incentives in France and an outright ban in Norway. As with import tariffs, each policy uses European buying power to reduce emissions abroad.

3.2 Data

Spatial panel data record palm oil plantations and mills from 1988 to 2016 using satellite imagery at a resolution of 30 arc-seconds – approximately 1 km². Figure 2 maps the expansion of production over this period. For plantations, Xu et al. (2020) analyze PALSAR and MODIS satellite data to capture plantation development from 2001 to 2016. Using data on tree cover loss from 1988 to 2016 from Song et al. (2018), who draw on Landsat and MODIS satellite data, I estimate the (positive) relationship between plantation development and tree cover loss, and I use this relationship to impute plantation development back to 1988. For mills, I rely on geocoded data on present-day mills from the World Resources Institute and the Center for International Forestry Research, and I use historical satellite data to manually identify construction dates back to 1988. The Indonesian data focus on Sumatra, Kalimantan, and Riau

¹ 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.

Figure 2: Palm oil plantations and mills over time

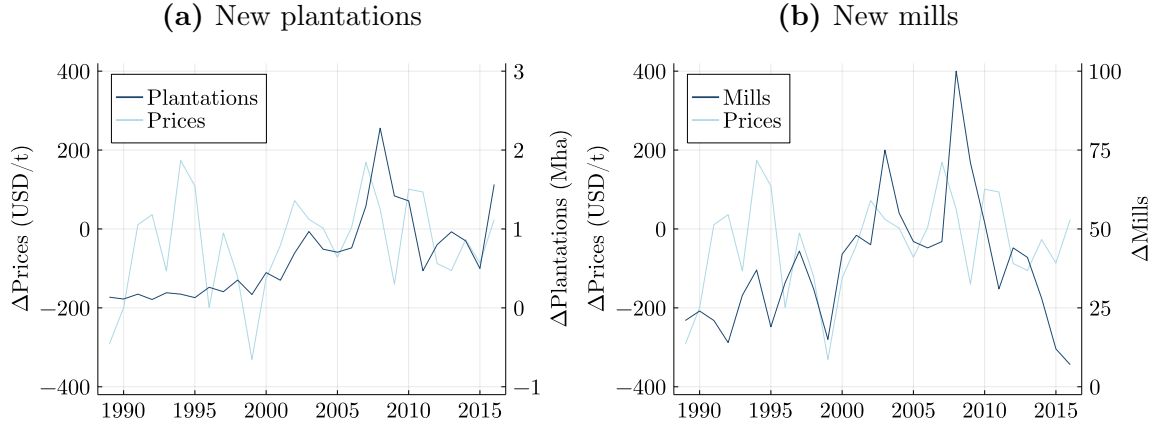


Data on plantations come from [Xu et al. \(2020\)](#) and [Song et al. \(2018\)](#), and data on mills come from the World Resources Institute and the Center for International Forestry Research. The study area is Sumatra, Kalimantan, and Riau of Indonesia and all of Malaysia, covering virtually all palm production in Indonesia and Malaysia.

but remain exhaustive, covering 97% of mills. I compare my measures of plantations and mills to aggregate government statistics and find that they align closely. Figure 3 compares investment in plantations and mills to fluctuations in world prices over time, with world price data from the International Monetary Fund and World Bank.

Figure 4 maps land characteristics, which I measure at a resolution of 30 arc-seconds. I use an agronomic model of the oil palm plant ([Hoffmann et al. 2014](#)) to compute potential palm oil yields as a function of climate. These potential yields are time-invariant but computable at high resolution, allowing me to downscale data on actual yields over time from provincial government statistics. Euclidean distances to the nearest major port, road, and urban district generate spatial heterogeneity via transport costs. I compute carbon stocks from geospatial data on tree biomass and peat deposits ([Zarin et al. 2016](#); [Gumbrecht et al. 2017](#)), which record how much

Figure 3: Palm oil production vs. world prices over time



Data on plantation development come from [Xu et al. \(2020\)](#) and [Song et al. \(2018\)](#), and data on mill construction from the Universal Mill List. Prices combine palm and palm kernel oil prices from the International Monetary Fund.

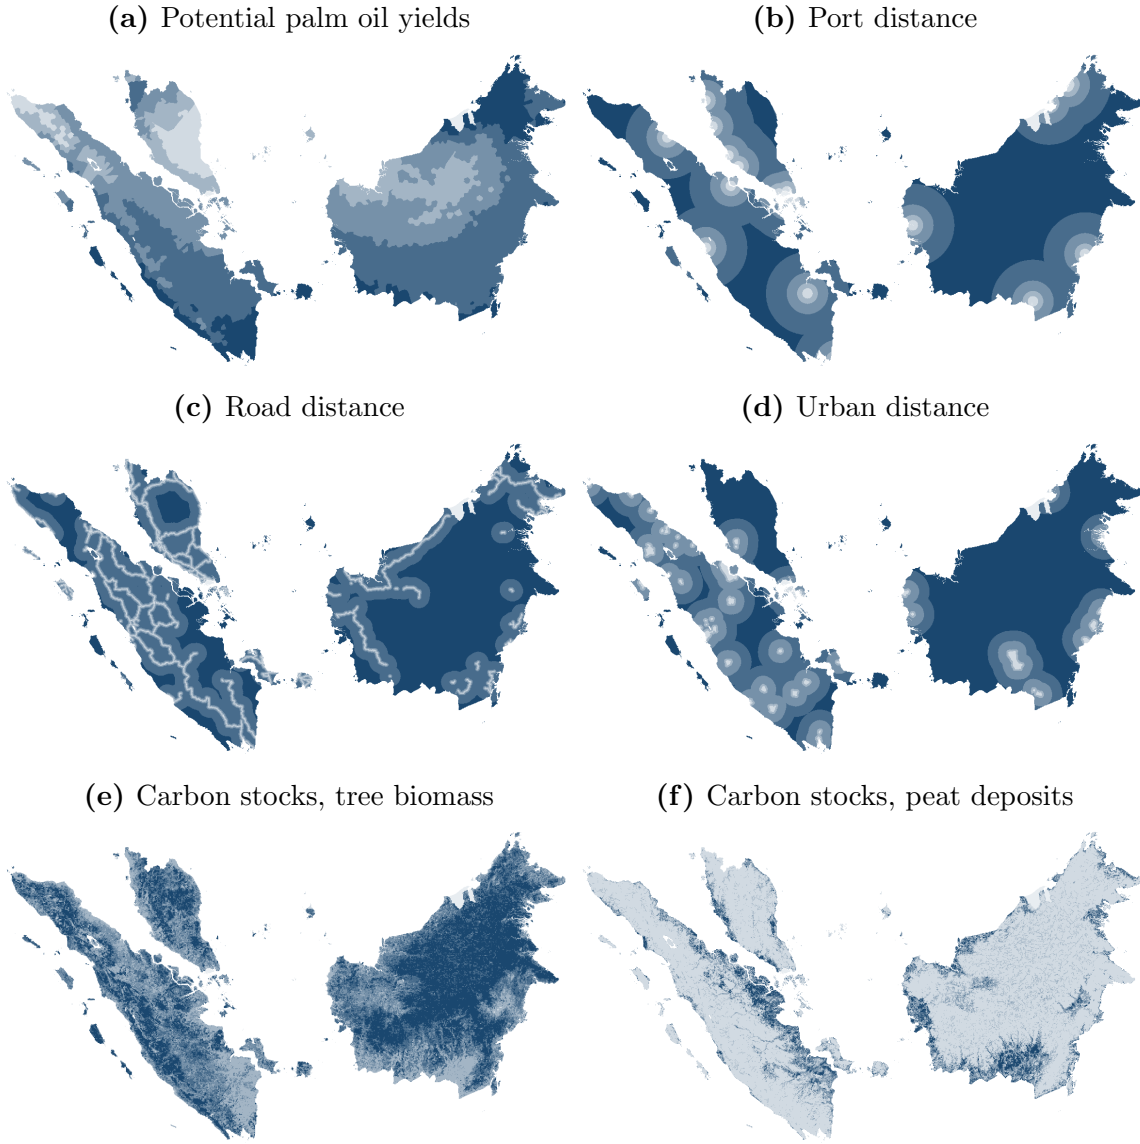
carbon would be released in developing any given plot of land and thus link counter-factual production to emissions.

For consumption, I compile annual panel data from 1988 to 2016 on palm oil and its substitutes. Consumption data by country come from the USDA Foreign Agricultural Service. Palm oils include palm and palm kernel, and other oils include coconut, olive, rapeseed, soybean, and sunflower. To address price endogeneity, I measure weather shocks to oil production. Rainfall and temperature data come from the Global Meteorological Forcing Dataset, which includes daily measures during the study period at 0.25° resolution. I identify producing regions – primarily states and provinces – with production data from the USDA Foreign Agricultural Service. For each crop, year, and region, I compute weather shocks as total absolute deviations from optimal levels during the growing season, with optimal levels given by the FAO Crop Ecological Requirements Database (ECOCROP). I then aggregate over regions, weighting by production, to obtain shocks by crop and year.

4 Empirical Model

This section describes empirical models that deliver palm oil demand and supply curves, which in turn correspond to $P_t^{Dr}(q)$, $P_t^{Du}(q)$, and $P_t^S(q)$ in section 2.

Figure 4: Land characteristics



Darker blue indicates high yields, farther distances, and larger carbon stocks. Yields are computed with the PALMSIM agronomic model ([Hoffmann et al. 2014](#)). Ports and roads are from the 2019 World Port Index, World Port Source, and Global Roads Inventory Project. Urban areas are administrative cities (*kota*) in Indonesia and federal territories in Malaysia. Carbon stocks are from [Zarin et al. \(2016\)](#) and [Gumbricht et al. \(2017\)](#).

4.1 Demand: an almost ideal demand system

I model aggregate demand for vegetable oils with a two-stage almost ideal demand system as in [Deaton and Muellbauer \(1980\)](#) and [Hausman et al. \(1994\)](#). Con-

sumers make an upper-level choice over total oil consumption. Given this total, they make a lower-level choice between palm and other oils, aggregated by Stone price index $\ln p_{it} = \sum_j \omega_{jt} \ln p_{jt}$. Relative to the characteristic-space approach, such as in [Berry et al. \(1995\)](#), this product-space approach allows for flexible substitution patterns and avoids the need to specify which product characteristics consumers value. Market-specific demand curves give me $P_t^{Dr}(q)$ and $P_t^{Du}(q)$ separately.

For each consumer market, the specifications are as follows. The lower level is

$$\omega_{it} = \alpha_i^0 + \alpha_i^1 t + \sum_j \gamma_{ij} \ln p_{jt} + \beta_i \ln \left(\frac{X_t}{P_t} \right) + \varepsilon_{it}, \quad (1a)$$

$$\ln P_t = \alpha_0 + \sum_j (\alpha_j^0 + \alpha_j^1 t) \ln p_{jt} + \frac{1}{2} \sum_j \sum_k \gamma_{jk} \ln p_{jt} \ln p_{kt}, \quad (1b)$$

for expenditure shares ω_{it} , palm and other oil prices p_{jt} , total oil expenditures $X_t = Q_t P_t$, and translog price index P_t . The upper level is

$$\ln Q_t = \alpha^0 + \alpha^1 t + \gamma \ln P_t + Z_t \beta + \varepsilon_t, \quad (2)$$

for quantity Q_t of total oil consumption and translog price index P_t . Demand shifters Z_t include GDP and the CPI, which capture overall income and prices.²

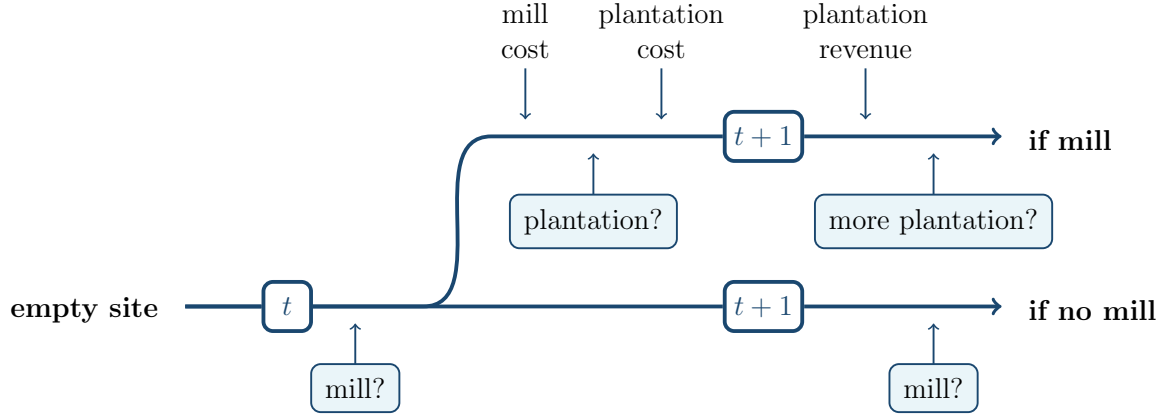
Both specifications are standard. For the upper level, an alternative is to specify total consumption in expenditure shares as in the lower level. However, vegetable oil expenditures are only 0.15% of GDP, and the resulting elasticities are unstable with expenditure shares so close to zero. Furthermore, the uncompensated price elasticities show why both levels are necessary.

$$e_{ijt} = \frac{\partial \ln q_{it}}{\partial \ln p_{jt}} = -\delta_{ij} + \frac{\gamma_{ij}}{\omega_{it}} + \left(\frac{\beta_i \gamma}{\omega_{it}} + \gamma + 1 \right) \left(\frac{\partial \ln P_t}{\partial \ln p_{jt}} \right), \quad (3)$$

where $\frac{\partial \ln P_t}{\partial \ln p_{jt}} = \alpha_j^0 + \alpha_j^1 t + \frac{1}{2} \sum_k (\gamma_{jk} + \gamma_{kj}) \ln p_{kt}$, Kronecker $\delta_{ij} = \mathbb{1}[i = j]$, and $q_{it} = \frac{\omega_{it} X_t}{p_{it}}$. The lower level allows substitution between palm and other oils (via γ_{ij}), and the upper level allows total category demand to respond to price changes (via γ).

² An important part of EU demand for palm oil is for biofuels. I do not include fossil fuels in the choice set because the EU has biofuel targets, such as for 14% of fuel for transportation to be renewable by 2030. Thus, higher palm oil prices arguably require substitution toward other vegetable oils rather than to fossil fuels. Including fossil fuels in the choice set would allow me to account for the substitution that occurs in the absence of these targets.

Figure 5: Supply model timeline



An empty site makes a binary choice over whether to construct a mill. If not, then the site faces the same choice 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 periods.

Prices are endogenous, as unobservables ε_{it} and ε_t shift demand and therefore affect equilibrium prices p_{jt} . I instrument with weather shocks to oil production as a supply shifter. The exclusion restriction is that these shocks affect vegetable oil demand only through their impact on prices. However, domestic shocks might also affect demand by impacting incomes or expenditures more broadly. I address this concern by isolating shocks to crops in producing states and provinces during the growing season, and also by directly testing for income and expenditure effects.

4.2 Supply: a dynamic model with sunk investment

Land is divided into sites, which I assume are small, independent, and managed by long-lived owners. Forward-looking sites generate profits by making sunk investments on two margins. On the extensive margin, sites make a binary choice over whether to build a mill. On the intensive margin, sites with mills make a continuous choice over how much land to develop into plantations.³ Figure 5 shows the timeline.

³ This model abstracts away from negotiations with smallholders, which account for 40% of production but are often vertically integrated into the production chain. In particular, smallholders are commonly bound by contracts that require selling harvests to specific mills in exchange for investment support (Cramb and McCarthy 2016). Even without vertical contracting, the intensive-margin model holds as long as investment is efficient, and the extensive-margin model holds as long as mills extract all surplus from plantations. Indeed, the perishability of harvest fruit gives mills spatial market power that helps in extracting rents.

