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

Allan Hsiao *
Massachusetts Institute of Technology

January 18, 2021
(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. In particular, I evaluate EU legislation targeting palm oil imports, primarily from Indonesia and Malaysia. I find coordinated, committed import tariffs to be effective, reducing carbon emissions relative to observed outcomes by 56% compared to 64% under a domestic palm oil tax. As coordination breaks down, emission reductions fall from 56% for action by all importers, to 17% for an EU-China-India coalition, to 2% for unilateral EU action, as tariff coverage falls from 80% to 35% to 12% of world consumption, respectively. As commitment breaks down, carbon reductions fall to as low as 0%. Finally, coordination and commitment interact. Achieving 95% of the full-commitment outcome requires a commitment period of only five years when importers coordinate, but more than twenty years when the EU acts unilaterally.

*Email: ajhsiao@mit.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, Tavneet Suri, and workshop audiences at MIT. 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 countries fail to regulate emissions domestically. Indeed, domestic regulation often faces significant challenges: low incentives from free riding and political constraints (Oates and Portney 2003), and 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 CO_2 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 4.7% of global CO_2 emissions from 1986 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 that import tariffs generally faces 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. As a result, the losses are disproportionate as the tariff coalition shrinks. A small coalition covers a small proportion of global consumption, and leakage concerns further constrain it to small tariffs.

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 leads to low marginal costs up to capacity. For example, once the costs of identifying, exploring, and drilling an oil well have already been paid, extraction is cheap and thus likely to proceed 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 in Indonesia and Malaysia threatens peatland forests that are particularly carbon-rich. Second, do-

Figure 1: CO_2 emissions from palm oil plantations

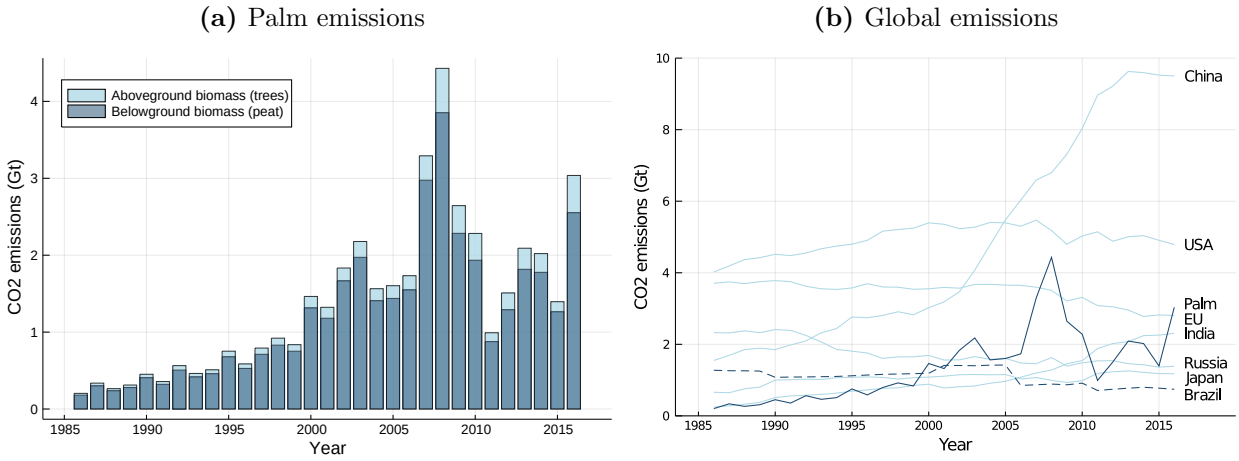


Figure 1a computes CO_2 emissions by combining data on palm oil plantations from Xu et al. (2020) and Song et al. (2018) with data on tree biomass from Zarin et al. (2016) and peat deposits from Gumbricht et al. (2017). Figure 1b uses data on global CO_2 emissions by country, including emissions from land-use change, from the World Resources Institute and Global Carbon Atlas. I show emissions for the top seven emitters from 1986 to 2016 alongside palm emissions, which account for 4.7% of global CO_2 emissions during this period. I highlight Brazil, which also generates significant emissions through land-use change. These calculations focus on CO_2 emissions, which account for 73% of total greenhouse gas emissions during the study period. I focus on CO_2 because the carbon content of peatlands is well documented, as detailed in appendix B.6. Palm oil production also involves the release of methane and nitrous oxide, but precise estimates of these emissions are not yet well established.

mestic incentives to regulate are limited. Despite its global consequences, palm oil is a major source of export revenue for Indonesia and Malaysia 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 the elasticity of palm oil demand in unregulated markets. Demand responses in turn depend on consumers' substitution between palm and other vegetable oils. The commitment problem depends on the elasticity of palm oil supply, and how it differs between short- and long-term tariffs. Supply responses 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). This product-space approach to demand estimation has two advantages: it allows for flexible patterns of substitution between palm and other vegetable oils, and it avoids the need to specify exactly which product characteristics consumers value. 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 using foreign rainfall shocks in oil-producing regions as instruments. I then estimate the extent to which world demand for palm oil shifts over time, and I use these demand shifts – driven, for example, by changes in total vegetable oil consumption – as price instruments in estimating supply.

I model palm oil supply with a dynamic model of land development for palm oil. In the model, forward-looking firms make sunk investment decisions 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 develop into plantations.[1] Data derived from satellite imagery allow me to observe these choices over time and at a high degree of spatial resolution. Firms' investments produce palm oil in each period and generate revenues as a function of world prices, which in turn depend on aggregate investment in palm oil production. Firms therefore play a dynamic competitive equilibrium as in the entry and investment game of Hopenhayn (1992). Modeling the dynamic investment decision allows me to infer firms' responses to hypothetical tariffs from their responses to observed price variation, while accounting for price expectations in a disciplined way. Intuitively, in the same way that price shocks today change both current revenues and expectations over future revenues, tariffs change revenues both today and in the future.

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 analyze the intertemporal trade-off in 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 investment levels fixed. Continuation values difference out, and estimation reduces to linear regression with instruments. Identification comes from two sources: exogenous variation in world palm oil prices over time, as induced by the demand shifters discussed above, and exogenous variation in palm oil yields over space, as induced by differences in sunlight and precipitation. Prices and yields interact because high prices raise revenues most for high-yield plantations. Furthermore,

[1] 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).

while a conventional full-solution approach would need to specify exactly how firms expect the state of the economy to evolve over the long term, the Euler approach relies instead on the weaker assumption of rational expectations. The computational advantage is that the Euler approach avoids solving the model for estimation, while the full-solution approach requires solving repeatedly.

For counterfactuals, specifying firms' expectations and solving the model are unavoidable, and so I solve by backward induction from the steady state. The model assumes no exit and therefore reaches a steady state when all feasible lands are exhausted. The computational challenge is that it takes many periods to reach this point, 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 the solutions converge. In the inner loop, I backward induct with a limited look-ahead window, then I update the starting point based on the solution and repeat until finding a fixed point. To quantify emissions, I combine spatial data on carbon stocks with the model's spatial predictions for plantation development, and I assume a social cost of carbon of $40 per ton. I also make the strong assumption that non-palm deforestation does not expand in response to palm oil tariffs. The primary threat to this assumption is substitution from palm to acacia plantations, but I assess this substitution and find it to be empirically small.

I evaluate how coordination and commitment, both individually and in combination, influence the effects of import tariffs on carbon emissions and social welfare, and I benchmark these effects against a domestic palm oil tax implemented by Indonesia and Malaysia. The domestic tax avoids the leakage problem because it covers all production, and it avoids the commitment problem because it can be imposed upfront with a license fee for new development. In my baseline analysis, all regulation is set to maximize social welfare and is uniform across units of palm oil, although I also present extensions that relax each condition. I find that import tariffs can be an effective substitute for domestic regulation. When coordination and commitment hold, import tariffs reduce carbon emissions by 56% relative to observed outcomes under business as usual. By comparison, the domestic tax reduces emissions by 64%. The loss arises because import tariffs cannot regulate domestic consumption in Indonesia and Malaysia. But the loss is not disproportionate because I find Indonesian and Malaysian demand to be quite inelastic, such that leakage is limited.

At the same time, emission reductions diminish as coordination and commitment weaken. Even under full commitment, relatively elastic demand among importers causes emission reductions to fall from 56% under full coordination among importers, to 17% under an EU-China-India coalition, to 2% under unilateral EU action. These 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. Time to build accounts for the stark no-commitment

result: it is statically optimal to eliminate tariffs because tariffs today do not affect new development, which does not generate taxable production until a later period. Thus, both coordination and commitment are necessary. When either fails, import tariffs are low and have little effect.

Furthermore, coordination and commitment interact, with weak coordination increasing the importance of commitment. As a intermediate between full and no commitment, I consider a limited commitment scenario in which importers commit to a tariff regime over a fixed number of periods at a time – 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. The interaction between leakage and commitment arises because, anticipating the temptation to reduce tariffs in future periods, importers wish to increase tariffs today. However, leakage makes doing so difficult. Producers facing large tariffs in regulated markets can make investments and focus sales on unregulated markets. Then as tariffs are reduced – because investment is sunk – producers can shift sales to regulated markets. The more severe the leakage problem, the more unregulated markets can absorb, and thus the more easily producers can skirt tariffs.

The division of surplus among countries reveals why coordination and commitment are difficult to achieve in practice. Coordination is difficult because own-surplus-maximizing coalition members have an incentive to defect. For example, the EU-China-India coalition becomes fragile if China and India ignore carbon damages and focus on their consumer surplus alone: China and India lose consumer surplus when they impose tariffs, but they gain when they do not because leakage allows defectors to free ride on lower world prices. Commitment is difficult when countries value their consumer surplus alone because longer commitment demands larger sacrifices of consumer surplus for the sake of reducing emissions. Lastly, for Indonesia and Malaysia, under most tariff scenarios I find that imposing the socially optimal domestic tax leads to lower surplus. However, Indonesia and Malaysia prefer domestic regulation if threatened with fully coordinated import tariffs. In this scenario, the domestic tax has low marginal impact on producer surplus because the outside option is tariffs that are already high, and the domestic tax raises government revenue that would otherwise go abroad.

The main contribution of this paper is to develop an empirical framework for assessing trade policy as a means of environmental regulation. While Shapiro (2020) establishes the negative outcomes of emission-inattentive trade policy, I show what emission-attentive trade policy can achieve, and I quantify the challenges in implementing such policy. In particular, I study two problems – leakage and commitment – that are well recognized individually, and I provide novel analysis of how the two interact in an empirical setting. A rich literature on environmental regulation in trade-exposed markets documents how supply-side leakage undermines domestic regulation as polluters move to unregulated markets, motivating 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). Similarly, demand-side leakage becomes a concern in my context, as free-riding makes the leakage problem fundamental and adds value to acting in coalition (Nordhaus 2015). I also build on a literature studying commitment problems in environmental regulation, in which the dynamic incentives to abate emissions depend critically on whether penalties are upheld over future periods (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).

Methodologically, my framework builds 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 Kalouptsidi et al. (2018), that develops Euler conditional choice probability (CCP) 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. In focusing on short-term perturbations in order to simplify dynamics, these Euler methods are closely related to moment-inequality techniques for revealed preference (Bajari et al. 2007; Pakes 2010; Pakes et al. 2015), with applications ranging from store placement to pension plans to export destinations (Holmes 2011; Illanes 2017; Morales et al. 2019). My contribution is to show how to combine both continuous and discrete Euler techniques in a single framework, with a model containing discrete entry choices on the extensive margin and continuous investment choices on the intensive margin. Indeed, many investment decisions involve a similar combination of extensive- and intensive-margin choices. I also show how to tractably solve my model in computing a set of counterfactuals unidentified by Euler methods alone.

More broadly, this paper contributes a quantitative analysis of environmental regulation for one of the world's largest sources of carbon emissions. Palm oil is ubiquitous, adding value to food and consumer products worldwide. But these benefits have come with severe costs: the industry accounts for an enormous 4.7% of global CO_2 emissions over the last three decades. Domestic regulations have failed to prevent these emissions, but trade policy offers an alternative set of tools for regulating this and other industries operating in low-regulation environments. Unlike the domestic programs evaluated in Burgess et al. (2019) and Souza-Rodrigues (2019), or the conservation contracting of Harstad (2012, 2016) and Harstad and Mideksa (2017), trade policy does not rely on a domestic government that is willing and able to enforce regulation. And unlike the payments for ecosystem services of Jayachandran et al. (2017) and Edwards et al. (2020), trade policy scales readily and does not rely on property rights that are well defined. Furthermore, swift action can still save vast swathes of forest that remain intact, particularly in Papua. Nonetheless, as with other forms of international climate action, coordination problems and dynamic concerns present fundamental challenges. This paper quantifies these challenges in an industry that is pivotal in the fight against climate change.

2 Illustrative Model of Emission-Based Trade Policy

This section studies optimal tariffs for an emission-intensive traded good in a setting with incomplete regulation and sunk investment. It discusses the leakage and commitment problems.

2.1 Import tariffs under incomplete regulation and sunk investment

Consider two markets: 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. Established plantations produce goods every period at zero marginal cost, do not depreciate, and have zero scrap value. Production begins one period after development.