Intensive margin (plantation development)

In each period t , sites i with mills make a continuous choice a_{it} over how much land to develop into plantations. Plantations have no scrap value and are sunk, such that plantation size s_{it} is given by law of motion $s_{it+1} = s_{it} + a_{it}$. Profits depend on observed state $\mathbf{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$ and unobserved state ε_{it} . Site-specific yields Y_{it} affect revenues, while site-specific cost factors x_i and shocks ε_{it} affect costs. Aggregate supply $s_t = \sum_i Y_{it}s_{it}$ and aggregate demand d_t affect world prices $P(s_t, d_t)$, which in turn affect revenues. Aggregate supply evolves endogenously: as in [Hopenhayn \(1992\)](#), atomistic sites affect world prices collectively but not individually, and firms play a dynamic competitive equilibrium in which collective action coincides with individual expectations. Aggregate demand evolves exogenously. Each period, sites with mills realize state $(\mathbf{w}_{it}, \varepsilon_{it})$ and make investment choice a_{it} , incurring costs in the current period to generate revenues in the following period and into the future.

The value, revenue, and cost functions are as follows, with $\mathbb{E}_{it}[\cdot] \equiv \mathbb{E}[\cdot | s_{it}, \mathbf{w}_{it}, \varepsilon_{it}]$.

$$V(s_{it}; \mathbf{w}_{it}, \varepsilon_{it}) = \max_{a_{it}} \{r(s_{it}; \mathbf{w}_{it}) - c(a_{it}; \mathbf{w}_{it}, \varepsilon_{it}) + \beta \mathbb{E}_{it}[V(s_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})]\}, \quad (4a)$$

$$r(s_{it}; \mathbf{w}_{it}) = Y_{it}P(s_t, d_t)s_{it}, \quad c(a_{it}; \mathbf{w}_{it}, \varepsilon_{it}) = \left(\frac{1}{2}\delta a_{it} + x_i\gamma + \kappa_m + \alpha_m t + \varepsilon_{it}\right)a_{it} \quad (4b)$$

Expectations are over next-period state $(\mathbf{w}_{it+1}, \varepsilon_{it+1})$, and I suppress function subscripts m . Linear revenues and convex costs ensure unique optima. Revenues are linear in plantation size and increasing in yields and world prices. Weather shocks ε_{it}^Y affect yields during production, but they are unrealized ex ante and thus do not enter here: sites invest based on climate and not weather. Costs are quadratic and convex in investment, spreading investment over time and reflecting credit and local factor market constraints, although upfront and future flow costs are not separately identified. Cost factors x_i capture observed heterogeneity by site, while fixed effects κ_m and time trends α_m accommodate unobserved heterogeneity by region. Cost shocks ε_{it} can be correlated across sites and over time.

Extensive margin (mill construction)

In each period t , sites i without mills make a binary choice a_{it}^e over whether to construct a mill. Plantations require mills because unmilled palm fruit decays

quickly after harvest and is not consumed directly. Mills have no scrap value and are sunk, with law of motion $s_{it+1}^e = s_{it}^e + a_{it}^e$. Profits depend on observed state $\mathbf{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$ and unobserved state ε_{it}^e , which captures mean-zero logit shocks $\{\varepsilon_{it0}^e, \varepsilon_{it1}^e\}$ with standard deviation σ^e . Each period, sites without mills realize state $(\mathbf{w}_{it}, \varepsilon_{it}^e)$ and make investment choice a_{it}^e . If they choose not to invest, then the period ends. If they choose to invest, then they immediately face the intensive-margin problem, realizing shock ε_{it} and making choice a_{it} before the period ends.

The ex-ante value, choice-specific conditional value, and cost functions are

$$V^e(\mathbf{w}_{it}) = \mathbb{E}_{it}^e[\max\{v^e(0; \mathbf{w}_{it}) + \varepsilon_{it0}^e, v^e(1; \mathbf{w}_{it}) + \varepsilon_{it1}^e\}], \quad (5a)$$

$$v^e(0; \mathbf{w}_{it}) = \beta \mathbb{E}_{it}^e[V^e(\mathbf{w}_{it+1})], \quad (5b)$$

$$v^e(1; \mathbf{w}_{it}) = -c^e(\mathbf{w}_{it}) + \mathbb{E}_{it}^e[V(0; \mathbf{w}_{it}, \varepsilon_{it})], \quad (5c)$$

$$c^e(\mathbf{w}_{it}) = x_i \gamma^e + \kappa_m^e + \alpha_m^e t, \quad (5d)$$

with $\mathbb{E}_{it}^e[\cdot] \equiv \mathbb{E}^e[\cdot | \mathbf{w}_{it}]$ and e superscripts referring to the extensive margin. In equation 5a, expectations are over logit shocks ε_{it}^e that imply mill construction probabilities

$$p^e(\mathbf{w}_{it}) = \frac{\exp[v^e(1; \mathbf{w}_{it})]}{\exp[v^e(0; \mathbf{w}_{it})] + \exp[v^e(1; \mathbf{w}_{it})]}. \quad (6)$$

In equation 5b, choosing not to build leads to the same decision in the following period given expectations over next-period state \mathbf{w}_{it+1} . The outside option is never constructing a mill, with utility normalized to zero given mean-zero shocks ε_{it}^e . In equation 5c, choosing to build incurs mill construction costs in return for the intensive-margin value of plantation development, where new plantations start with size $s_{it} = 0$. Expectations are over intensive-margin shock ε_{it} . In equation 5d, cost factors x_i capture observed heterogeneity by site, while fixed effects κ_m^e and time trends α_m^e accommodate unobserved heterogeneity by region. Logit cost shocks ε_{it}^e are uncorrelated across sites and over time, and also uncorrelated with intensive-margin shocks ε_{it} .

Unobserved heterogeneity and endogeneity

Both margins only allow for regional unobserved heterogeneity. Sites otherwise differ only in observables and shocks. Identifying site-level unobserved would require multiple plantation development and mill construction decisions over time. The for-

mer exists only for early sites, and the latter is inconsistent with a model in which sites construct no more than one mill each.

There is also an endogeneity problem on the intensive margin: both prices P_t and yields Y_{it} are correlated with cost shocks ε_{it} . First, collectively low costs induce entry, raising supply and lowering prices. Second, attained yields depend on unobserved, costly effort. Assuming uncorrelated cost shocks across sites addresses the first concern, but this assumption is strong. Instead, I instrument for prices with demand shifters d_t and for yields with potential yields Y_i^p . From estimated world demand

$$\ln p_t = -\hat{\phi} \ln q_t + \hat{d}_t,$$

I obtain shifters \hat{d}_t that capture changes in the level of demand over time. Potential yields are a function of climate, which is exogenous, and instrumenting also mitigates bias from mismeasured yields. These concerns do not arise on the extensive margin because mills themselves do not affect prices or yields, and because extensive- and intensive-margin cost shocks are assumed to be uncorrelated with each other.

I take cost factors x_i as exogenous. Port distance is to major ports, which predate plantations. Road distance is to major roads, and not to small roads that develop endogenously around plantations. Urban distance is to major cities, which do not include palm oil settlements. Carbon stocks are predetermined.

5 Estimation

This section describes estimation of the empirical model specified in section 4.

5.1 Demand: iterated linear least squares

I estimate the lower-level demand system with iterated linear least squares as in [Blundell and Robin \(1999\)](#). I start by estimating a linear approximate version, using a Stone price index instead of translog. I then construct the translog price index with the resulting estimates and iterate until convergence, thereby avoiding nonlinear estimation. Each iteration imposes the standard adding-up, homogeneity, and symmetry restrictions. Given the lower-level estimates, I estimate the upper

level by linear IV. Throughout, I instrument for prices with weather shocks to oil production, and Newey-West standard errors account for serial correlation. I compute demand elasticities by market-year and obtain standard errors with the delta method.

5.2 Supply: Euler methods

I estimate the supply model with Euler methods as in [Hall \(1978\)](#) and [Scott \(2013\)](#). On the intensive margin, I form Euler equations from the first order conditions for investment; on the extensive margin, I compare discrete, short-term perturbations that hold long-term investments fixed. Continuation values difference out.

Step 1: defining sites

I divide land into sites using observed mills and plantations as a guide. I identify the palm oil industry's most developed provinces and imagine bringing all provinces to this level of development. By several metrics, I obtain a target density of one mill per 521 km². I then define sites by k -means clustering on geographic coordinates, with the number of clusters in each province chosen to reach this target density. I impose that clusters separate observed mills and that observed plantations be assigned to clusters with observed mills. This procedure yields 2,135 contiguous sites, of which 1,467 have one observed mill and some observed plantations by 2016.

Step 2: intensive margin (plantation development)

The first order condition for investment delivers an Euler equation.⁴ Investing in period t instead of period $t + 1$ increases revenues, as production begins earlier, but it also increases costs, which are otherwise delayed and discounted.

$$c'(a_{it}; \mathbf{w}_{it}, \varepsilon_{it}) = \beta \mathbb{E}_{it}[r'(s_{it+1}; \mathbf{w}_{it+1}) + c'(a_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})] \quad (7)$$

I set $\beta = 0.9$, as the discount factor is typically unidentified in dynamic discrete choice models ([Magnac and Thesmar 2002](#)). By equation [4b](#), the Euler equation becomes

$$a_{it} - \beta \mathbb{E}_{it}[a_{it+1}] = \frac{\beta}{\delta} \mathbb{E}_{it}[Y_{it+1} P_{t+1}] - \frac{1 - \beta}{\delta} x_i \gamma - \frac{1 - \beta}{\delta} \kappa_m - \frac{1}{\delta} \alpha_m \tilde{t} - \frac{1}{\delta} \varepsilon_{it} + \frac{\beta}{\delta} \mathbb{E}_{it}[\varepsilon_{it+1}],$$

⁴ By equation [4a](#), the first order condition for a_{it} is $c'(a_{it}; \mathbf{w}_{it}, \varepsilon_{it}) = \beta \mathbb{E}_{it}[V'(s_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})]$, and the envelope theorem gives $V'(s_{it}; \mathbf{w}_{it}, \varepsilon_{it}) = r'(s_{it}; \mathbf{w}_{it}) + \beta \mathbb{E}_{it}[V'(s_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})]$.

with $P_t \equiv P(s_t, d_t)$ and $\tilde{t} \equiv t - \beta(t+1)$. In using the first order condition, I implicitly assume an interior solution. Indeed, 99.5% of observed intensive-margin decisions are interior: 0.5% involve zero development, and 0% exceed sites' total area.

I take realized values as noisy measures of expectations, which are unobserved, subject to expectational errors η_{it} as in [Hall \(1978\)](#). I obtain the regression equation

$$a_{it} - \beta a_{it+1} = \frac{\beta}{\delta} Y_{it+1} P_{t+1} - \frac{1-\beta}{\delta} x_i \gamma - \frac{1-\beta}{\delta} \kappa_m - \frac{1}{\delta} \alpha_m \tilde{t} - \frac{1}{\delta} \varepsilon_{it} + \frac{\beta}{\delta} \varepsilon_{it+1} + \eta_{it}. \quad (8)$$

Rational expectations are correct on average and use all available information. Thus, expectational errors are mean-zero and orthogonal to sites' period- t information sets.

$$\begin{aligned} \eta_{it} &= \beta \mathbb{E}_{it}[a_{it+1}] - \beta a_{it+1} + \frac{\beta}{\delta} \mathbb{E}_{it}[Y_{it+1} P_{t+1}] - \frac{\beta}{\delta} Y_{it+1} P_{t+1} + \frac{\beta}{\delta} \mathbb{E}_{it}[\varepsilon_{it+1}] - \frac{\beta}{\delta} \varepsilon_{it+1} \\ &= \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta} \left(\mathbb{E}_{it}[Y_{it+t'} P_{t+t'}] - \mathbb{E}_{it+1}[Y_{it+t'} P_{t+t'}] \right) \end{aligned} \quad (9)$$

The last line follows from equation 8.⁵ Investment choices, yields, prices, and cost factors are data, where cost factors include port, road, and urban distances, as well as carbon stocks. I instrument for yields and prices with lagged period- t values for potential yields and the demand shifters discussed above. Figures 3 and 4a plot the time-series variation in world prices and spatial variation in yields. Identification relies on both sources of variation: intuitively, price increases are more valuable for sites that produce more palm oil. I cluster by region given correlated cost shocks. Since revenues $Y_{it+1} P_{t+1}$ are measured directly, parameters γ , κ_m , and α_m are interpretable in dollar terms. Production begins one period after investment in this exposition, but in estimation I impose the typical three years to reach crop maturity.⁶

Step 3: extensive margin (mill construction)

As before, I compare investing today and tomorrow. In this discrete case, differencing stands in for the first order condition, and finite dependence for the envelope theorem. My comparison is between two sequences of actions: $(1, a_{it}^*, a_{it+1}^*)$ and $(0, 1, a'_{it+1})$ for $a'_{it+1} = a_{it}^* + a_{it+1}^*$. The first constructs a mill today, then develops

⁵ Telescoping implies $a_{it} = \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta} \mathbb{E}_{it}[Y_{it+t'} P_{t+t'}] - \frac{1}{\delta} x_i \gamma - \frac{1}{\delta} \kappa_m - \frac{1}{\delta} \alpha_m \tilde{t} - \frac{1}{\delta} \varepsilon_{it}$.

⁶ Each year is one period. Y_{t+1} terms become Y_{it+3} , and P_{t+1} terms become P_{t+3} . But a_{it+1} does not change because the intertemporal comparison is still between developing today or tomorrow.

a_{it}^* plantations today and a_{it+1}^* plantations tomorrow; the second constructs a mill tomorrow, then develops a'_{it+1} plantations tomorrow. Finite dependence holds when actions lead to common states – and thus common payoffs – in all future periods (Arcidiacono and Miller 2011). It holds here because, for both sequences, by period $t + 2$ the mill is constructed and the plantation is of size $a_{it}^* + a_{it+1}^*$. The payoffs are

$$\begin{aligned} v^e(1, a_{it}^*, a_{it+1}^*; \mathbf{w}_{it}) &= -c^e(\mathbf{w}_{it}) + \mathbb{E}_{it}^e[-c(a_{it}^*; \mathbf{w}_{it}, \varepsilon_{it}) + \beta r(a_{it}^*; \mathbf{w}_{it+1}) - \beta c(a_{it+1}^*; \mathbf{w}_{it+1}, \varepsilon_{it+1})] \\ &\quad + \beta^2 \mathbb{E}_{it}^e[V(a_{it}^* + a_{it+1}^*; \mathbf{w}_{it+2}, \varepsilon_{it+2})], \\ v^e(0, 1, a'_{it+1}; \mathbf{w}_{it}) &= -\beta \mathbb{E}_{it}^e[c^e(\mathbf{w}_{it+1}) + c(a'_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})] + \beta^2 \mathbb{E}_{it}^e[V(a'_{it+1}; \mathbf{w}_{it+2}, \varepsilon_{it+2})]. \end{aligned}$$

The continuation values align: $V(a_{it}^* + a_{it+1}^*; \mathbf{w}_{it+2}, \varepsilon_{it+2}) = V(a'_{it+1}; \mathbf{w}_{it+2}, \varepsilon_{it+2})$. By the Hotz-Miller inversion (Hotz and Miller 1993), equation 6 implies

$$\ln \left(\frac{p^e(\mathbf{w}_{it})}{1 - p^e(\mathbf{w}_{it})} \right) = v^e(1; \mathbf{w}_{it}) - v^e(0; \mathbf{w}_{it}). \quad (10)$$

While $v^e(1; \mathbf{w}_{it}) = v^e(1, a_{it}^*, a_{it+1}^*; \mathbf{w}_{it})$ by definition, $v^e(0; \mathbf{w}_{it})$ and $v^e(0, 1, a'_{it+1}; \mathbf{w}_{it})$ generally differ because the latter imposes particular choices.⁷ The difference is

$$v^e(0; \mathbf{w}_{it}) - v^e(0, 1, a'_{it+1}; \mathbf{w}_{it}) = \frac{1}{2} \beta \mathbb{E}_{it}^e[c''(a'_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})(a_{it+1}^* - a'_{it+1})^2] - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})].$$

Substituting into equation 10 and applying the revenue and cost functions, I obtain an Euler equation in which continuation values difference out.

$$\ln \left(\frac{p^e(\mathbf{w}_{it})}{1 - p^e(\mathbf{w}_{it})} \right) - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})] = \mathbb{E}_{it}^e[I_{it+1}] - (1 - \beta)x_i\gamma^e - (1 - \beta)\kappa_m^e - \alpha_m^e \tilde{t},$$

for $\tilde{t} = t - \beta(t + 1)$ and $I_{it+1} = [\beta Y_{it+1} P_{t+1} - (1 - \beta)x_i\gamma - (1 - \beta)\kappa_m - \alpha_m \tilde{t}]a_{it}^* + \delta[-\frac{1}{2}a_{it}^{*2} + \beta a_{it}^* a_{it+1}^*]$.

Substituting expectational errors and estimated values yields regression equation

$$\ln \left(\frac{\hat{p}^e(\mathbf{w}_{it})}{1 - \hat{p}^e(\mathbf{w}_{it})} \right) - \beta \ln \hat{p}^e(\mathbf{w}_{it+1}) = \hat{I}_{it+1} - (1 - \beta)x_i\gamma^e - (1 - \beta)\kappa_m^e - \alpha_m^e \tilde{t} + \eta_{it}^e. \quad (11)$$

I estimate conditional choice probabilities $\hat{p}^e(\mathbf{w}_{it})$ non-parametrically by regressing observed investment choices on a flexible set of basis terms: piecewise linear splines in

⁷ $v^e(1, a_{it}^*, a_{it+1}^*; \mathbf{w}_{it})$ has $a_{it}^*(0; \mathbf{w}_{it}, \varepsilon_{it})$ and $a_{it+1}^*(a_{it}^*; \mathbf{w}_{it+1}, \varepsilon_{it+1})$, while $v^e(0; \mathbf{w}_{it})$ has $a_{it+1}^*(0; \mathbf{w}_{it+1}, \varepsilon_{it+1})$. But $a_{it+1}^*(a_{it}^*; \mathbf{w}_{it+1}, \varepsilon_{it+1}) = a_{it+1}^*(0; \mathbf{w}_{it+1}, \varepsilon_{it+1})$ for linear revenues.

Y_{it+1} , P_{t+1} , x_i , and \tilde{t} , as well as their interactions. I do so separately for each region to account for regional unobserved heterogeneity, and I estimate intensive-margin choices \hat{a}_{it}^* similarly. Dollar-denominated intensive-margin profits \hat{I}_{it+1} provide a scale normalization that allows parameters γ^e , κ_m^e , and α_m^e to be interpreted in dollar terms, with intercepts κ_m^e identified relative to the outside option.

Discussion

Euler estimation has several advantages. First, I can address endogeneity concerns using standard instrumental variable techniques because estimation reduces to linear regression. Second, while I need to assume rational expectations, I do not need to specify what expectations are. By contrast, the conventional full-solution approach requires explicit structure on expectations over the long-run horizon. Rational expectations remains a strong assumption, but regional effects κ_m absorb expectational bias to the extent that it is fixed within regions. Third, the full-solution approach requires solving the model repeatedly and computing continuation values with each iteration. The Euler approach sidesteps this computational burden because it estimates the model without solving it. Other methods have similar computational advantages in the discrete case, but they cannot accommodate the non-stationarity of my setting (Aguirregabiria and Mira 2007; Bajari et al. 2007; Pakes et al. 2007; Pesendorfer and Schmidt-Dengler 2008).⁸

Estimation relies on several other assumptions. First, I compare investing today or tomorrow, but weak property rights may promote land grabbing and thus bias toward investing today. Regional effects κ_m help by absorbing regional variation in property rights. Second, sites are independent and atomistic. Otherwise, finite dependence does not hold: if price-makers delay investment, then competitors respond, altering the evolution of the economy such that continuation values do not align. Indeed, world production is unconcentrated, with the largest producer accounting for 4% and the largest ten for 21%. But I must rule out spatial competition, including in local factor markets, as spatial interaction makes estimation intractable. Third, plantation age does not affect profits. Otherwise, delayed investment affects profits in all future periods, and finite dependence again does not hold.

⁸ Besides Scott (2013), recent applications of discrete Euler methods include Diamond et al. (2017), De Groote and Verboven (2019), Traiberman (2019), and Almagro and Domínguez-Iino (2020).

Table 2: Weather shocks as price instruments

	All	All	Palm	Other
Rainfall shocks (100 mm)	0.208*** (0.0317)	0.212*** (0.0278)	0.139*** (0.0325)	0.224*** (0.0318)
Temperature shocks (°C)	0.297*** (0.0335)	0.308*** (0.0315)	0.681 (0.804)	0.315*** (0.0334)
Oil FE	x	x		
Oil-year trend		x		
Year trend			x	x
Observations	174	174	29	145
F-statistic	40.94	49.25	10.56	48.90

Each column is a regression, and the outcome variable is log prices. Data are annual and 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 regions. Oil fixed effects and oil-specific time trends differentiate between palm and other oils. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6 Estimates


This section presents parameter estimates. Demand estimates suggest inelastic Indonesian and Malaysian demand, and supply estimates quantify production costs.

6.1 Demand

In the first stage, table 2 shows that both rainfall and temperature shocks significantly increase world oil prices. The first two columns pool across oil products and show that weather shocks have significant price effects. The last two columns show estimates for palm and other oils separately. For palm oil, a smaller sample size means less precision, but the point estimates are relatively close to those of the pooled specifications, and the instruments remain strong. Temperature effects are perhaps imprecisely estimated because palm oil is grown in tropical climates with limited year-to-year variation in temperatures. Toward assessing the exclusion restriction, I also find that weather shocks do not directly affect overall incomes or expenditures.

Figure 6 plots the estimated demand curves. Price-responsive demand among non-EU importers suggests that unilateral EU action is susceptible to leakage, partic-

Figure 6: Palm oil demand



Demand estimation draws on annual data covering coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016, and it instruments for prices with weather shocks to oil production. The left figure averages over the study period for each market. The right figure shows demand over time, and each faded line is one year between 1988 and 2016.

ularly because non-EU importers account for a substantial 68% of global consumption (table 1). By contrast, Indonesia and Malaysia have quite inelastic demand, limiting leakage if importers coordinate on tariffs. As producers, Indonesia and Malaysia consume far more palm oil than any other oil – consistent with home bias. Few effective substitutes leads to inelastic palm oil demand. Furthermore, demand for palm oil is rising rapidly over time as the demand curve shifts rightward in log scale.

One limitation is that the demand model is static. I can rule out significant bias from stockpiling because I observe oil stocks held in storage facilities and find that they are small. However, demand may be sticky because switching between oils requires reformulating recipes and rewriting contracts with suppliers. If so, then larger long-run elasticities imply more leakage. Another concern is that palm oil tariffs may encourage unregulated markets to supply regulated markets with palm oil in other forms, such as palm-based biofuels. Static demand estimates do capture the short-term responses of existing industries, as the consumption data include palm oil used as inputs. However, static demand cannot capture the long-term response of new industries short of modeling them. A simple solution is for import tariffs to cover both palm oil and palm oil content, although most palm oil is exported in raw form and thus there may be limited gains from such a policy.

Table 3: Intensive-margin supply regressions

	OLS	IV	First stage
	$a_{it} - \beta a_{it+1}$	$a_{it} - \beta a_{it+1}$	$Y_{it+3}P_{t+3}$
Yield \times price ($Y_{it+3}P_{t+3}$)	0.105*** (0.00664)	0.180*** (0.0683)	
Potential yield \times demand ($Y_i^p d_t$)			19.15*** (0.845)
Province FE	x	x	x
Province-year trend	x	x	x
Observations	20,046	20,046	20,046
F-statistic			513

Each column is one regression, and each observation is a site-year. Column headings show the specifications and dependent variables. I instrument for the interaction of yields and prices with the interaction of potential yields and demand shifters. Potential yields are computed with the agronomic model of [Hoffmann et al. \(2014\)](#), and demand shifters are computed in demand estimation. Prices combine palm and palm kernel oil prices measured in inflation-adjusted, year-2000 dollars. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6.2 Supply

Tables 3 and 4 present supply estimates. Table 3 shows that higher revenues increase development, with a larger IV estimate and a strong instrument. Table 4 shows the estimated model parameters, which are interpretable in dollar terms. I estimate average lifetime costs of \$6,000 per hectare of plantation and \$36 million per mill, both of which decrease over time. Plantation costs include quadratic costs with 542 hectares of new development on average. Accounting estimates suggest \$9,000 and \$20 million ([Fairhurst and McLaughlin 2009](#); [Man and Baharum 2011](#)), implying similar total costs given 8,265 hectares of plantation per average mill.

Cost factor estimates reveal what producers internalize. On the intensive margin, producers treat all land similarly and simply develop around constructed mills. On the extensive margin, producers internalize transport costs. Distance from major ports, roads, and urban centers meaningfully discourages production. By contrast, producers do not internalize the cost of emissions. Tree biomass is a modest deterrent, as developing heavily forested areas requires additional effort and may face local scrutiny. But peat deposits – the main source of emissions – have little effect and even slightly encourage production, as peatlands have few alternative uses.

Table 4: Supply model parameter estimates

	Intensive-margin		Extensive-margin	
	Mean	SE	Mean	SE
Province-specific costs ($\bar{\kappa}_m, \bar{\kappa}_m^e$)	4,657**	(1,915)	36,140,684***	(2,173,896)
Province-specific cost trends ($\bar{\alpha}_m, \bar{\alpha}_m^e$)	-844***	(119)	-647,242***	(116,021)
Cost factors (γ, γ^e)				
Log port distance, km	-531	(425)	6,727,143***	(1,134,264)
Log road distance, km	-458	(310)	6,328,282***	(643,226)
Log urban distance, km	-19	(221)	8,765,560***	(936,047)
Log carbon in tree biomass, t	447	(441)	3,708,482***	(840,238)
Log carbon in peat deposits, t	-137	(88)	-398,526**	(162,985)
Quadratic costs (δ)	5***	(2)	—	—
Logit scale (σ^e)	—	—	6,496,442***	(492,005)

Estimates are interpretable in terms of inflation-adjusted, year-2000 dollars. I report mean province-specific parameters weighted by number of sites. Province-specific costs are for a mean year at mean values for cost factors. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7 Counterfactuals

This section evaluates counterfactual tariff policies. I find that import tariffs are effective in place of domestic regulation, but only if coordinated and committed.

7.1 Tariffs

Emissions respond to tariffs. First, tariffs τ_t affect world prices P_t as a function of demand curves $P_t^{Dr}(q)$ and $P_t^{Du}(q)$. Prices net of tariffs equalize across markets, as tariffs in regulated markets cause producers to shift sales elsewhere.

$$P_t^{Dr}(Q_t^{ro}) - \tau_t = P_t^{Du}(Q_t^{uo})$$

The demand model characterizes both the regulated and unregulated demand curves. Second, world prices P_t affect the investments that produce supply $Q_t^{ro} + Q_t^{uo}$ as characterized by the supply model. Finally, a carbon map connects changes in supply to emission reductions. The baseline analysis considers tariffs that are uniform and maximize social welfare, although extensions relax both conditions.

I quantify the effects of coordination and commitment by studying three tariff coalitions, each under full, no, and limited commitment. Coalitions include all importers together, an EU-China-India partnership, and the EU alone. Under full commitment, tariffs are set and upheld in perpetuity. Under no commitment, tariffs are reset every period as in sequential static optimization. Under limited commitment, L -period plans specify tariffs for L periods at a time before getting revised.

7.2 Solving the model

Unlike estimation, counterfactuals require solving the model. For a given set of tariffs, I can solve by backward induction from steady-state period S . At this point, all lands are exhausted and there is no further development, but plantations continue to generate revenues over the infinite horizon. Cost shocks and finite land ensure that such a steady state exists, although it may take many periods to reach.

The challenge is that backward induction over long horizons is computationally intensive. I address this challenge in two ways. First, I start with an arbitrary period $T < S$ and solve as if it were the steady state. This approach is biased if substantial development occurs after T , as future entry by competitors makes it less appealing to enter today. But future entry diminishes concavely as T approaches S , and so I increase T and re-solve until convergence. Second, I solve each subproblem using an iterative algorithm that approximates the solution with a limited look-ahead horizon. This algorithm breaks the typical curse of dimensionality in which the state space grows exponentially with the length of the horizon considered.

7.3 Quantifying emissions

I quantify carbon emissions by combining the model’s site-level predictions for plantation development with site-level data on carbon stored both aboveground in trees and belowground in peat. The underlying assumption is that plantation development releases all carbon stocks. Indeed, trees must be cut to make space for plantations, and peat must be cleared to access soil beneath the peat layer.

On the demand side, I ignore the carbon effects of substitution to other oil products. The primary threat is South American soybean oil, which contributes to Amazonian deforestation. I argue that 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 one of several substitutes at 13% of total oil consumption.

On the supply side, I ignore the carbon effects of substitution to other drivers of deforestation. The primary threat is substitution to acacia (paper pulp) plantations, which also destroy peatlands. I argue that the resulting bias is small because palm oil is seven times more profitable than acacia, which requires replanting after harvest, such that switching to acacia is unappealing for many palm oil producers (Sofiyuddin et al. 2012). Indeed, I do not find that palm displaces acacia development in the data, at least in partial equilibrium. Other settings might require multi-product tariffs and a multi-industry model in which producers first choose among products, then invest as in the baseline model. Other drivers of deforestation include mining and selective logging, but mining relies on the exogenous distribution of deposits, and selective logging does not destroy peatlands.

7.4 Coordinated, committed tariffs are effective

Table 5 shows that import tariffs can be effective in reducing carbon emissions. When importers coordinate on import tariffs, and when they commit to upholding them, carbon emissions are reduced by 56%. By comparison, the socially optimal domestic tax reduces carbon emissions by 64%. The difference arises from leakage to domestic consumption in Indonesia and Malaysia, which is not exported and therefore not subject to import tariffs. However, the loss is not disproportionate because demand in Indonesia and Malaysia is quite inelastic. Indeed, importers impose tariffs nearly as large as the domestic tax given limited leakage concerns. Furthermore, the magnitude of the emission externality leads to a domestic tax that is itself quite large at several times observed prices.

Emission reductions diminish as coordination and commitment weaken. Figure 7a plots emission reductions under each of the scenarios in table 5. First, weak coordination lowers emission reductions because importers have relatively elastic demand. Emissions fall by up to 56% under full coordination among importers, 17% under an EU-China-India coalition, and 2% under unilateral EU action. These emission reductions fall disproportionately more than tariff coverage – 80%, 35%, and 12% of world

Table 5: Counterfactual experiments

Experiment	$\$/t$ NPV	$\Delta\%$	$\Delta\%$ surplus				$\$/t$ CO ₂
	Tax	CO ₂	EU	China India	Other	Indo Malay	Avg cost
Domestic regulation	20,487	-64	-93	-65	-31	-61	20
Import tariffs: full coordination							
Full commitment	19,718	-56	-86	-58	-25	-71	24
Limited commitment (20 years)	19,665	-56	-86	-58	-24	-71	24
Limited commitment (10 years)	19,476	-55	-85	-57	-24	-70	24
Limited commitment (5 years)	18,639	-53	-80	-54	-22	-67	24
Import tariffs: EU, China, India							
Full commitment	11,573	-17	-49	-32	45	-21	16
Limited commitment (20 years)	11,156	-16	-47	-30	43	-20	16
Limited commitment (10 years)	9,882	-14	-40	-25	38	-18	16
Limited commitment (5 years)	6,445	-9	-23	-13	25	-12	15
Import tariffs: EU only							
Full commitment	6,785	-2	-11	10	5	-3	10
Limited commitment (20 years)	6,445	-2	-10	10	5	-3	9
Limited commitment (10 years)	5,466	-2	-7	8	4	-2	9
Limited commitment (5 years)	3,197	-1	-3	5	2	-1	8

Columns show the net present value of tariffs or taxes, percentage changes in emissions or surplus relative to observed values, and average social surplus losses for averted emissions. Figures for Indonesia and Malaysia combine consumer and producer surplus, and all figures include government tax or tariff revenue where applicable. The first panel is for a socially optimal domestic tax in Indonesia and Malaysia. The second, third, and fourth panels are for foreign import tariffs under full coordination among importers, under an EU-China-India coalition, and for the EU alone. Full commitment is over all future periods, and limited commitment is for five, ten, or twenty years at a time. Without any commitment, tariffs are not imposed at all. The discount factor is $\beta = 0.9$.

consumption, respectively – because leakage concerns lead to smaller tariffs. Second, weak commitment lowers emission reductions, especially when the commitment period does not exceed time to build. In this case, tariffs are set to zero because they have no effect on new development, which does not generate taxable production during the commitment period. Third, coordination and commitment interact: figure 7b shows how coordination lessens the burden of commitment. Under full coordination, a five-year commitment period achieves 95% of full-commitment outcomes. But under an EU-China-India coalition or unilateral EU action, only twenty-year commitment periods approach full-commitment outcomes.