I study tariffs on regulated consumption, where tariffs are set to maximize social welfare – i.e., consumer and producer surplus net of emission damages. Appendix A.1 provides derivations. Social welfare depends on old development Q_t^o, the path of new development $\{Q_t^n, Q_{t+1}^n, \ldots\}$, and how the resulting production is allocated across markets. Given discrete time, discount factor β, new development $Q_t^n = Q_t^{rn} + Q_t^{un}$, and old development $Q_{t+1}^o = Q_t^o + Q_t^n$,

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

Domestic regulation

The first best is a domestic Pigouvian tax that reflects the full magnitude of the externality.

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

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 target new development with a license fee and thus impose the full tax upfront.

The leakage problem

Regulation is incomplete because import tariffs miss unregulated consumption. To isolate the leakage problem, suppose importers can commit to upholding tariffs. The optimal tariff is

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

where $\varepsilon_t^S > 0$ and $\varepsilon_{t+1}^{Du} < 0$ are elasticities of supply and unregulated demand, and "C" indicates full commitment. Even within the regulated market, the tariff is smaller than the first-best tax. First, leakage lowers the benefits of the tariff relative to the first best. Although tariffs decrease regulated consumption, net emission reductions are smaller because tariffs also increase unregulated consumption as they lower world prices. Second, leakage raises the costs of the tariff. Tariffs shift consumption from higher willingness-to-pay consumers in the regulated market to lower willingness-to-pay consumers in the unregulated market, and in doing so produce allocative inefficiency.

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: under leakage, they create allocative inefficiency in distorting consumption between markets. In combination, these forces make it statically optimal to set tariffs to zero. In the no-commitment case, importers follow these static incentives in each period and never levy tariffs at all.

Under limited commitment, I assume that importers can commit to upholding tariffs for L periods at a time. In other words, they revise tariffs every L periods. I consider a special case with time-invariant demand and supply curves in order to highlight intuition and solve for tariffs in closed form. The empirical exercise avoids these assumptions by solving numerically. Importers remove tariffs at the beginning of each L-period regime, and they set tariffs in other periods anticipating these periodic breaks. Tariffs have net present value

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

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 $L \to \infty$.

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

In the more general case, the statically optimal tariff also decreases over time because tariffs do less to reduce emissions as the stock of sunk investment grows. At the extreme, tariffs are set to zero when all lands are exhausted because tariffs cannot reduce emissions when there are no forests left to save. The above formula nests this case in which the elasticity of supply is zero.

How leakage and commitment interact

The key mechanism is that producers shift sales across markets as tariffs change. That is, producers focus on the unregulated market when tariffs are high, and shift toward the regulated

market when tariffs are low. As a result, leakage and commitment interact. Intuitively, the regulator can only compensate for low future tariffs by imposing high tariffs while tariffs are in place. But these high tariffs suffer from leakage, and so the regulator cannot compensate fully. In particular, incomplete regulation allows producers to skirt high tariffs in any given period by directing sales to the unregulated market until tariffs fall. Thus, leakage exacerbates the commitment problem. The greater the leakage problem, the more the unregulated market can absorb, and thus the greater the loss from failures of commitment.

2.2 Extensions

Appendix A.2 considers heterogeneous emissions across goods, which I also allow for in the empirical exercise. Appendix A.3 studies terms-of-trade effects, which the baseline model shuts down because tariffs are set to maximize social welfare. Finally, the baseline model treats emissions as released upon development, as is appropriate for deforestation. When emissions are instead released over time, either in production or consumption, the same framework applies if sunk investment in a brown technology leads to permanently low marginal costs of production, such that production continues in each period. Emissions are then committed upon investment, and externality e becomes the net present value of emission damages.

3 Empirical Setting and Data

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

3.1 Empirical setting

Palm oil is among the most widely used plant products in the world. High yields drive its low price point, with oil palm producing more oil per hectare of land than any comparable oilseed. Palm oil is used as a cooking oil, particularly in Asia, and is a common ingredient in processed foods, where it has replaced trans fats. Palm oil also has non-food uses ranging from soaps to cosmetics to biofuels. At the country level, table 1 shows that Indonesia and Malaysia account for 84% of global production, 90% of exports, and 20% of consumption, with the European Union, China, and India accounting for another 35% of global consumption. At the firm level, the market is unconcentrated: the largest producer (FGV Holdings Berhad) accounts for 4% of global production (POA 2017), and the largest consumer (Unilever) accounts for 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 such 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

Table 1: Palm oil production and consumption 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 and each sum to one.

deposits exceed those from tree biomass by five to ten times.[2] Figure 1b shows that palm emissions account for more CO_2 from 1986 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, with the goal of promoting domestic efforts to curb deforestation. As a case study, consider Indonesia's primary response: a 2011 moratorium on new forest concessions. Busch et al. (2015) cite problems of weak regulation and weak enforcement. The moratorium failed to regulate forests within existing concessions, and regulating all concessions would still have been insufficient because most deforestation occurred (illegally) outside of concessions, including in protected areas.