Figure 7: Counterfactual emissions

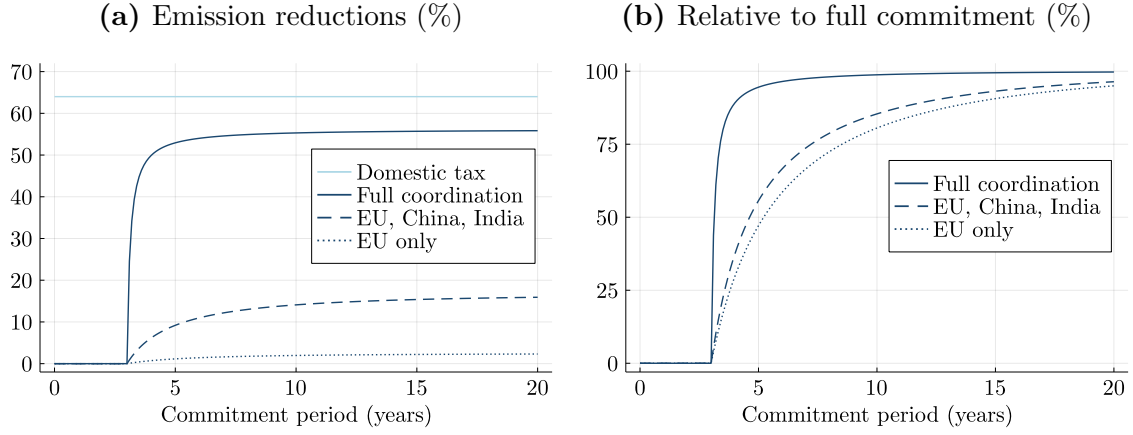


Figure 7a shows emission reductions under several scenarios. The light blue line is the socially optimal domestic tax. The solid line is full coordination among all importers, excluding domestic consumption in Indonesia and Malaysia, the dashed line is an EU-China-India coalition, and the dotted line is unilateral EU action. Each line shows emission reductions for each of the commitment periods listed on the x -axis. Emission reductions are zero when the commitment period does not exceed time to build because otherwise tariffs do not influence new development. Figure 7b rescales emission reductions for the tariff scenarios relative to their levels under full commitment.

The division of surplus highlights why coordination and commitment are difficult in practice. For coordination, importers that defect from the tariff coalition can free-ride on the emission reductions and lower world prices induced by the coalition. For example, focusing on full commitment, other importers have 25% lower consumer surplus when they join an EU-China-India coalition, but 45% higher consumer surplus when they defect. For commitment, moving to lower commitment levels leads to smaller tariffs and smaller sacrifices in consumer surplus. For example, focusing on full coordination, all importers have higher surplus under five-year commitment than they do under full commitment. Tariff coordination and commitment may require transfers aimed at counteracting these incentives.

More broadly, the same considerations underscore why Indonesia and Malaysia have not implemented the socially optimal domestic tax, even independent of enforcement issues. Only the threat of fully coordinated tariffs is enough to motivate the domestic tax. Uncoordinated tariffs are small and reduce producer surplus only modestly, leaving Indonesia and Malaysia better off accepting the tariffs. By contrast, fully coordinated tariffs are only slightly lower than the domestic tax, which generates government revenue for Indonesia and Malaysia that otherwise goes abroad.

Table 6: Robustness and extensions, carbon emission reductions ($\Delta\%$ CO₂)

	Coordination:	All importers		EU-China-India		EU alone	
	Commitment:	20-year	5-year	20-year	5-year	20-year	5-year
Baseline		-56	-53	-16	-9	-2	-1
Discount factor							
$\beta = 0.8$		-75	-71	-21	-12	-3	-2
$\beta = 0.85$		-65	-61	-18	-11	-3	-1
$\beta = 0.95$		-48	-46	-14	-8	-2	-1
Demand elasticity, Indonesia/Malaysia							
$\varepsilon^{DI}, \varepsilon^{DM} = 0.2$		-50	-43	-13	-7	-2	-1
$\varepsilon^{DI}, \varepsilon^{DM} = 0.4$		-44	-34	-10	-5	-2	-1
$\varepsilon^{DI}, \varepsilon^{DM} = 0.6$		-39	-28	-8	-4	-1	-1
Objective function, own surplus only		-57	-54	-3	-2	-0	-0
Conditioning on unit-specific emissions		-80	-75	-22	-12	-2	-1
Static supply		-5	-4	-1	-0	-0	-0

Each cell is one counterfactual. The first panel corresponds to table 5. The second panel changes the discount factor. The third panel changes the elasticities of Indonesian and Malaysian demand. The fourth panel sets tariff to maximize the surplus of the acting coalition, net of its own costs of carbon as computed by [Ricke et al. \(2018\)](#). The fifth panel allows import tariffs to condition on the emissions specific to each unit of palm oil. The last panel assumes a static supply model.

Nonetheless, each scenario has net negative effects for Indonesia and Malaysia, perhaps motivating compensatory transfers that this quantification exercise can inform.

Finally, while uncoordinated action has limited effects, these effects are still positive. Unilateral EU action reduces carbon emissions by up to 2% and at low average cost. Import tariffs allow governments to do more in the fight against climate change, even if they have exhausted low-cost options domestically.

7.5 Robustness and extensions

Table 6 presents robustness checks and extensions. First, lower discount factors increase the benefits of delaying development. Second, more elastic Indonesian and Malaysian demand decreases carbon reductions, although coordinated tariffs still have large effects. Third, I set tariffs to maximize the coalition's own surplus instead

of social surplus.⁹ Coordination raises tariffs as leakage falls, but also as the coalition internalizes more of the social cost of carbon and, to a lesser extent, as it has more power to improve terms of trade. Fourth, conditioning tariffs on unit-specific emissions leads to larger carbon reductions by more efficiently targeting peatland destruction. In practice, however, such policies rely on green certification, which is itself subject to commitment problems: since deforestation is sunk, it is statically optimal to certify previously deforested land as green. Fifth, a static supply model leads to low supply elasticities and much smaller effects. Dynamics matter quantitatively.

8 Conclusion

The conventional approach to environmental regulation focuses on domestic intervention, but domestic regulation can face major challenges. Governments may prioritize local profits over global consequences or lack the capacity to enforce regulation. Trade policy offers the international community a set of tools to intervene when domestic policies fail. This paper argues that trade policy requires both coordination and commitment to be effective. Without coordination, tariffs are undermined by leakage to unregulated markets. Without commitment, tariffs are reduced over time as importers give in to static incentives.

I develop a dynamic empirical framework for quantifying these forces, and I apply it to studying proposed EU tariffs on imports of palm oil, which accounts for more CO₂ emissions over the last three decades than the entire economy of India. Using data from satellite imagery, the framework delivers predictions of plantation development – and therefore deforestation – at a fine level of spatial disaggregation. I find that EU tariffs are most effective when coordinated with other major importers like China and India, and when regulators can commit to upholding them over the long term. Coordinated, committed tariffs are comparable to domestic regulation, reducing carbon emissions by 56% compared to 64%. At the same time, coordination and commitment are complements: when either fails, EU action has limited effects. These findings underscore the significance of the Paris Agreement, as well as the implications of US withdrawal.

⁹ The EU, China/India, other importers, and Indonesia/Malaysia bear 1%, 17%, 80%, and 2% of the social costs of carbon, respectively, based on pooled estimates from [Ricke et al. \(2018\)](#).

I leave several directions open for future work. First, trade policy may interact with domestic regulation in settings where domestic regulation is at least partially feasible. Second, dynamic bargaining considerations could influence the formation and stability of tariff coalitions in ways that I do not explicitly capture here. Third, spatial interaction among plantations might create path dependence, which in turn can amplify the effects of import tariffs on palm oil. Each direction invites more work on this important topic, particularly as expansion along new frontiers threatens vast swathes of forest that until now have remained untapped.

References

- Abrego, Lisandro and Carlo Perroni. Investment Subsidies and Time-Consistent Environmental Policy. Oxford Economic Papers, 54(4):617–635, 2002.
- Acemoglu, Daron and Will Rafey. Mirage on the Horizon: Geoengineering and Carbon Taxation Without Commitment. 2019.
- Aguirregabiria, Victor and Arvind Magesan. Euler Equations for the Estimation of Dynamic Discrete Choice Structural Models. Advances in Econometrics, 31:3–44, 2013.
- Aguirregabiria, Victor and Pedro Mira. Sequential Estimation of Dynamic Discrete Games. Econometrica, 75(1):1–53, 2007.
- Almagro, Milena and Tomás Domínguez-Iino. Location Sorting and Endogenous Amenities: Evidence from Amsterdam. 2020.
- Araujo, Rafael, Francisco Costa, and Marcelo Sant’Anna. Efficient Forestation in the Brazilian Amazon: Evidence from a Dynamic Model. 2021.
- Arcidiacono, Peter and Paul Ellickson. Practical Methods for Estimation of Dynamic Discrete Choice Models. Annual Review of Economics, 3:363–394, 2011.
- Arcidiacono, Peter and Robert Miller. Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity. Econometrica, 79(6):1823–1867, 2011.
- Assunção, Juliano, Robert McMillan, Joshua Murphy, and Eduardo Souza-Rodrigues. Optimal Environmental Targeting in the Amazon Rainforest. 2021.
- Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. Estimating Dynamic Models of Imperfect Competition. Econometrica, 75(5):1331–1370, 2007.
- Battaglini, Marco and Bård Harstad. Participation and Duration of Environmental Agreements. Journal of Political Economy, 124(1):160–204, 2016.
- Berry, Steven, James Levinsohn, and Ariel Pakes. Automobile Prices in Market Equilibrium. Econometrica, 63(4):841–890, 1995.
- Blundell, Richard and Jean Marc 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, Jared Carbone, and Thomas Rutherford. The Strategic Value of Carbon Tariffs. American Economic Journal: Economic Policy, 8(1):28–51, 2016.
- Brunner, Steffen, Christian Flachsland, and Robert Marschinski. Credible Commitment in Carbon Policy. Climate Policy, 12(2):255–271, 2012.

- Burgess, Robin, Matthew Hansen, Benjamin Olken, Peter Potapov, and Stefanie Sieber. The Political Economy of Deforestation in the Tropics. Quarterly Journal of Economics, 127(4):1707–1754, 2012.
- Burgess, Robin, Francisco Costa, and Benjamin Olken. The Brazilian Amazon’s Double Reversal of Fortune. 2019.
- Busch, Jonah, Kalifi Ferretti-Gallon, Jens Engelmann, Max Wright, Kemen Austin, Fred Stolle, Svetlana Turubanova, Peter Potapov, Belinda Margono, Matthew Hansen, and Alessandro Baccini. Reductions in Emissions from Deforestation from Indonesia’s Moratorium on New Oil Palm, Timber, and Logging Concessions. Proceedings of the National Academy of Sciences of the United States of America, 112(5):1328–1333, 2015.
- Collard-Wexler, Allan. Demand Fluctuations in the Ready-Mix Concrete Industry. Econometrica, 81(3):1003–1037, 2013.
- Copeland, Brian and M. Scott Taylor. North-South Trade and the Environment. Quarterly Journal of Economics, 109(3):755–787, 1994.
- Copeland, Brian and M. Scott Taylor. Trade and Transboundary Pollution. American Economic Review, 85(4):716–737, 1995.
- Cramb, Rob and John McCarthy, editors. The Oil Palm Complex: Smallholders, Agribusiness and the State in Indonesia and Malaysia. NUS Press, Singapore, 2016.
- Davis, Steven, Glen Peters, and Ken Caldeira. The Supply China of CO2 Emissions. Proceedings of the National Academy of Sciences, 108(45):18554–18559, 2011.
- De Groote, Olivier and Frank 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 John Muellbauer. An Almost Ideal Demand System. American Economic Review, 70(3):312–326, 1980.
- Diamond, Peter. Consumption Externalities and Imperfect Corrective Pricing. The Bell Journal of Economics and Management Science, 4(2):526–538, 1973.
- Diamond, Rebecca, Tim McQuade, and Franklin Qian. The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco. 2017.
- Domínguez-Iino, Tomás. Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains. 2021.
- Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas 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, Walter Falcon, Gracia Hadiwidjaja, Matthew Higgins, Rosamond Naylor, and Sudarno Sumarto. Fight Fire with Finance: A Randomized Field Experiment to Curtail Land-clearing Fire in Indonesia. 2020.
- Elliott, Joshua, Ian Foster, Samuel Kortum, Todd Munson, Fernando Pérez Cervantes, and David Weisbach. Trade and Carbon Taxes. American Economic Review: Papers & Proceedings, 100:465–469, 2010.
- Erdem, Tülin, Susumu Imai, and Michael Keane. Brand and Quantity Choice Dynamics Under Price Uncertainty. Quantitative Marketing and Economics, 1:5–64, 2003.
- Ericson, Richard and Ariel Pakes. Markov-Perfect Industry Dynamics: A Framework for Empirical Work. Review of Economic Studies, 62(1):53–82, 1995.
- Fairhurst, Thomas and David McLaughlin. Sustainable Oil Palm Development on Degraded Land in Kalimantan. Technical report, World Wildlife Fund, 2009.
- Farrokhi, Farid and Ahmad Lashkaripour. Can Trade Policy Mitigate Climate Change? 2021.
- Fowlie, Meredith. Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage. American Economic Journal: Economic Policy, 1(2):72–112, 2009.
- Fowlie, Meredith, Mar Reguant, and Stephen Ryan. Market-Based Emissions Regulation and Industry Dynamics. Journal of Political Economy, 124(1):249–302, 2016.
- Gaveau, David, Bruno Locatelli, Mohammad Salim, Husna Yaen, Pablo Pacheco, and Douglas Sheil. Rise and Fall of Forest Loss and Industrial Plantations in Borneo (2000–2017). Conservation Letters, 12(3):e12622, 2019.
- Gumbrecht, Thomas, Rosa Maria Roman-Cuesta, Louis Verchot, Martin Herold, Florian Wittmann, Ethan Householder, Nadine Herold, and Daniel Murdiyarso. 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 Kenneth Singleton. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. Econometrica, 50(5):1269–1286, 1982.
- Hansen, M. C., P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend. High-Resolution Global Maps of 21st-Century Forest Cover Change. Science, 342(6160):850–853, 2013.

- Harstad, Bård. Buy Coal! A Case for Supply-Side Environmental Policy. Journal of Political Economy, 120(1), 2012.
- Harstad, Bård. The Market for Conservation and Other Hostages. Journal of Economic Theory, 166:124–151, 2016.
- Harstad, Bård. Technology and Time Inconsistency. Journal of Political Economy, 128(7): 2653–2689, 2020.
- Harstad, Bård. Trade and Trees. 2021.
- Harstad, Bård and Torben Mideksa. Conservation Contracts and Political Regimes. Review of Economic Studies, 84:1708–1734, 2017.
- Hausman, Jerry, Gregory Leonard, and J. Douglas Zona. Competitive Analysis with Differentiated Products. Annales d'Economie et de Statistique, 34:159–180, 1994.
- Helm, Dieter, Cameron Hepburn, and Richard Mash. Credible Carbon Policy. Oxford Review of Economic Policy, 19(3):438–450, 2003.
- Hendel, Igal and Aviv 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, Alba Castañeda-Vera, Mark van Wijk, Ken Giller, Thomas Oberthuer, Christopher Donough, and Anthony Whitbread. 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 Robert Miller. Conditional Choice Probabilities and the Estimation of Dynamic Models. Review of Economic Studies, 60(3):497–529, 1993.
- Jayachandran, Seema, Joost de Laat, Eric Lambin, Charlotte Stanton, Robin Audy, and Nancy Thomas. Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation. Science, 357:267–273, 2017.
- Johnson, Harry. Optimum Tariffs and Retaliation. Review of Economic Studies, 21(2): 142–153, 1953.
- Kalouptside, Myrto, Paul Scott, and Eduardo Souza-Rodrigues. Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models. Journal of Econometrics, 222(1):778–804, 2021.
- Kortum, Samuel and David Weisbach. The Design of Border Adjustments for Carbon Prices. National Tax Journal, 70(2):421–446, 2017.