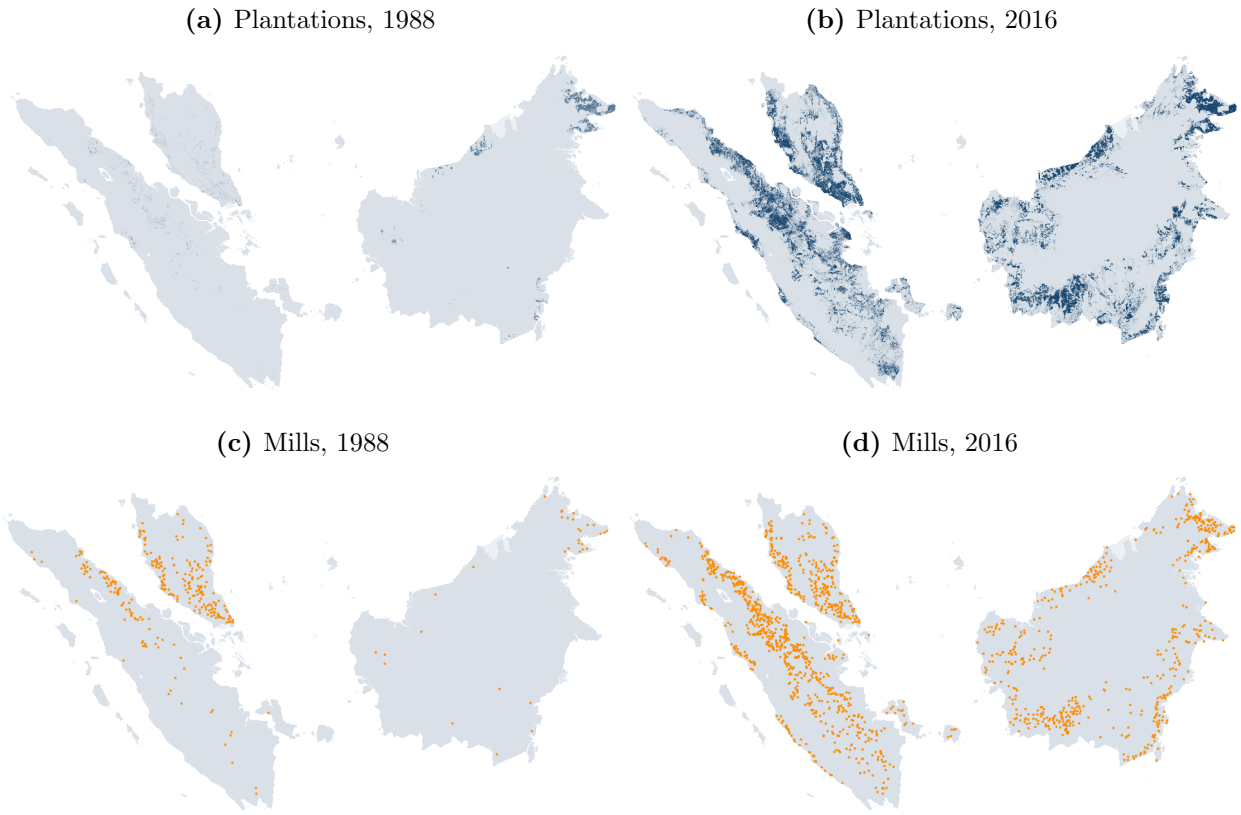
Third, foreign governments are actively discussing trade-policy interventions, particularly in Europe. French parliament debated a "Nutella" tax on palm oil in food products in 2016, although it failed to pass. Furthermore, the European Union initially provided green subsidies for palm-based biofuels, but policymakers later recognized the consequences of palm-driven deforestation. Recent policy therefore moves to eliminate green subsidies for palm-based biofuels, cap production, and achieve a complete phase-out by 2030. As well, palm-based biofuels face the further loss of green tax incentives in France and an outright ban in Norway, both by 2020. While none of these policies explicitly imposes tariffs across all palm oil imports, they all leverage European buying power to influence emissions abroad in the same way that tariffs do.

3.2 Data

I compile data on palm oil production, consumption, and world prices, with data sources and construction detailed in appendix B. For production, I assemble a spatial panel dataset at a resolution of 30 arc-seconds – approximately 1 km^2 – that measures the annual extent of plantations

[2] Converting peatlands to croplands involves draining peatlands and clearing the land with fire, releasing large amounts of carbon. Even without clearing by fire, unsubmerged peat releases carbon as it decomposes. Furthermore, fire spreads quickly on dried-out peat, and in 2015 slash-and-burn practices combined with dry El Niño conditions caused an estimated 100,000 deaths and $16 billion in damages (Koplitz et al. 2016; World Bank 2016).

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 and Kalimantan of Indonesia and all of Malaysia, covering 83% of global palm oil production.

and mills from 1988 to 2016 using satellite imagery. Figure 2 maps their widespread expansion over this period. For plantations, Xu et al. (2020) analyze PALSAR and MODIS satellite data to measure the expansion of palm oil plantations 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 (appendix B.2). 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 manually cross-reference historical satellite data to identify construction dates back to 1988 (appendix B.3). For both plantations and mills, I find that the data align closely with aggregate government statistics.

Because plantation development and mill construction are interdependent, I lightly harmonize the plantation and mill data to ensure consistency. The industry standard is that plantations be within 50 kilometers of a mill because oil palm crops deteriorate rapidly, losing value between harvest and milling. I impose this standard on the data: for plantation development without a mill within 50 kilometers in the current period, I either delay development until such a mill is

constructed, or I drop it entirely if no such mill is constructed by 2016. This procedure affects only 7% of plantation development, with 5% delayed and 2% dropped (appendix B.4).

I supplement these data with spatial data on land characteristics, which I map in figure 3. I compute palm oil yields from two sources. First, I compute potential yields at a disaggregated level using an agronomic model of the oil palm plant (Hoffmann et al. 2014). The agronomic model delivers potential yields as a function of solar radiation and precipitation, which I measure at the grid-cell level. These potential yields vary over space but not over time, and furthermore differ from the yields that farmers actually attain. Second, I compute yield gaps: one minus the proportion of potential yields attained. I do so by province-year with data on actual yields from government statistics. Assuming yield gaps are homogeneous within province-years, I combine both sources of data to obtain yearly estimates of attained yields at a disaggregated level (appendix B.5). Other sources of spatial heterogeneity include Euclidean distances to the nearest port, road, and urban district, all of which affect transport costs.