- Kortum, Samuel and David Weisbach. Optimal Unilateral Carbon Policy. 2021.
- Magnac, Thierry and David Thesmar. Identifying Dynamic Discrete Decision Processes. Econometrica, 70(2):801–816, 2002.
- Man, Elaine Lau Ying and Adam 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 Thomas Renström. Time Inconsistency in Environmental Policy: Tax Earmarking as a Commitment Solution. Economic Journal, 110(462):C123–C138, 2000.
- Nordhaus, William. Climate Clubs: Overcoming Free-riding in International Climate Policy. American Economic Review, 105(4):1339–1370, 2015.
- Oates, Wallace and Paul Portney. The Political Economy of Environmental Policy. Handbook of Environmental Economics, 1:325–354, 2003.
- Official Journal of the European Union. European Parliament Resolution of 4 April 2017 on Palm Oil and Deforestation of Rainforests, 2018.
- Oliva, Paulina. Environmental Regulations and Corruption: Automobile Emissions in Mexico City. Journal of Political Economy, 123(3):686–724, 2015.
- Pakes, Ariel, Michael Ostrovsky, and Steven Berry. Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples). RAND Journal of Economics, 38(2):373–399, 2007.
- Palm Oil Analytics. Essential Palm Oil Statistics. Technical report, 2017.
- Parks, Richard. Efficient Estimation of a System of Regression Equations when Disturbances are Both Serially and Contemporaneously Correlated. Journal of the American Statistical Association, 62(318):500–509, 1967.
- Pesendorfer, Martin and Philipp Schmidt-Dengler. Asymptotic Least Squares Estimators for Dynamic Games. Review of Economic Studies, 75:901–928, 2008.
- Peters, Glen, Jan Minx, Christopher Weber, and Ottmar 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, Oxford, 1997.
- Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni. Country-Level Social Cost of Carbon. Nature Climate Change, 8:895–900, 2018.

- Ryan, Stephen. The Costs of Environmental Regulation in a Concentrated Industry. Econometrica, 80(3):1019–1061, 2012.
- Scott, Paul. Dynamic Discrete Choice Estimation of Agricultural Land Use. 2013.
- Shapiro, Joseph. The Environmental Bias of Trade Policy. 2020.
- Sofiyuddin, Muhammad, Arief 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, Matthew Hansen, Stephen Stehman, Peter Potapov, Alexandra Tyukavina, Eric Vermote, and John Townshend. 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 2019. Technical report, 2019a.
- USDA Foreign Agricultural Service. Malaysia Biofuels Annual Report 2019. Technical report, 2019b.
- Wagstaff, Kiri, Claire Cardie, Seth Rogers, and Stefan Schroedl. Constrained K-means Clustering with Background Knowledge. Proceedings of the Eighteenth International Conference on Machine Learning, 1:577–584, 2001.
- Warren, Matthew, Kristell Hergoualc’h, J. Boone Kauffman, Daniel Murdiyarso, and Randall Kolka. 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 Wildlife Fund. Palm Oil Buyers Scorecard: Measuring the Progress of Palm Oil Buyers. Technical report, 2016.
- Xu, Yidi, Le Yu, Wei Li, Philippe Ciais, Yuqi Cheng, and Peng Gong. Annual Oil Palm Plantation Maps in Malaysia and Indonesia from 2001 to 2016. Earth System Science Data, 12:847–867, 2020.
- Zarin, Daniel, Nancy Harris, Alessandro Baccini, Dmitry Aksenov, Matthew Hansen, Claudia Azevedo-Ramos, Tasso Azevedo, Belinda Margono, Ane Alencar, Chris Gabris, Adrienne Allegretti, Peter Potapov, Mary Farina, Wayne Walker, Varada Shevade, Tatiana Loboda, Svetlana Turubanova, and Alexandra Tyukavina. Can Carbon Emissions from Tropical Deforestation Drop by 50% in 5 Years? Global Change Biology, 22:1336–1347, 2016.

ONLINE APPENDIX

A Theory

Domestic regulation

Let the total inverse demand curve be $P_t^D(q)$. Social welfare depends on the path of new development $\{Q_t^n, Q_{t+1}^n, \dots\}$, as well as prior, old development Q_t^o , which is sunk. New development becomes old development by law of motion $Q_{t+1}^o = Q_t^n + Q_t^o$. For discount factor β ,

$$W_t(Q_t^n, Q_{t+1}^n, \dots; Q_t^o) = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t \left[\int_0^{Q_{t+s}^o} P_{t+s}^D(q) dq - \int_{Q_{t+s}^o}^{Q_{t+s}^o + Q_{t+s}^n} \left(P_{t+s}^S(q) + e \right) dq \right],$$

where $Q_{t+s}^o = Q_t^o + Q_t^n + Q_{t+1}^n + \dots + Q_{t+s-1}^n$. Domestic regulation targets new development with upfront tax $\tilde{\tau}_t$. In equilibrium, new development equalizes marginal cost and expected revenue.

$$P_t^S(Q_{t+1}^{o*}(\tilde{\tau}_t)) = \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t [P_{t+s}^D(Q_{t+s}^{o*}(\tilde{\tau}_t))] - \tilde{\tau}_t.$$

Assuming an interior solution $Q_t^{n*}(\tilde{\tau}_t) > 0$, the first order condition and resulting tax are

$$\frac{dW_t}{d\tilde{\tau}_t} = (\tilde{\tau}_t - e) \frac{dQ_t^n}{d\tilde{\tau}_t} = 0, \quad \tilde{\tau}_t^{\text{FB}} = e.$$

Upfront tax $\tilde{\tau}_t$ only directly affects new development Q_t^n , and I apply the envelope theorem to ignore second-order effects on future development.

The leakage problem

Suppose importers can impose target new development with upfront tariff $\tilde{\tau}_t$. Producers choose between regulated market r and unregulated market u . Social welfare is

$$\begin{aligned} W_t(Q_t^{rn}, Q_{t+1}^{rn}, \dots, Q_t^{un}, Q_{t+1}^{un}, \dots; Q_t^o) \\ = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t \left[\int_0^{Q_{t+s}^{ro}} P_{t+s}^{Dr}(q) dq + \int_0^{Q_{t+s}^{uo}} P_{t+s}^{Du}(q) dq - \int_{Q_{t+s}^o}^{Q_{t+s}^o + Q_{t+s}^n} \left(P_{t+s}^S(q) + e \right) dq \right]. \end{aligned}$$

New development equalizes marginal cost and revenue and is indifferent across markets.

$$P_t^S(Q_{t+1}^{o*}(\tilde{\tau}_t)) = \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t [P_{t+s}^{Dr}(Q_{t+s}^{ro*}(\tilde{\tau}_t))] - \tilde{\tau}_t = \sum_{s=1}^{\infty} \beta^s \mathbb{E}_t [P_{t+s}^{Du}(Q_{t+s}^{uo*}(\tilde{\tau}_t))]$$

Assuming an interior solution, the first order condition and resulting tariff are

$$\frac{dW_t}{d\tilde{\tau}_t} = (\tilde{\tau}_t - e) \frac{dQ_t^{rn}}{d\tilde{\tau}_t} - e \frac{dQ_t^{un}}{d\tilde{\tau}_t} = 0, \quad \tilde{\tau}_t^{\text{L}} = \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du}} \right) e < \tilde{\tau}_t^{\text{FB}}, \quad (12)$$

for elasticities $\varepsilon_t^S > 0$ and $\varepsilon_{t+1}^{Du} < 0$ evaluated at quantities $Q_{t+1}^o \equiv Q_{t+1}^{o*}(\tilde{\tau}_t^L)$ and $Q_{t+1}^{uo} \equiv Q_{t+1}^{uo*}(\tilde{\tau}_t^L)$, respectively. Elasticity of regulated demand $\varepsilon_{t+1}^{Dr} < 0$ does not enter the tariff itself, although tariffs do have smaller effects on quantities and welfare as ε_{t+1}^{Dr} shrinks. If $Q_{t+1}^{uo} = 0$, then $\tilde{\tau}_t^L = \tilde{\tau}_t^{FB}$.

Elastic unregulated demand leads to severe leakage and pushes tariffs to zero. The same is true for inelastic supply, as tariffs produce allocative inefficiency without reducing emissions. Conversely, leakage is limited under inelastic unregulated demand (or a small unregulated share of consumption), as unregulated demand does not increase meaningfully as world prices fall. Similarly, leakage is limited under elastic supply, as world prices do not fall to begin with.

The commitment problem

In reality, importers cannot target new development with upfront tariffs. Rather, they can only target consumption over time. This constraint has two consequences. First, under time to build, new development does not begin production until the next period. Thus, new development is unaffected by tariffs on consumption today, and instead governed by future tariffs $\{\tau_{t+1}, \tau_{t+2}, \dots\}$. Second, producers shift sales away from the regulated market as tariffs rise and toward it as tariffs fall. This shifting does not occur with upfront tariff $\tilde{\tau}_t^L$ because producers that have paid upfront have no further cost of selling to the regulated market and thus no incentive to shift sales.

To see the implications, it becomes convenient to rewrite social welfare as

$$\begin{aligned} W_t(Q_t^{ro}, Q_{t+1}^{ro}, \dots, Q_t^{uo}, Q_{t+1}^{uo}, \dots; Q_t^o) \\ = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t \left[\int_0^{Q_{t+s}^{ro}} P_{t+s}^{Dr}(q) dq + \int_0^{Q_{t+s}^{uo}} P_{t+s}^{Du}(q) dq - \int_{Q_{t+s}^o}^{Q_{t+s+1}^o} \left(P_{t+s}^S(q) + e \right) dq \right], \end{aligned}$$

with the following equilibrium conditions for all $s \geq 0$.

$$P_{t+s}^S(Q_{t+s+1}^{o*}(\tau)) = \sum_{s'=1}^{\infty} \beta^{s'} \mathbb{E}_t [P_{t+s+s'}^{Du}(Q_{t+s+s'}^{uo*}(\tau))], \quad P_{t+s}^{Dr}(Q_{t+s}^{ro*}(\tau)) - \tau_{t+s} = P_{t+s}^{Du}(Q_{t+s}^{uo*}(\tau)).$$

The first order condition and resulting tariff for $s = 0$ show the source of the commitment problem.

$$\frac{dW_t}{d\tau_t} = \tau_t \frac{dQ_t^{ro}}{d\tau_t} = 0, \quad \tau_t^{\text{NC}} = 0$$

From the perspective of time t , tariffs τ_t have no effect on new development because of time to build, and no effect on prior development because it is sunk. In the presence of leakage, tariffs distort the allocation of consumption across markets, and as such are set to zero. Importers that sequentially choose static optima in a no-commitment scenario will therefore never impose tariffs.

$$\tilde{\tau}_t^{\text{NC}} = 0$$

Without leakage, there is no such problem: $\frac{dQ_t^{uo}}{d\tau_t} = \frac{dQ_t^{ro}}{d\tau_t} = 0$, which satisfies the first order condition without setting tariffs to zero.

Under limited commitment, importers commit to tariffs for L periods at a time before resetting them. Indeed, Indonesia and China both conduct national planning under “five-year plans,” and the US revises many policies based on decennial census results. In each new commitment regime,

importers treat prior development as sunk and thus set the regime's initial tariffs to zero.

$$\tau_t^{\text{LC}} = \tau_{t+L}^{\text{LC}} = \tau_{t+2L}^{\text{LC}} = \dots = 0$$

The remaining tariffs are set anticipating these periodic breaks. I obtain closed-form expressions by assuming that the demand and supply curves are constant over time, but I relax this simplifying assumption in the empirical implementation by solving numerically.

Time-invariant demand and supply curves imply time-invariant tariffs. To see why, note that

$$\frac{dW_t}{d\tau_{t+s}} = [\beta\tau_{t+s} - (1-\beta)e] \frac{dQ_{t+s}^{ro}}{d\tau_{t+s}} - (1-\beta)e \frac{dQ_{t+s}^{uo}}{d\tau_{t+s}} = 0,$$

nesting $\frac{dW_t}{d\tau_{t+s}} = \tau_{t+s} \frac{dQ_{t+s}^{ro}}{d\tau_{t+s}} = 0$ given $\frac{dQ_{t+s}^{ro}}{d\tau_{t+s}} = 0$ for $s \in \{0, L, 2L, \dots\}$. But $\frac{dQ_{t+s}^{ro}}{d\tau_{t+s}}$ and $\frac{dQ_{t+s}^{uo}}{d\tau_{t+s}}$ are time-invariant because the demand and supply curves are time-invariant, and thus $\tau_{t+s} = \tau$ for all $s \notin \{0, L, 2L, \dots\}$. Furthermore, any response to announced tariffs will occur in the initial period. To see why, suppose not. Development in a later period must be profitable in that period, but if so then developing in the first period is more profitable: fixed demand, supply, and tariffs imply fixed flow profits, such that developing earlier increases profits.

Social welfare therefore depends only on two allocations of consumption across markets: that under zero-tariff periods and that under non-zero-tariff periods. The key mechanism is that these allocations differ because producers shift sales away from the regulated market where tariffs are in place, and toward the regulated market when they are not.

$$\begin{aligned} W_t(Q_{t+1}^{ro}, Q_{t+L}^{ro}, Q_{t+1}^{uo}, Q_{t+L}^{uo}; Q_t^o) \\ = \left(\frac{\beta}{1-\beta} - \frac{\beta^L}{1-\beta^L} \right) \left[\int_0^{Q_{t+1}^{ro}} P^{Dr}(q) dq + \int_0^{Q_{t+1}^{uo}} P^{Du}(q) dq \right] \\ + \frac{\beta^L}{1-\beta^L} \left[\int_0^{Q_{t+L}^{ro}} P^{Dr}(q) dq + \int_0^{Q_{t+L}^{uo}} P^{Du}(q) dq \right] - \int_{Q_t^o}^{Q_{t+1}^o} (P^S(q) + e) dq, \end{aligned}$$

with $(Q_{t+1}^{ro}, Q_{t+1}^{uo})$ when tariffs are in place and $(Q_{t+L}^{ro}, Q_{t+L}^{uo})$ when they are not. In equilibrium,

$$\begin{aligned} P^S(Q_{t+1}^{o*}(\tau)) &= \left(\frac{\beta}{1-\beta} - \frac{\beta^L}{1-\beta^L} \right) P^{Du}(Q_{t+1}^{uo*}(\tau)) + \frac{\beta^L}{1-\beta^L} P^{Du}(Q_{t+L}^{uo*}(\tau)), \\ P^{Dr}(Q_t^{ro*}(\tau)) - \tau_t &= P^{Du}(Q_t^{uo*}(\tau)) \quad \forall t, \text{ given } \tau_{t+s} = \begin{cases} 0 & \text{for } s \in \{0, L, 2L, \dots\} \\ \tau & \text{otherwise} \end{cases}, \end{aligned}$$

and $Q_{t+1}^{ro} + Q_{t+1}^{uo} = Q_{t+L}^{ro} + Q_{t+L}^{uo}$. The first order condition is

$$\frac{dW_t}{d\tau} = \left[\left(\frac{\beta}{1-\beta} - \frac{\beta^L}{1-\beta^L} \right) \tau - e \right] \frac{dQ_{t+1}^{ro}}{d\tau} - e \frac{dQ_{t+1}^{uo}}{d\tau}.$$

Assuming an interior solution, tariffs have net present value

$$\hat{\tau}_t^{\text{LC}}(L) = \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du} \left[1 + \frac{(1-\beta)\beta^L}{\beta-\beta^L} \frac{P_{t+L}^{Du}}{P_{t+1}^{Du}} \left(1 - \frac{Q_{t+1}^o \varepsilon_t^S}{Q_{t+L}^{ro} \varepsilon_{t+1}^{Dr} + Q_{t+L}^{uo} \varepsilon_{t+1}^{Du}} \right) \right]} \right) e,$$

for elasticities $\varepsilon_t^S > 0$ and $\varepsilon_{t+1}^{Dr}, \varepsilon_{t+L}^{Dr}, \varepsilon_{t+1}^{Du}, \varepsilon_{t+L}^{Du} < 0$, and quantities and prices evaluated at τ^{LC} .