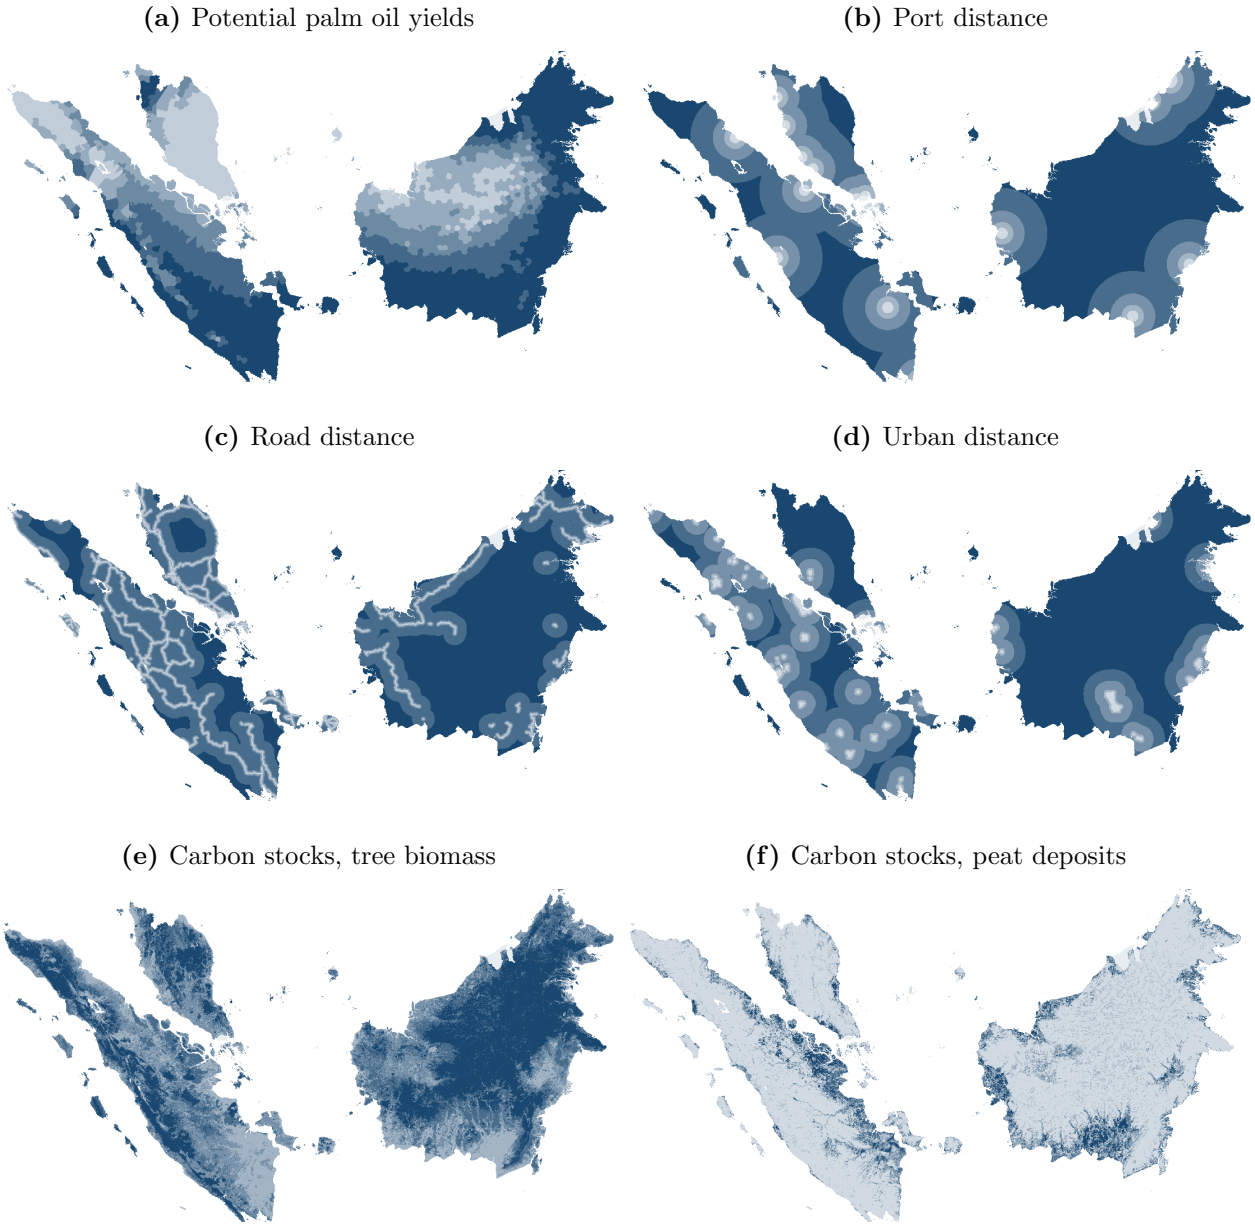
Figure 3 also maps the spatial distribution of carbon stocks. These data provide a direct link between counterfactual production and emissions because I observe how much carbon would be released by developing any given plot of land. I construct these data at a resolution of 30 arc-seconds by combining geospatial data on tree biomass and peat deposits from Zarin et al. (2016) and Gumbricht et al. (2017), respectively. I use conversion factors of 0.5 for biomass to carbon, 65.1 kg C/m^3 for peat to carbon, and 3.67 for carbon to CO_2. Furthermore, I treat these carbon stocks as predetermined. The tree biomass data are measured in 2000, so I use the data to estimate tree biomass in 1988 instead of using the 2000 values directly. I do by combining the tree biomass data in 2000 with the Song et al. (2018) data on tree cover loss, which provide the extent of tree loss between 1988 and 2000. The peat deposits data are measured in 2011 but largely rely on predetermined features like precipitation and topography to identify wetland areas (appendix B.6).

For consumption and world prices, I compile annual panel data from 1980 to 2016 on palm oil and its substitutes. Consumption data come from the USDA Foreign Agricultural Service and are measured at the country level, while world price data come from the International Monetary Fund and the World Bank. Palm oil includes palm and palm kernel oils, and substitutes include coconut, olive, rapeseed, soybean, and sunflower oils. I omit cottonseed and peanut oil – 5% of vegetable oil consumption by volume in 2016 – given a lack of historical price data. Figure 4 plots the time-series variation in world prices against palm oil production capacity.

To address price endogeneity, I measure rainfall shocks in oil-producing regions. Rainfall data come from the Global Meteorological Forcing Dataset, which records rainfall in millimeters per day from 1980 to 2016 at 0.25° resolution. I identify producing regions with production data from the USDA Foreign Agricultural Service.[3] For each crop, year, and region, I compute rainfall shocks

[3] A region is either a country or, in the case of large countries, a subnational region. I use subnational regions – namely, states or provinces – for Argentina, Brazil, Canada, China, Indonesia, Malaysia, Russia, and the US.

Figure 3: Land characteristics

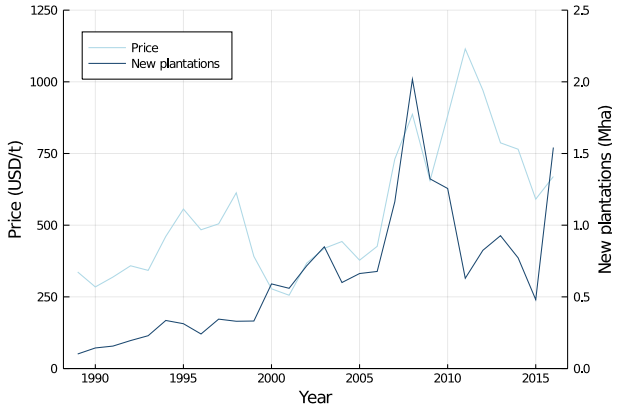


Darker blue indicates high yields, farther distances, and larger carbon stocks. Yields are generated from the PALMSIM agronomic model (Hoffmann et al. 2014). Major ports are from the 2019 World Port Index and World Port Source, and major roads are from the Global Roads Inventory Project. Urban areas include administrative cities in Indonesia (*kota*) and federal territories in Malaysia. Carbon stocks for tree biomass and peat deposits are from Zarin et al. (2016) and Gumbricht et al. (2017), respectively.

as the total absolute deviation of actual monthly rainfall from optimal rainfall levels during the growing season. I define optimal levels for each crop as the midpoint of the optimal window recorded in the FAO Crop Ecological Requirements Database (ECOCROP). I then compute rainfall shocks at the crop-year level by aggregating over regions, weighting by regional production over the study period. I do not weight by yearly production because it is a direct function of yearly rainfall.

Figure 4: Palm oil production vs. world prices

(a) Prices vs. plantation development



(b) Prices vs. mill construction



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.

4 Empirical Model

This section specifies empirical models of palm oil demand and supply. The resulting demand and supply curves correspond to the functions $P_t^{Dr}(q)$, $P_t^{Du}(q)$, and $P_t^S(q)$ of section 2.

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). First, consumers make an upper-level decision over total vegetable oil consumption. Second, given this total, they make a lower-level choice between palm oil and "other" oils, namely coconut, olive, rapeseed, soybean, and sunflower oils.[4] 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. To obtain both $P_t^{Dr}(q)$ and $P_t^{Du}(q)$, I specify separate demand curves for regulated and unregulated consumer markets.

For a given consumer market, the specifications are as follows. For the lower level,

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

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

[4] 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.

for expenditure shares ω_{it}, palm and other oil prices p_{jt}, total vegetable oil expenditures X_t, and translog price index P_t. For the upper level,

$$\ln Q_t = \alpha^0 + \alpha^1 t + \gamma \ln P_t + Z_t \beta + \varepsilon_t \,, \tag{2}$$

where Q_t is the quantity of total vegetable oil consumption, and P_t is the price index above. Demand shifters Z_t include GDP, which captures total expenditures, and the CPI, which captures prices of other consumption goods. Finally, the two levels are linked by $X_t = Q_t P_t$.

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.

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

where $\frac{\partial \ln P_t}{\partial \ln p_{jt}} = \alpha_j^0 + \alpha_j^1 t + \sum_k \gamma_{kj} \ln p_{kt}$, and δ_{ij} is the Kronecker delta (one for $i = j$ and zero otherwise). The first three terms arise from the lower-level decision, with i-specific parameters capturing substitution between palm and other oils. Ignoring these substitution patterns would impose that palm oil have no close substitutes. The last term comes from the upper-level decision, with the γ parameter governing how category demand for vegetable oils responds to prices. Ignoring this response would hold total vegetable oil expenditures fixed over changes in prices.

Endogeneity

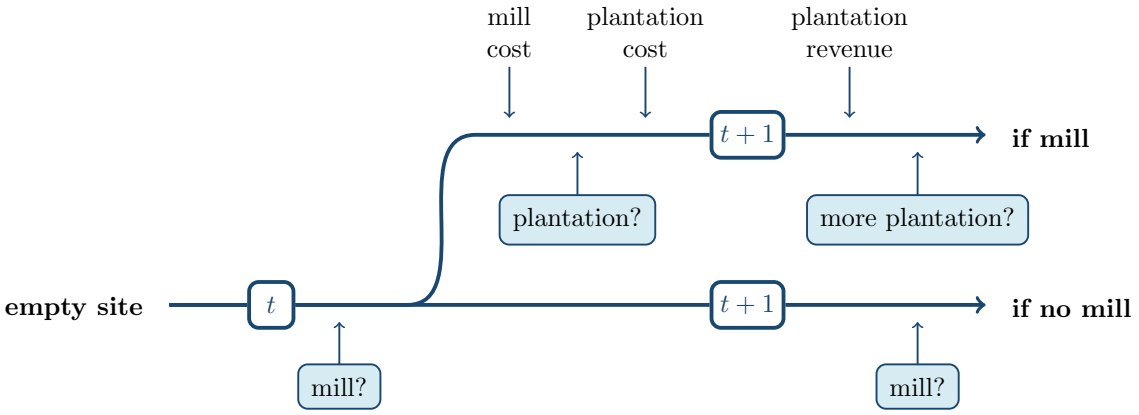
As is typical, prices are endogenous. Unobservables ε_{it} and ε_t shift demand and therefore affect equilibrium prices p_{jt}. I instrument with rainfall shocks in producing regions as a supply shifter. The exclusion restriction is that these shocks affect vegetable oil demand only through their impact on prices. Domestic rainfall shocks may impact demand directly, and so I focus on foreign rainfall shocks, which arguably satisfy the exclusion restriction absent macroeconomic consequences that affect incomes or consumption patterns abroad.

4.2 Supply: a dynamic model with sunk investment

Forward-looking firms generate profits by making sunk investments in palm oil mills and plantations. Land is divided into sites, which I assume are small, independent, and managed by long-lived owners. Sites make investment decisions 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.[5] Figure 5 shows the timeline.

[5] The model assumes that the same agent makes both decisions. In practice, about half of plantations are managed by smallholders that contract with mills. Smallholders do not present a problem for the intensive-margin model,

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 binary choice in the following period. If so, then the site makes a continuous choice over how much land to develop into plantations. In future periods, the site faces more continuous choices over plantation expansion.

Intensive margin (plantation development)

In each period t, sites i that have mills choose how much land area a_{it} to develop into plantations. Development is an investment, increasing revenues tomorrow at some cost today. Plantations have no scrap value and are sunk, and plantation size s_{it} evolves by deterministic law of motion

$$s_{it+1} = s_{it} + a_{it} \, .$$

Dynamics enter because development a_{it} today affects plantation sizes in all future periods. A static model would eliminate these dynamics with law of motion $s_{it+1} = a_{it}$.

Profits depend on publicly observed state $\boldsymbol{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$ and privately observed state ε_{it}, where these states include both site-specific and aggregate elements. Profits vary across sites as a function of site-specific yields Y_{it}, which affect revenues, and site-specific cost factors x_i and shocks ε_{it}, which affect costs. Furthermore, world prices $P(s_t, d_t)$ depend on aggregate supply s_t and aggregate demand d_t. Aggregate supply s_t measures total production over all plantations.

$$s_t = \sum_i Y_{it} s_{it}$$

Higher supply lowers world prices and profits. As in Hopenhayn (1992), sites are atomistic: collective action affects world prices but individual actions do not, and firms play a dynamic competitive equilibrium in which collective action coincides with individual expectations. Aggregate demand

as long as their choices are efficient. On the extensive margin, the underlying assumption is that the mill extracts all surplus from its associated plantations. Indeed, mills have spatial market power over plantations, which must sell to nearby mills because palm fruit decays in transport. Mills are therefore well positioned to extract surplus. The assumption breaks down if smallholders have bargaining power that allows them to extract rents.

d_t captures world demand for palm oil, with higher demand raising world prices and profits. The timing of each period is that sites with mills realize state $(\boldsymbol{w}_{it}, \varepsilon_{it})$ then choose investment a_{it}, which begins generating revenues in the following period.