$L \rightarrow \infty$ corresponds to full commitment, and $L = 2$ to the minimum binding level of commitment.

$$\tilde{\tau}_t^{\text{LC}}(L) < \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du}} \right) e = \lim_{L \rightarrow \infty} \tilde{\tau}_t^{\text{LC}}(L) = \tilde{\tau}_t^{\text{C}} = \tilde{\tau}_t^{\text{L}}.$$

The same mechanism applies in the more general case if tariffs are declining over time. Indeed, tariffs set with a sequential static approach will be governed by equations 12, which imply declining tariffs if the elasticity of supply is declining over time. Such will be the case when the marginal costs of development are increasing as development progresses from more suitable lands to less suitable lands. At the extreme, tariffs are set to zero once all lands are exhausted: at this point, tariffs cannot reduce emissions because prior development is sunk and new development is impossible.

How leakage and commitment interact

Leakage is most severe when unregulated demand is elastic or supply is inelastic. Leakage in turn affects the role of commitment L . First, $\tilde{\tau}_t^{\text{LC}}(L)$ increases more rapidly in L for smaller $|\varepsilon_{t+1}^{Du}|$.

$$\lim_{\varepsilon_{t+1}^{Du} \rightarrow 0} \tilde{\tau}_t^{\text{LC}}(L) = e > 0 = \lim_{\varepsilon_{t+1}^{Du} \rightarrow -\infty} \tilde{\tau}_t^{\text{LC}}(L)$$

Second, $\tilde{\tau}_t^{\text{LC}}(L)$ increases more rapidly in L for larger ε_t^S .

$$\lim_{\varepsilon_t^S \rightarrow 0} \tilde{\tau}_t^{\text{LC}}(L) = 0 < \left(\frac{1}{1 + \frac{(1-\beta)\beta^L}{\beta-\beta^L} \frac{P_{t+L}^{Du}}{P_{t+1}^{Du}} \frac{Q_{t+1}^{uo}}{Q_{t+1}^o} \varepsilon_{t+1}^{Du} \left(\frac{Q_{t+1}^o}{Q_{t+L}^{ro} \varepsilon_{t+1}^{Dr} + Q_{t+L}^{uo} \varepsilon_{t+1}^{Du}} \right)} \right) e = \lim_{\varepsilon_t^S \rightarrow \infty} \tilde{\tau}_t^{\text{LC}}(L)$$

Extensions

The baseline model treats emissions as homogeneous over space, but in reality there is spatial variation in carbon stocks. In the absence of leakage, the first-best regulation is Pigouvian, with higher tariffs for higher-emission goods. In practice, however, tracing goods to their emissions is difficult and imperfect. I therefore focus on a uniform tariff that treats all goods equally. These “second-best” uniform tariffs are weighted averages over emission levels as in [Diamond \(1973\)](#), with weights given by level-specific supply elasticities.

The baseline model also rules out terms-of-trade effects. This classic motivation for import tariffs arises because tariffs in large markets can change world prices and therefore improve terms of trade at the expense of other countries ([Johnson 1953](#)). The objective function in the baseline model is global social welfare, and so the regulator fully internalizes these effects by construction. Suppose instead that the regulator considers only consumer surplus in the regulated market alongside the emissions externality. In this case, the expression for tariffs includes an additional terms-of-trade term, although this term is dominated when the emissions externality is large.

B Data

B.1 Sources

Table B1: Palm oil plantations and mills

Source	Period	Coverage	Description
Xu et al. (2020)	2001-2016	Indonesia, Malaysia	Palm oil plantations over time, 100m resolution
Song et al. (2018)	1982-2016	World	Land cover change over time, 5.6km resolution
WRI Universal Mill List	2018	Indonesia, Malaysia	List of mill coordinates
CIFOR mill list	2017	Indonesia	List of mill coordinates
Economic census	2016	Indonesia	Palm oil firms by village
Malaysian Palm Oil Board	2016	Malaysia	Palm oil mills by region
Google Earth	1987-2018	Indonesia	Historical satellite images of mill coordinates

Table B2: Yields

Source	Period	Coverage	Description
WorldClim	1970-2000	World	Average monthly solar radiation and precipitation
World Bank INDO-DAPOER	1996-2010	Indonesia	Annual yields by province
Indonesian Ministry of Agriculture	2011-2017	Indonesia	Annual yields by province
Malaysian Palm Oil Board	1990-2018	Malaysia	Annual yields by state

Table B3: Land characteristics

Source	Period	Coverage	Description
World Port Index	2019	World	Port coordinates
World Port Source	2020	World	Port coordinates
Global Roads Inventory Project	2018	World	Road networks
Gumbricht et al. (2017)	2011	World	Peatlands and depth, 231m resolution
Zarin et al. (2016)	2000	World	Aboveground biomass, 30m resolution
Hansen et al. (2013)	2001-2018	World	Tree cover loss, 30m resolution

Table B4: Consumption and world prices

Source	Period	Coverage	Description
USDA Foreign Agricultural Service	1960-2019	World	Annual consumption, production, area harvested, imports, and exports by country and oilcrop
IMF, World Bank	1980-2019	World	Monthly prices by oilcrop
World Bank	1980-2019	World	Inflation
Global Meteorological Forcing Dataset	1980-2016	World	Daily precipitation and temperature, 28km resolution
Database of Global Administrative Areas	2018	World	GIS maps of administrative boundaries

B.2 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 a resolution of 100 meters 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 palm as well as both 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 the 30-arc-second resolution (approximately 1 km²) by averaging. As discussed in [Xu et al. \(2020\)](#), I impose that development is uni-directional, such that the proportion of development for each tile is non-decreasing over time. [Xu et al. \(2020\)](#) restrict their attention to Sumatra, Kalimantan, Riau, and Malaysia, and I do the same in my analysis. These regions cover virtually all palm production in Indonesia and Malaysia during the period of study, although Papua and Sulawesi remain important frontiers for future expansion.

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 plantation development, I estimate the empirical relationship between plantation development and tree cover loss during the period of overlap, and I use these estimates to impute plantation development prior to 2001. For tiles i in years t ,

$$\Delta\text{Plantation}_{it} = \sum_{s=0}^3 \beta_s \Delta\text{Tree cover}_{it-s} + \varepsilon_{it}, \quad (13)$$

where $\Delta\text{Plantation}_{it}$ is new plantation development and $\Delta\text{Tree cover}_{it-s}$ terms are tree cover loss in the preceeding periods. The [Song et al. \(2018\)](#) data are at 5.6-km resolution, so I downscale them to match the 1-km resolution of the aggregated [Xu et al. \(2020\)](#) data. Table B5 shows the resulting estimates: negative coefficients indicate that more plantation development corresponds to higher tree cover loss, especially over the preceeding two years. For each tile, I combine the predicted changes in plantation development with the observed levels in 2001 to impute pre-2001 plantation development, 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 a resolution of 5.6km, and indeed my core analysis analyzes aggregated sites and not individual tiles.

Table B5: Xu et al. (2020) plantation vs. Song et al. (2018) tree cover data, 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

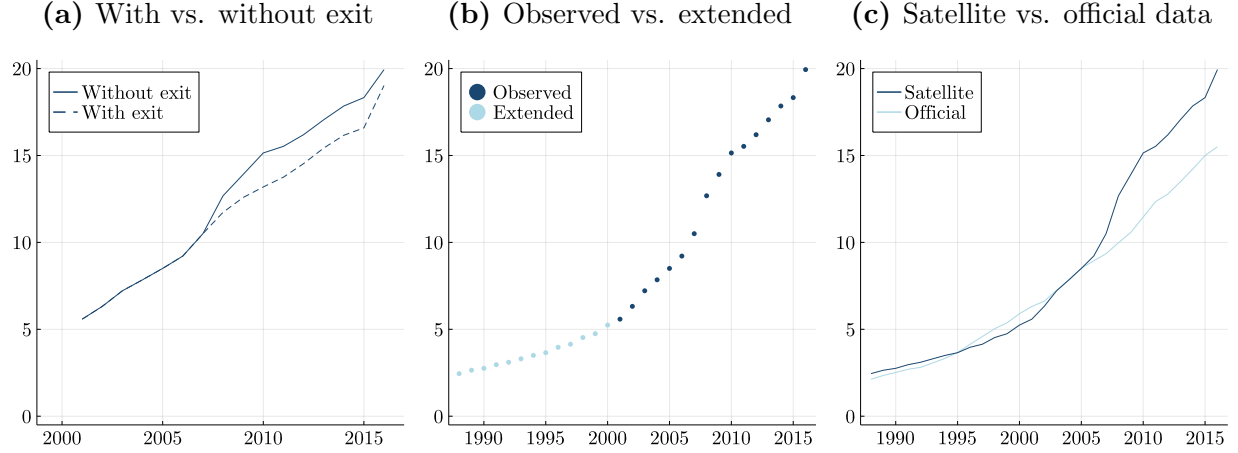
Each observation is a 30-arc-second tile in a given year, and each column is a regression. The dependent variable is from Xu et al. (2020), which measures the ratio of each tile that has been developed into palm oil plantations over time. The independent variables come from Song et al. (2018), which measures the ratio of each tile that is covered by tree canopy over time. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure B1 plots the resulting data. First, imposing uni-directional development rules out exit. Indeed, there is little exit in the data to begin with, and in any case plantation development releases carbon emissions irreversibly. 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 figures from government statistics. 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 (UML), a joint effort led by the World Resources Institute and Rainforest Alliance that collects data from palm oil processors, traders, corporate consumers, and NGOs. Mills are geocoded and manually verified by satellite. I combine these data with the 2017 Center for International Forestry Research (CIFOR) database, an independent effort that combs traceability reports for major palm oil processors and also verifies coordinates manually by satellite. I merge the datasets spatially, matching 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, and I use historical satellite images from Google Earth to determine the timing of mill construction. For each mill, I record the first year in which I observe mill construction.

In this way, I identify 1,521 palm mills as of 2016. I verify the data by comparing them to official government data from the Indonesian economic census and Malaysian Palm Oil Board.¹⁰ Table B6 shows that the total number of mills matches well, as does the overall spatial distribution. Discrepancies in regional counts are concentrated in the Indonesian data, where the census often records firm locations based on administrative offices and not milling facilities.

¹⁰ The 2016 Indonesian economic census contains 1,248 palm-oil establishments, of which 1,154 are located outside of Java. Focusing on firms involved in extracting crude oil from crops, I obtain 1,070 firms that produce either crude palm or palm kernel oil (KBLI codes 10431 and 10432, respectively).

Figure B1: Total plantations over time (Mha)

The left figure shows how imposing no exit affects the [Xu et al. \(2020\)](#) data. The middle figure shows the [Xu et al. \(2020\)](#) data in navy and the extended data in light blue, where I extend the data using tree cover data from [Song et al. \(2018\)](#). The right figure compares the satellite data to USDA FAS data.


Table B6: Mill counts by region, mill data vs. government figures

	Mill data	Government figures
Indonesia	1,050	1,070
Kalimantan	328	260
Central Sumatra	264	358
North Sumatra	225	237
South Sumatra	204	178
Sulawesi	21	30
Papua	8	7
Malaysia	471	453
Peninsular Malaysia	266	247
Sabah	132	129
Sarawak	73	77
Total	1,521	1,523

Mill data and government figures are both for 2016. Mill data come from the Universal Mill List and CIFOR. Indonesia government data come from the economic census, and Malaysian government data come from the Malaysian Palm Oil Board. Regions are in descending order by number of mills. Central Sumatra includes West Sumatra, Riau, and Kepulauan Riau; North Sumatra includes North Sumatra and Aceh; South Sumatra includes South Sumatra, Bangka Belitung, Bengkulu, Jambi, and Lampung.

I lightly harmonize to ensure consistency between the plantation and mill data. First, I assign plantations to the nearest mill in 2016, and I assume these assignments are consistent over time. Second, I drop plantations and mills that do not meet industry standards. Plantations must be within 50 kilometers of a mill, as oil palm fruit deteriorates rapidly after harvest and thus cannot be processed without nearby mills. Mills must have at least 1,000 hectares of plantations, which is the

Figure B2: Harmonized plantation and mill data over time



Light blue lines show unharmonized data, and navy lines harmonized data. Harmonization drops plantations and mills inconsistent with each other, and dashed light blue lines show the effects of dropping these data.

Table B7: Proportion of data impacted by harmonization

	All		Within province	
	Plantations	Mills	Plantations	Mills
Dropped (%)	1.83	0.91	2.06	1.06
Adjusted (%)	11.98	12.23	11.95	11.95
Total (%)	13.80	13.14	14.00	13.01

The table shows the proportion of plantations and mills affected by harmonization. I assign plantations to the nearest existing mill within 50 kilometers and – in the last two columns – within the same province. Harmonization adjusts the timing of plantation and mill investment to avoid plantations that predate their assigned mills, dropping data that cannot be reconciled in this way.

minimum required to run a small mill at capacity.¹¹ Third, I adjust the data to avoid plantations that pre-date their assigned mills.¹² I weight the plantation and mill data equally, which balances delaying plantation development against advancing mill construction.

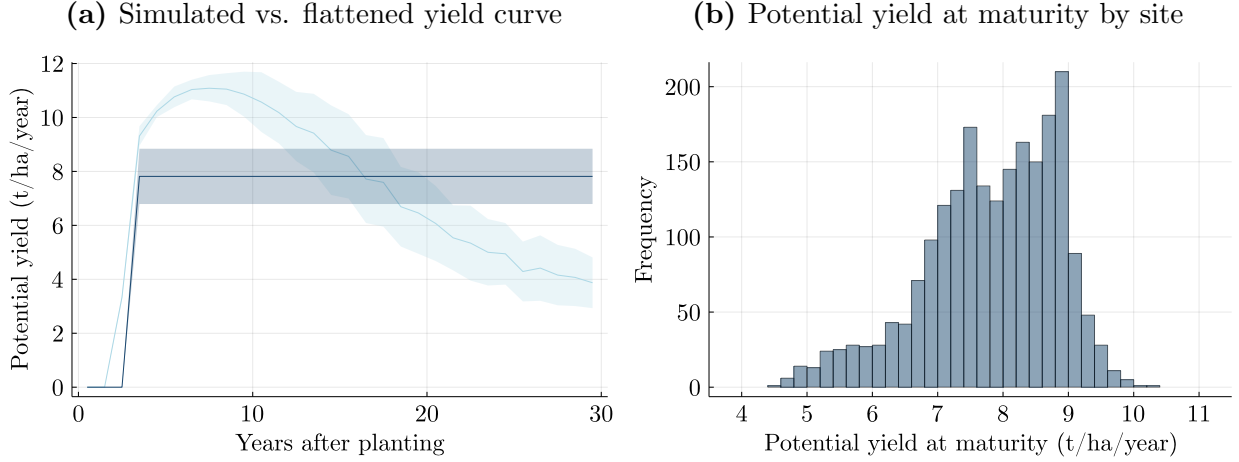
Figure B2 and table B7 show the modest impacts of harmonization. I further impose that plantations be linked to mills within the same province (Indonesia) or state (Malaysia). This assumption simplifies computation in defining potential sites because it allows me to define sites separately by region, and table B7 shows that it has little marginal effect.¹³

¹¹ Each year, 1,000 hectares with a yield of 3 tons of palm oil per hectare will produce 3,000 tons, matching the capacity of a small mill that processes 1 ton per hour for 10 hours per day for 300 days per year.

¹² The plantation data record when young palm trees have been established, and the mill data record when mill construction begins. Proximity to an under-construction mill ensures that young palm trees will have access to an operational mill by the time they reach maturity and begin to bear fruit.

¹³ There is also anecdotal support for plantations' staying within these borders to avoid licensing with multiple regional governments. Kuala Lumpur, Labuan, Perlis, and Putrajaya are small and lack yields data, so I combine them with neighboring states Selangor, Sabah, Kedah, and Selangor, respectively.

Figure B3: Potential palm oil yields



Yield curves are computed from the PALMSIM model (Hoffmann et al. 2014) using field-level average monthly solar radiation and precipitation from WorldClim. On the left, the light blue 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.

B.3 Yields

I construct data on palm oil yields by site over time by combining cross-sectional, site-level data on potential yields from the PALMSIM model of Hoffmann et al. (2014) with panel, province-level data on attained yields from government statistics.


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 from WorldClim, which measures these variables at a resolution of 30 arc-seconds. To facilitate computation, I aggregate climate inputs and run the PALMSIM model by site, as defined in section D.1. Figure B3a shows the resulting 30-year yield curve, which starts at zero before increasing steeply then declining gradually. Because the data on attained yields distinguish only between “immature” and “mature” crops, I flatten the curve to these levels while holding fixed the average yield over time. Figure B3b shows the variation in the flattened yields at maturity. These data are time-invariant because yields under optimal conditions are an inherent characteristic of the oil palm plant.

Second, I compile data on attained yields by province and year from government statistics, namely the Indonesian Ministry of Agriculture, the World Bank INDO-DAPOER database (via the Indonesian MoA), and the Malaysian Palm Oil Board. Each reports yields for mature crops, omitting immature crops that do not yet produce fruit. Figure B4a shows that, on average, these yields are increasing over time as technology improves, although attained yields fall far short of the maximum potential yields in all provinces and years.¹⁴ Across provinces and years, the average observed annual yield per hectare is 3.30 tons.

Lastly, I combine these data to produce estimates of attainable yields by site and year. Suppose

¹⁴ Crop age mix also affects yields over time. However, two effects potentially offset each other: young crops approaching their peak have increasing yields, while aging crops past their peak have decreasing yields.

Figure B4: Attained and attainable palm oil yields over time (t/ha/year)



On the left, each observation is the annual attained yield for a given province (Indonesia) or state (Malaysia) as recorded in government statistics. Data come from the Indonesian Ministry of Agriculture, World Bank INDO-DAPOER, and Malaysian Palm Oil Board. On the right, each observation is the annual attainable yield for a given site computed by combining site-level potential yields from PALMSIM with province-year-level attained yields from government statistics. For both, fitted lines show common time trends accounting for province/state fixed effects.

the desired attainable yields Y_{it} in sites i and years t are products of site-specific, time-invariant potential yields Y_i^p and province-specific, time-varying yield gaps γ_{mt} .

$$Y_{it} = (1 - \gamma_{mt})Y_i^p \quad (14)$$

The underlying restriction is that, while potential yields are allowed to vary by site, yield gaps are fixed across sites in a given province-year. Yield gaps are a function of known quantities.

$$\frac{\sum_{i \in \mathcal{I}_m} Y_{it} d_{it}}{\sum_{i \in \mathcal{I}_m} d_{it}} = Y_{mt} \quad \Rightarrow \quad \gamma_{mt} = 1 - Y_{mt} \left(\frac{\sum_{i \in \mathcal{I}_m} Y_i^p d_{it}}{\sum_{i \in \mathcal{I}_m} d_{it}} \right)^{-1},$$

where attained yields Y_{mt} , potential yields Y_i^p , and plantation development d_{it} are known. I isolate the underlying levels and trends of these yield gaps with the specification


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

and I use the fitted values to estimate attainable yields Y_{it} with equation 14. In doing so, I extrapolate back before 1990 for Malaysia and 1996 for Indonesia. I also extrapolate past 2016 to obtain future yields, which I use in computing counterfactuals. Figure B4b shows the resulting estimates, which combine the uptrend of figure B4a with the site-level dispersion of figure B3b.

B.4 Carbon stocks

I compute carbon stocks over space using two datasets, which I aggregate to a resolution of 30 arc-seconds: [Zarin et al. \(2016\)](#) measures aboveground tree biomass at a resolution of 30m, and [Gumbricht et al. \(2017\)](#) measures belowground peat biomass at a resolution of 231m. Plantation

Figure B5: Plantation development vs. CO₂ emissions over time



Data on plantation development come from [Xu et al. \(2020\)](#) and [Song et al. \(2018\)](#), and data on carbon emissions from [Zarin et al. \(2016\)](#) and [Gumbricht et al. \(2017\)](#).

development releases both. To convert aboveground biomass to carbon, I use a biomass-to-carbon conversation factor of 0.5. To convert belowground biomass, I use the conversation factor of 65.1 kg C/m³ peat in [Warren et al. \(2017\)](#). I convert carbon to carbon dioxide emissions with a molecular-weight conversion factor of 3.67. 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. Palm oil production also involves the release of methane and nitrous oxide, but precise estimates of these emissions are not yet well established.

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 in the [Song et al. \(2018\)](#) data. For peat deposits, bias is limited because [Gumbricht et al. \(2017\)](#) rely primarily on precipitation and topography – predetermined features – in order 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 B5 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.

B.5 Weather shocks to 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-pixel 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 B8 lists these countries for each oil (aggregating EU countries).

Table B8: 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%

Data are from the USDA Foreign Agricultural Service. Production pools over the study period (1988-2016), and for each oil I omit producers accounting for less than 5% of world production.

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 both the USDA and local government sources. I define the growing season for rapeseed, soybean, and sunflower oils to be those specified by country-specific crop calendars from the USDA, 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-pixel level. For rainfall, I first aggregate from daily to monthly values for each pixel, as daily variation in rainfall is not detrimental to crop growth in the same way that daily variation in temperatures can be. I then compute shocks as absolute deviations from optimal levels for each crop. The FAO Crop Ecological Requirements Database records optimal windows by crop for both rainfall and temperature, and I take the midpoint of these windows as optimal levels. The FAO database 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 pixel.

Third, I aggregate to obtain annual weather shocks by oil. I do so by averaging over pixels for each oil-producing region, then averaging 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, as annual production is a direct function of annual weather. In this final step, I can isolate foreign shocks for each consumer market by omitting shocks to domestic oil-producing regions, and I do so in checking robustness.

C Demand

In estimating the lower-level demand system, I impose the standard adding-up, homogeneity, and symmetry restrictions. The adding-up restrictions are $\sum_i \alpha_i^0 = 1$, $\sum_i \alpha_i^1 = 0$, $\sum_i \beta_i = 0$, $\sum_i \gamma_{ij} = 0 \forall j$ and are automatically satisfied since expenditure shares sum to one. Homogeneity imposes $\sum_j \gamma_{ij} = 0 \forall i$, such that proportional changes in prices and income have no impact on demand. Symmetry imposes $\gamma_{ij} = \gamma_{ji} \forall i, j$. Given a choice between two products – palm vs. other oils – imposing homogeneity imposes symmetry, and vice versa. A choice between two products also allows me to estimate the demand system on palm oil expenditure shares alone, as the adding-up restriction requires the dropping of one product. Thus, I apply linear IV and use Newey-West standard errors to account for serial correlation in the error terms. The typical case with more than two products applies seemingly unrelated regression to estimate a system of regression equations. Serial correlation can then be accounted for with a Prais-Winsten transformation as in [Parks \(1967\)](#).

On instruments, table [C1](#) shows that weather shocks do not affect domestic incomes or expenditures for any consumer market. Such effects would influence demand directly – as opposed to through the channel of oil prices – and therefore violate the exclusion restriction. These results also provide reassurance that the instruments do not simply capture idiosyncratic fluctuations in macroeconomic conditions. Table [C2](#) shows the first stage for foreign weather shocks, which are also strong instruments. In omitting domestic shocks within a given consumer market, these instruments go one step further toward avoiding violations of the exclusion restriction. However, the baseline analysis favors the use of all weather shocks because they greatly simplify the construction of aggregate demand curves.¹⁵ Furthermore, the baseline instruments already target oil producers explicitly, and they pass the test above.

Table [C3](#) presents demand elasticities for palm oil by market. Table [C4](#) shows the lower- and upper-level parameter estimates that I use to compute these elasticities, and table [C5](#) shows demand elasticities for vegetable oils in general. I obtain reasonable estimates with negative own-price elasticities that are statistically significant and positive cross-price elasticities. For Malaysia, elasticities for other oils have larger standard errors because other oils account for only 3% of consumption in the data. Table [C6](#) shows elasticities computed without price instruments, indicating clear bias in the form of positive own-price and negative cross-price elasticities, some of which are statistically significant. Figure [C1](#) indicates why, with co-movement in palm and other oil prices effectively dampening observed price variation. Instruments leverage differential weather shocks across oils, and indeed the instrumented price series are much less correlated.

Finally, I observe oil stockpiles and find that they are limited in this context. In particular, stockpiles are 12.5% of average annual consumption by volume, compared to an estimated 342% of average weekly consumption for ketchup in [Erdem et al. \(2003\)](#) and 188% of median weekly consumption for laundry detergent in [Hendel and Nevo \(2006\)](#). Temporal aggregation explains the difference: the vegetable oil data measure annual consumption, and substitution across years may be less salient than substitution across weeks for consumer products sold for regular discounts. As well, national consumption aggregates over the stockpiling of individual consumers.

¹⁵ For example, to estimate demand for the combined Indonesian-Malaysian market, I can aggregate their consumption data then estimate an aggregate curve directly. With foreign weather shocks, I must estimate separate curves for Indonesia and Malaysia then aggregate the curves themselves. Each demand curve relies heavily on the AIDS functional form at its extremes, and aggregating curves exacerbates this problem, particularly for markets with different consumption levels. Aggregating curves is also theoretically inconsistent with AIDS microfoundations.

Table C1: Weather shocks vs. incomes and expenditures

Market	Outcome	Rainfall		Temperature		Obs
		Estimate	SE	Estimate	SE	
European Union	CPI	0.00362	(0.00275)	0.00264	(0.00245)	174
	GDP	0.00530	(0.00762)	0.00408	(0.00736)	174
	GDE	0.00587	(0.00783)	0.00437	(0.00748)	174
	GDE (hh)	0.000190	(0.000257)	0.000147	(0.000245)	174
	GDE (gov)	0.000241	(0.000303)	0.000169	(0.000292)	174
China/India	CPI	0.00632	(0.0109)	0.00346	(0.0113)	174
	GDP	8.10e-05	(0.0103)	-0.00344	(0.00986)	174
	GDE	-0.00163	(0.00969)	-0.00434	(0.00922)	174
	GDE (hh)	-5.51e-05	(0.000343)	-0.000148	(0.000327)	174
	GDE (gov)	4.56e-05	(0.000281)	-6.68e-05	(0.000263)	174
Other importers	CPI	0.00571	(0.00776)	0.000995	(0.00787)	174
	GDP	0.00360	(0.00448)	0.00180	(0.00411)	174
	GDE	0.00429	(0.00415)	0.00235	(0.00373)	174
	GDE (hh)	0.000138	(0.000130)	8.12e-05	(0.000117)	174
	GDE (gov)	0.000181	(0.000182)	9.07e-05	(0.000162)	174
Indonesia/Malaysia	CPI	-0.0231	(0.0246)	-0.0221	(0.0242)	174
	GDP	0.0113	(0.0154)	0.00539	(0.0157)	174
	GDE	0.00920	(0.0147)	0.00424	(0.0152)	174
	GDE (hh)	0.000384	(0.000536)	0.000202	(0.000555)	174
	GDE (gov)	0.000283	(0.000769)	5.96e-05	(0.000798)	174

Each row is a regression. Data are annual and cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. For outcome variables, GDPs and GDEs are in logs, GDEs measure total, household, and government expenditures, 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. I control for oil fixed effects and oil-specific time trends, which differentiate between palm and other oils. Newey-West standard errors account for serial correlation. Significance levels: ** $p < 0.05$, * $p < 0.1$.

Table C2: Foreign weather shocks as price instruments

	European Union	China	India	Other importers	Indonesia	Malaysia
Rainfall shocks (100 mm)	0.000499 (0.0419)	0.217*** (0.0179)	0.197*** (0.0307)	0.111** (0.0443)	0.185*** (0.0236)	0.199*** (0.0297)
Temperature shocks (°C)	0.150*** (0.0523)	0.343*** (0.0178)	0.275*** (0.0356)	0.240*** (0.0514)	0.295*** (0.0302)	0.300*** (0.0327)
Observations	174	174	174	174	174	174
F-statistic	12.76	200.5	30.12	12.22	48.22	45.83

Each column is a regression, and the outcome variable is log prices. Data are annual and cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. Foreign weather shocks are absolute deviations from optimal conditions during the growing season, aggregated over foreign producing regions. I control for oil fixed effects and oil-specific time trends, which differentiate between palm and other oils. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C3: Demand elasticities for palm oil

Market	Estimate	SE	Market	Estimate	SE
European Union (E)	-0.510***	(0.181)	Importers (ECNR)	-0.555***	(0.134)
China/India (CN)	-0.667***	(0.210)	Producers (IM)	-0.026	(0.171)
Other importers (R)	-0.558***	(0.134)	EU/China/India (ECN)	-0.437***	(0.164)
Indonesia/Malaysia (IM)	-0.026	(0.171)	Not EU/China/India (RIM)	-0.482***	(0.129)
World (ECNRIM)	-0.447***	(0.133)	Not EU (CNRIM)	-0.602***	(0.113)

Each row of each table shows the palm oil demand elasticity for an individual or group of consumer markets. I present mean elasticities over the study period, and I compute standard errors with the delta method. The demand estimation underlying these elasticities draws on annual data that cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. It instruments for prices with weather shocks to oil production, and it accounts for serial correlation with Newey-West standard errors. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C4: Demand parameter estimates

Parameter	European Union		China/India		Other importers		Indonesia/Malaysia	
	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
α_1^0	0.162	(0.155)	0.328*	(0.168)	0.345***	(0.069)	0.662***	(0.127)
α_1^1	0.003***	(0.001)	0.004	(0.003)	0.005***	(0.001)	0.009***	(0.002)
γ_{11}	0.038	(0.026)	0.027	(0.030)	0.017	(0.016)	0.022	(0.029)
β_1	0.012	(0.029)	0.035	(0.032)	0.033***	(0.011)	-0.025	(0.024)
γ	-0.198	(0.122)	-0.416	(0.339)	-0.090	(0.235)	0.215	(0.159)

Each pair of columns is a demand system, and subscript $i = 1$ refers to palm oil. The first four rows describe the lower level of demand, and the last row the upper level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C5: Demand elasticities for vegetable oils


Market		Estimates		SEs	
		Palm	Other	Palm	Other
European Union	Palm	-0.510***	0.290	(0.181)	(0.206)
	Other	0.105	-0.301*	(0.148)	(0.171)
China/India	Palm	-0.667***	0.172	(0.210)	(0.302)
	Other	0.187	-0.584***	(0.159)	(0.224)
Other importers	Palm	-0.558***	0.454**	(0.134)	(0.180)
	Other	0.350***	-0.436***	(0.113)	(0.149)
Indonesia/Malaysia	Palm	-0.026	0.234*	(0.171)	(0.120)
	Other	0.707	-0.416	(0.474)	(0.478)

Each panel shows uncompensated price elasticities for a consumer market. I present mean elasticities over the study period, and I compute standard errors with the delta method. The demand estimation underlying these elasticities draws on annual data that cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. It instruments for prices with weather shocks to oil production, and it accounts for serial correlation with Newey-West standard errors. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C6: Demand elasticities for vegetable oils without price instruments


Market		Estimates		SEs	
		Palm	Other	Palm	Other
European Union	Palm	-0.075	0.018	(0.116)	(0.150)
	Other	-0.347**	0.196	(0.149)	(0.184)
China/India	Palm	0.606	-0.113	(0.693)	(0.556)
	Other	0.850**	-0.617*	(0.342)	(0.359)
Other importers	Palm	-0.484***	0.224	(0.051)	(0.143)
	Other	-0.279**	-0.139	(0.135)	(0.221)
Indonesia/Malaysia	Palm	0.730*	-0.685*	(0.424)	(0.403)
	Other	0.417	-0.477	(0.576)	(0.518)

Each panel shows uncompensated price elasticities for a consumer market. I present mean elasticities over the study period, and I compute standard errors with the delta method. The demand estimation underlying these elasticities draws on annual data that cover coconut, olive, palm, rapeseed, soybean, and sunflower oils from 1988 to 2016. It does not instrument for prices, but it does account for serial correlation with Newey-West standard errors. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure C1: Vegetable oil prices over time (USD/t)

Data on vegetable oil prices come from the International Monetary Fund and the World Bank. Palm oils aggregate palm and palm kernel, and other oils aggregate coconut, olive, rapeseed, soybean, and sunflower. I aggregate with a Stone price index, drawing on expenditure shares computed with data from the USDA Foreign Agricultural Service. The left figure shows observed prices, and the right figure shows predicted prices using weather shocks to oil production as instruments.