The value function and my specifications for the revenue and cost functions are

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

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

where investment increases future revenues while incurring some cost. Expectations are over next-period state $(\boldsymbol{w}_{it+1}, \varepsilon_{it+1})$, with shorthand $\mathbb{E}_{it}[\cdot] \equiv \mathbb{E}[\cdot | s_{it}, \boldsymbol{w}_{it}, \varepsilon_{it}]$. Revenues are linear in plantation size and increasing in yields and world prices. Costs are quadratic and convex in investment, reflecting diseconomies of scale such as credit and local factor market constraints, and encouraging investment to be spread over time. Linear revenues and convex costs together ensure unique optima. Cost factors x_i capture observed heterogeneity at the site level, and I accommodate regional unobserved heterogeneity κ_m and regional time trends α_m. Unobserved cost shocks ε_{it} are mean-zero and IID over time, but can be correlated across sites.

Sites evaluate investments on a net-present-value basis and therefore do not distinguish between upfront and future flow costs. Thus, the cost function can also be interpreted as capturing flow costs realized over time. Similarly, yields Y_{it} are subject to rainfall shocks ε_{it}^Y that affect production in each period. But these weather shocks do not enter sites' investment decisions because they average to zero over time. That is, forward-looking sites invest based on climate and not weather.

Extensive margin (mill construction)

In each period t, sites i that do not have mills make a binary choice a_{it}^e over whether to construct a mill. Mill construction is also an investment: it leads to palm oil revenues tomorrow at some cost today. Plantations require a mill because oil palm fruit decays quickly after harvest if not milled into oil, and the fruit is not consumed directly. Like plantations, mills have no scrap value and are sunk, with law of motion $s_{it+1}^e = s_{it}^e + a_{it}^e$. State variables include publicly observed $\boldsymbol{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$ as above and privately observed binary choice shocks $\varepsilon_{it}^e = \{\varepsilon_{it0}^e, \varepsilon_{it1}^e\}$, which are mean-zero, IID over time and across sites, and logit-distributed with standard deviation σ^e. For sites without mills, the timing of each period is as follows. First, sites realize state $(\boldsymbol{w}_{it}, \varepsilon_{it}^e)$. Second, they make choice a_{it}^e over whether to construct a mill. Third, if they choose not to do so then the period ends. Fourth, if they choose to do so then they realize intensive-margin cost shock ε_{it}. Fifth, they make intensive-margin choice a_{it}, and the period ends.

The ex-ante value function, the choice-specific conditional value functions, and my specification

for the cost function are

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

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

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

$$c^e(\boldsymbol{w}_{it}) = x_i\gamma^e + \kappa^e_m + \alpha^e_m t\,, \tag{5d}$$

where the e superscript indicates the extensive margin, with shorthand $\mathbb{E}^e_{it}[\cdot] \equiv \mathbb{E}^e[\cdot|\boldsymbol{w}_{it}]$. In equation 5a, expectations are over logit shocks ε^e_{it} that imply mill construction probabilities

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

I suppress regional subscripts m in the notation, but these probabilities are more precisely $p^e_m(\boldsymbol{w}_{it})$ given regional heterogeneity in the cost function. In equation 5b, choosing not to build leads to the same decision in the following period, subject to expectations over next-period state $\boldsymbol{w}_{it+1}$. In equation 5c, choosing to build incurs mill construction costs in return for the value of plantation development on the intensive margin, 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 at the site level, and I accommodate regional unobserved heterogeneity κ^e_m and regional time trends α^e_m. It is isomorphic to think of the difference in logit shocks ε^e_{it} as a cost shock. Finally, the outside option is never constructing a mill, with utility normalized to zero given mean-zero shocks ε^e_{it}.

Unobserved heterogeneity and endogeneity

The primary restriction on both margins is that unobserved heterogeneity is allowed only at the regional level. Within regions, sites can receive differential shocks but otherwise have no persistent heterogeneity beyond that explained by observables. On the intensive margin, identifying site-level unobserved heterogeneity would require a long panel of plantation development decisions, which I only have for sites where development began earliest. On the extensive margin, identifying site-level unobserved heterogeneity would require multiple mill construction decisions per site, but the model assumes that sites construct no more than one mill each.[6]

Even with this restriction, there remains an endogeneity concern on the intensive margin: both prices P_t and yields Y_{it} may be correlated with unobserved costs ε_{it}. First, low costs induce entry, raising supply and lowering prices. Second, attained yields depend on unobserved, costly effort, and furthermore low costs induce entry disproportionately from high-yield sites. I therefore instrument for prices with demand shifters d_t and for yields with potential yields Y^p_i. Demand shifters come

[6] Thus, the expectation-maximization approach of Arcidiacono and Miller (2011) for accommodating unobserved heterogeneity does not apply in this setting.

from demand estimation, which delivers the world demand curve for palm oil in each period.

$$\ln p_t = \widehat{\phi} \ln q_t + \widehat{d}_t$$

The intercept captures the level of demand over time, which I interpret as a demand shifter. Variation in total oil consumption $\ln Q_t$ drives this demand shifter, and indeed instrumenting directly with $\ln Q_t$ leads to similar results. Potential yields by site are functions of solar radiation and precipitation, which are exogenous, and instrumenting also mitigates bias from mismeasured yields. These concerns do not arise on the extensive margin because mills do not themselves affect prices or yields, and mill shocks ε^e_{it} are assumed to be uncorrelated with plantation shocks ε_{it}.

I take cost factors x_i to be exogenous. I measure port distance using only major ports, which largely predate the palm oil industry. I measure road distance using only major roads, and not the small roads that develop endogenously as land is cleared for palm oil. I measure urban distance using officially designated urban districts, which cover only major cities and do not include palm oil settlements. Finally, carbon stocks are predetermined as discussed in section 3.2.

5 Estimation

This section describes how I estimate the demand and supply models specified in section 4. I take an iterated linear least squares approach for demand and an Euler approach for supply.

5.1 Demand: iterated linear least squares

I adopt the iterated linear least squares approach of Blundell and Robin (1999) to estimate the lower-level specification. I start by estimating the linear approximate version of the demand system with instruments, using the Stone price index ($\ln p_{it} = \sum_j \omega_{jt} \ln p_{jt}$) in place of the translog price index. I then construct the translog price index with the resulting estimates and iterate until convergence. In this way, I avoid nonlinear estimation. In each iteration, I estimate the system by seemingly unrelated regression, accounting for serial correlation in the error terms with a Prais-Winsten transformation as in Parks (1967). I impose the standard adding-up restriction, but not homogeneity or symmetry.[7] I aggregate palm and palm kernel prices with a fixed 1:8 ratio, reflecting the composition of the oil palm fruit, and I aggregate coconut, olive, rapeseed, and sunflower oils with a Stone price index. I compute standard errors with the delta method, and I evaluate expenditure shares, prices, and the time trend at their averages over the study period. Given the lower-level estimates, I estimate the upper-level specification by linear IV.