Figure D1: Potential sites



Blue shading indicates different potential sites, and gray shading indicates omitted regions. Orange dots are palm oil mills observed by 2016. There are 2,135 sites and 1,467 observed mills.

D Supply

D.1 Defining sites

To divide land into sites, I first compute the maximum number of sites \bar{k} for each province: $\bar{k} = \max\{\text{floor}(\text{area}/521), \text{number of observed mills}\}$. I use a benchmark site size of 521 km², which I obtain as the average of three calculations. First, I consider provinces with high mill density. At the 75th percentile, there is one mill per 455 km². Second, I consider provinces without mill construction in the last five years of the study period, reflecting plateaued expansion. The median such province has one site per 553 km². These two methods thus imagine bringing site density for all provinces to that of the most developed provinces. A third method considers circular sites that reflect the upper end of plantation-mill distances observed in the data. The 75th percentile of these distances implies radii of 13.3 km and site sizes of 553 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 do so with a version of the constrained k -means clustering algorithm described in [Wagstaff et al. \(2001\)](#), and 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 significant plantations, reassign points to clusters with mills.

Step (2) ensures consistency with observed mills, and step (6) observed plantations. In step (6), I define clusters with more than 10 30-arc-second tiles of plantations as having “significant” plantations. I drop the 0.3% of plantations that remain unassigned to clusters with mills. A lower cutoff would drop fewer plantations at the cost of losing more clusters. This procedure results in 2,135 sites, of which 1,467 contain an observed mill by 2016. Figure D1 plots the potential sites.

D.2 Extensive margin (mill construction)

Lemma 1. $v^e(0; \mathbf{w}_{it}) - v^e(0, 1; \mathbf{w}_{it}) = -\beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})]$.

$$\begin{aligned} v^e(0; \mathbf{w}_{it}) - v^e(0, 1; \mathbf{w}_{it}) &= \beta \mathbb{E}_{it}^e[\ln(\exp(v^e(0; \mathbf{w}_{it+1})) + \exp(v^e(1; \mathbf{w}_{it+1})))] - \beta \mathbb{E}_{it}^e[v^e(1; \mathbf{w}_{it+1})] \\ &= \beta \mathbb{E}_{it}^e[v^e(1; \mathbf{w}_{it+1}) - \ln p^e(\mathbf{w}_{it+1})] - \beta \mathbb{E}_{it}^e[v^e(1; \mathbf{w}_{it+1})] \\ &= -\beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})]. \end{aligned}$$

The first line applies the logit log-sum formula for expected utilities, and the second line applies the expression for logit choice probabilities. Arcidiacono and Ellickson (2011) document this result as the logit special case of Arcidiacono and Miller (2011) Lemma 1.

Lemma 2. $v^e(1; \mathbf{w}_{it}) - v^e(1, a_{it}; \mathbf{w}_{it}) = \frac{1}{2} \mathbb{E}_{it}^e[c''(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it})^2]$.

$$\begin{aligned} &v^e(1; \mathbf{w}_{it}) - v^e(1, a_{it}; \mathbf{w}_{it}) \\ &= \mathbb{E}_{it}^e[-c(a_{it}^*; \mathbf{w}_{it}, \varepsilon_{it}) + c(a_{it}; \mathbf{w}_{it}, \varepsilon_{it}) + \beta V(a_{it}^*; \mathbf{w}_{it+1}, \varepsilon_{it+1}) - \beta V(a_{it}; \mathbf{w}_{it+1}, \varepsilon_{it+1})] \\ &= \mathbb{E}_{it}^e\left[-c'(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it}) - \frac{1}{2}c''(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it})^2 + \beta V'(a_{it}; \mathbf{w}_{it+1}, \varepsilon_{it+1})(a_{it}^* - a_{it})\right] \\ &= \mathbb{E}_{it}^e\left[-c'(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it}) - \frac{1}{2}c''(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it})^2 + c'(a_{it}^*; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it})\right] \\ &= \frac{1}{2} \mathbb{E}_{it}^e[c''(a_{it}; \mathbf{w}_{it}, \varepsilon_{it})(a_{it}^* - a_{it})^2], \end{aligned}$$

where $a_{it}^* \equiv a_{it}^*(0; \mathbf{w}_{it}, \varepsilon_{it})$. The first equality is definitional. The second equality applies that costs are quadratic and revenues linear. The third equality applies the linearity of revenues and the first order condition that holds at a_{it}^* . The last equality again applies that costs are quadratic, and thus that c' is linear. For convex costs, the last line is positive, and indeed $v^e(1; \mathbf{w}_{it}) \geq v^e(1, a_{it}; \mathbf{w}_{it})$.

Result. $v^e(0; \mathbf{w}_{it}) - v^e(0, 1, a'_{it+1}; \mathbf{w}_{it}) = \frac{1}{2} \beta \mathbb{E}_{it}^e[c''(a'_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})(a_{it+1}^* - a'_{it+1})^2] - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})]$.

$$\begin{aligned} v^e(0; \mathbf{w}_{it}) - v^e(0, 1, a'_{it+1}; \mathbf{w}_{it}) &= v^e(0, 1; \mathbf{w}_{it}) - v^e(0, 1, a'_{it+1}; \mathbf{w}_{it}) - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})] \\ &= \beta \mathbb{E}_{it}^e[v^e(1; \mathbf{w}_{it+1})] - \beta \mathbb{E}_{it}^e[v^e(1, a'_{it+1}; \mathbf{w}_{it+1})] - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})] \\ &= \frac{1}{2} \beta \mathbb{E}_{it}^e[c''(a'_{it+1}; \mathbf{w}_{it+1}, \varepsilon_{it+1})(a_{it+1}^* - a'_{it+1})^2] - \beta \mathbb{E}_{it}^e[\ln p^e(\mathbf{w}_{it+1})], \end{aligned}$$

where $a_{it+1}^* \equiv a_{it+1}^*(0; \mathbf{w}_{it+1}, \varepsilon_{it+1})$. The first line substitutes Lemma 1, the second line is definitional, and the third line substitutes Lemma 2.

Table D1: Supply model parameter estimates by region

	Intensive-margin		Extensive-margin	
	κ_m	α_m	κ_m^e	α_m^e
Indonesia				
Aceh	6,200	-495	73,929,931	-533,741
Bengkulu	7,106	-528	107,774,316	1,308,477
Jambi	5,796	-550	27,216,255	413,225
Kalimantan Barat	-889	-955	14,705,348	-1,430,757
Kalimantan Selatan	1,185	-770	40,092,506	-1,176,381
Kalimantan Tengah	-638	-1,278	6,931,542	-1,733,589
Kalimantan Timur	-267	-1,212	70,702,346	-209,553
Lampung	-414	-1,104	28,299,191	-3,974,791
Riau, Kepulauan Riau	6,374	-667	39,297,091	-78,075
Sumatera Barat	5,659	-519	57,185,480	1,317,423
Sumatera Selatan, Bangka Belitung	5,061	-691	24,672,070	-980,322
Sumatera Utara	6,422	-717	42,374,773	-141,493
Malaysia				
Johor, Kuala Lumpur, Melaka, Negeri Sembilan, Putrajaya, Selangor	3,463	-977	75,255,698	1,461,261
Kedah, Perak, Penang, Perlis	5,913	-934	49,142,760	2,372,118
Kelantan, Pahang, Terengganu	3,677	-866	25,994,696	595,175
Sabah, Labuan	8,699	-837	26,272,959	-383,121
Sarawak	721	-1,113	-14,342,872	-1,908,711

Estimates are interpretable in terms of inflation-adjusted, year-2000 dollars. Region-specific costs κ_m and κ_m^e are for a mean year at mean values for cost factors. Parameters α_m and α_m^e denote region-specific cost trends. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D.3 Estimates

Table D1 presents the regional estimates underlying the means in table 4. Regions include provinces in Indonesia and states in Malaysia, grouped to maintain at least 25 sites each. I combine island provinces Kepulauan Riau and Bangka Belitung with nearby mainland provinces in Indonesia, and I do the same for island territory Labuan in Malaysia. I also group the peninsular Malaysian states, which are relatively small, into three regions. Low and falling costs rationalize high production in Sarawak despite relative isolation and low yields, as well as the particularly rapid expansion of production across Kalimantan.

E Counterfactuals

E.1 Solving the model

Firms play a dynamic entry game taking tariffs as given. Boldfaced $\mathbf{s}_{it} = \{s_{it}, s_{it}^e\}$ and $\mathbf{a}_{it} = \{a_{it}, p_{it}^e\}$ denote site-specific supply and entry. Total supply and entry in period t are

$$s_t = \sum_i Y_{it} s_{it}, \quad a_t = \sum_i \left(s_{it}^e a_{it} + (1 - s_{it}^e) p_{it}^e a_{it} \right). \quad (15)$$

For sites without mills ($s_{it}^e = 0$), entry depends on mill construction probability p_{it}^e and plantation development a_{it} . Entry affects future supply and thus prices.

$$s_{t+1}(a_t, s_t) = s_t + a_t, \quad P(s_{t+1}(a_t, s_t), d_{t+1}, \tau_{t+1}) \quad (16)$$

Prices depend on supply, demand, and tariffs. Only total supply enters; tracking supply over space is much harder computationally. Plantation development and mill construction probabilities are

$$a_{it} = \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta} \mathbb{E}_{it} \left[Y_{it+t'} P(s_{t+t'}, d_{t+t'}, \tau_{t+t'}) \right] - \frac{1}{\delta} x_i \gamma - \frac{1}{\delta} \kappa_m - \frac{1}{\delta} \alpha_m t - \frac{1}{\delta} \varepsilon_{it}, \quad (17)$$

$$p_{it}^e = \frac{\exp \left(-x_i \gamma^e - \kappa_m^e - \alpha_m^e t + \mathbb{E}_{it}^e [V(0; \mathbf{w}_{it}, \varepsilon_{it})] \right)}{1 + \exp \left(-x_i \gamma^e - \kappa_m^e - \alpha_m^e t + \mathbb{E}_{it}^e [V(0; \mathbf{w}_{it}, \varepsilon_{it})] \right)}, \quad (18)$$

subject to bounds $a_{it} \in [0, \bar{s}_i - s_{it}]$ and an outside option normalized to zero. Entry depends on expected future supply, demand, and yields, as well as current cost shocks.

The residuals of equation 8 capture future expectations and current cost shocks. For residuals $v_{it} = -\frac{1}{\delta} \varepsilon_{it} + \frac{\beta}{\delta} \varepsilon_{it+1} + \eta_{it}$, I compute the discounted sum over time and apply equation 9 to obtain

$$\tilde{v}_{it} \equiv \sum_{t'=0}^{\infty} \beta^{t'} v_{it+t'} = -\frac{1}{\delta} \varepsilon_{it} + \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta} \left(\mathbb{E}_{it} [Y_{it+t'} P_{t+t'}] - Y_{it+t'} P_{t+t'} \right).$$

I can then rewrite equation 17 in terms of observed values.

$$a_{it} = \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta} Y_{it+t'} P(s_{t+t'}, d_{t+t'}, \tau_{t+t'}) - \frac{1}{\delta} x_i \gamma - \frac{1}{\delta} \kappa_m - \frac{1}{\delta} \alpha_m t + \tilde{v}_{it}.$$

I approximate long-term yields, demand, and residuals by extrapolating from the study period. Discounting limits the resulting bias. I assume that yields grow at fixed, province-specific rates, that demand evolves along a sigmoid curve fit to observed values, and that out-of-sample residuals are given by their site-specific, in-sample means. I hold each fixed in computing counterfactuals. On the extensive margin, firms invest before realizing intensive-margin cost shocks, which are uncorrelated with extensive-margin shocks.

The entry game determines supply. In steady-state period S , all lands are exhausted and there is no further entry. But the steady state is reached only asymptotically, and full backward induction is infeasible for large S . I address this challenge by iterating on two dimensions. In an outer loop, I solve assuming no further entry after $T < S$. Doing so lightens computation but introduce bias

Table E1: Acacia vs. palm oil plantation development

	Acacia	Acacia	Acacia	Acacia
Palm development (ha)	0.0221*** (0.00682)	0.0155** (0.00776)	0.0113 (0.00782)	0.0115 (0.00792)
Site FE		x	x	x
Year FE			x	
Province-year FE				x
Observations	6,018	6,018	6,018	6,018

Data on acacia and palm plantations come from [Gaveau et al. \(2019\)](#) and cover the island of Borneo. Standard errors are clustered by site. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

in ignoring post- T entry. My solution is to increase T and re-solve until convergence. Intuitively, entry today is less appealing when future competitors have a longer window of opportunity to enter, but discounting leads to a diminishing marginal impact of extending this window.

In an inner loop, I solve iteratively for entry a_t (and site-specific $\{a_{it}\}$) given supply s_t and T . Rather than looking ahead to T , as in full backward induction, I solve with limited look-ahead.

1. Initialize by forward iteration. Compute $\{a_t, s_{t+1}\}, \{a_{t+1}, s_{t+2}\}, \dots, \{a_T, s_{T+1}\}$ sequentially, at each point accounting for preceding periods but ignoring future entry.
2. Revise by backward iteration. Re-compute $\{a_{T-1}, s_T\}$ given $\{s_{T-1}, s_T(a_{T-1}), s_{T+1}\}$, then $\{a_{T-2}, s_{T-1}\}$ given $\{s_{T-2}, s_{T-1}(a_{T-2}), s_T, s_{T+1}\}$, and so on until $\{a_t, s_{t+1}\}$. Then re-compute $\{a_T, s_{T+1}\}$ given $\{s_T, s_{T+1}(a_T)\}$, noting $s_{T+t'} = s_{T+t'}(a_T)$ for all $t' \geq 1$.
3. Repeat (2) until convergence.

To compute future supply s_{t+1} is simply to apply law-of-motion equation 16. To compute entry a_t is to solve equations 15, 17, and 18, which characterize a fixed point in which beliefs are consistent with actions in equilibrium. If all firms enter today, then future prices will be low and some firms are better off not entering; if no firm enters, then future prices will be high and some firms are better off entering. Prices decrease monotonically in total entry, such that the solution is unique and standard root-finding algorithms work well. Furthermore, as in [Hopenhayn \(1992\)](#), atomistic firms and the law of large numbers lead to expected extensive-margin entry that equals probability p_{it}^e . Otherwise, solving requires integrating over the distribution of potential entry outcomes.

E.2 Quantifying emissions

Data from [Gaveau et al. \(2019\)](#) measure acacia and palm oil plantations for the island of Borneo in five-year intervals from 1990 to 2015. For sites i , provinces m , and years t , I compare new palm and acacia development with the specification

$$\text{acacia}_{it} = \beta \text{palm}_{it} + \alpha_i + \alpha_{mt} + \varepsilon_{it}. \quad (19)$$

Table E1 shows that palm development has very small effects on acacia development. Palm does not displace acacia and if anything slightly increases acacia investment, perhaps in opening up new lands. That is, palm and acacia are not substitutes, but rather weak complements. I can isolate intensive-margin investments by focusing on sites with nonzero initial development, and I can allow for cross-site effects by aggregating over sites within provinces. Both lead to similar results.