[7] The adding-up restrictions are $\sum_i \alpha^0_i = 1$, $\sum_i \alpha^1_i = 0$, $\sum_i \beta_i = 0$, $\sum_i \gamma_{ij} = 0 \; \forall j$ 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 $\gamma_{ij} = \gamma_{ji} \; \forall i, j$.

5.2 Supply: Euler approach

I take an Euler approach for estimation, focusing on the timing of observed investment 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 investment levels fixed. Continuation values difference out. Estimation proceeds in three steps: (1) defining site boundaries, (2) estimating the intensive-margin model, and (3) estimating the extensive-margin model. I assume a discount factor of $\beta = 0.9$, as the discount factor is typically unidentified in dynamic discrete choice models (Magnac and Thesmar 2002).

Step 1: defining site boundaries

The model associates each mill with an individual site, and so I draw boundaries for potential sites using observed mills as a guide. I identify the palm oil industry's most developed provinces by several metrics, and in each case I observe approximately one mill per 500 km^2. I use this density as a cutoff. For provinces with observed mill density above the cutoff, I assume the market is saturated. Without additional entrants, the number of potential sites is simply the number of observed mills, and I draw site boundaries by assigning land to the closest mill. For provinces with observed mill density below the cutoff, I assume further entry is possible until the cutoff is reached. I assign land to sites by k-means clustering, where the number of clusters k is chosen using the cutoff, and where I impose that clusters separate observed mills. The end result is a set of contiguous potential sites, each containing either zero or one observed mill. In total, I obtain 2,805 potential sites. I restrict attention to Sumatra, Kalimantan, and Malaysia, which account for 96% of palm oil area harvested in Indonesia and Malaysia, and where the quality of the mill and plantation data is highest. Appendix D.1 provides further detail.

Step 2: estimating the intensive-margin model (plantation development)

The first order condition for investment and the envelope theorem deliver an Euler equation.

$$\begin{aligned} c'(a_{it}; \boldsymbol{w}_{it}, \varepsilon_{it}) &= \beta \mathbb{E}_{it}[V'(s_{it+1}; \boldsymbol{w}_{it+1}, \varepsilon_{it+1})] \,, \\ V'(s_{it}; \boldsymbol{w}_{it}, \varepsilon_{it}) &= r'(s_{it}; \boldsymbol{w}_{it}) + \beta \mathbb{E}_{it}[V'(s_{it+1}; \boldsymbol{w}_{it+1}, \varepsilon_{it+1})] \,, \end{aligned}$$

where the first line is the first order condition for a_{it} and the second line applies the envelope theorem to equation 4a. Together, these equations imply the Euler equation

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

which captures the intertemporal trade-off in investing in period t compared to $t+1$. With the functional form assumptions of equation 4b, the Euler equation specializes to

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

with shorthand $P_t \equiv P(s_t, d_t)$. Intuitively, it compares developing today and tomorrow. The left-hand side is the marginal increase in costs, and the right-hand side is the marginal increase in revenues. The Euler equation implicitly assumes an interior solution, otherwise the first order condition may not hold. Indeed, 99.2% of observed intensive-margin decisions are interior: 0.8% involve zero development, and 0% hit the upper bound of land available within a site.

Because I do not observe expectations of next-period values, I take realized values as noisy measures of these expectations as in Hall (1978). In particular, I replace expected values with realized values and expectational errors η_{it}, and I rearrange to obtain the regression equation

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

for $\tilde{t} \equiv t - \beta(t+1)$, and where expectational errors η_{it} are

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

The first line is definitional, and the second line (derived in appendix D.2) shows how expectational errors depend on expectations over future state variables. If sites have rational expectations, then these expectational errors are mean-zero and orthogonal to sites' period-t information sets. That is, rational expectations are correct on average and use all available information.

Operationally, I regress the left-hand-side variables of equation 8, which I read directly from the data, on the interaction of yields and prices, cost factors, regional fixed effects, and regional time trends. I instrument for yields and prices with potential yields and demand shifters as discussed above, and cost factors include port, road, and urban distances, as well as carbon stocks. I instrument with lagged values – variables in sites' period-t information sets – because contemporaneous values are mechanically correlated with the expectational error. This exposition assumes that production begins one period after investment, but in estimation I impose the typical three-year lag for palm maturity.[8] Figure 4 plots the time-series variation in world prices, and figure 3a plots the spatial variation in yields. The regression combines both sources of variation for identification: intuitively, higher prices are more valuable for sites that produce more palm oil. And since revenues $Y_{it+1} P_{t+1}$ are measured directly, parameters γ, κ_m, and α_m are interpretable in dollar terms.

Step 3: Estimating the extensive-margin model (mill construction)

Discreteness precludes the use of a first order condition and the envelope theorem. Instead, I obtain an Euler equation by differencing and finite dependence. I compare sequences of actions,

[8] 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 between developing today and tomorrow. Time to build is exogenous in this setting and thus simpler to accommodate than it is in Kalouptsidi (2014).

with differences in likelihoods reflecting differences in payoffs. Finite dependence facilitates the comparison: under finite dependence, I can choose sequences that lead to common states – and therefore common payoffs – in all future periods (Arcidiacono and Miller 2011).

As before, I compare investing today and tomorrow. More precisely, I compare two sequences of extensive- and intensive-margin 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 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 because, for both sequences, by period $t+2$ the mill is constructed and plantation size is $a^*_{it} + a^*_{it+1}$. To form the Euler equation, I first evaluate the payoffs for each sequence.

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

The continuation values align: $\beta^2 \mathbb{E}^e_{it}[V(a^*_{it} + a^*_{it+1}; \boldsymbol{w}_{it+2}, \varepsilon_{it+2})] = \beta^2 \mathbb{E}^e_{it}[V(a'_{it+1}; \boldsymbol{w}_{it+2}, \varepsilon_{it+2})]$ because $a'_{it+1} = a^*_{it} + a^*_{it+1}$. I then write these payoffs in terms of choice-specific conditional value functions $v^e(1; \boldsymbol{w}_{it})$ and $v^e(0; \boldsymbol{w}_{it})$, which the Hotz-Miller inversion links to choice probabilities.

$$\ln\left(\frac{p^e(\boldsymbol{w}_{it})}{1 - p^e(\boldsymbol{w}_{it})}\right) = v^e(1; \boldsymbol{w}_{it}) - v^e(0; \boldsymbol{w}_{it})\,, \tag{9}$$

as follows from equation 6 (Hotz and Miller 1993).

For the first sequence, $v^e(1; \boldsymbol{w}_{it}) = v^e(1, a^*_{it}, a^*_{it+1}; \boldsymbol{w}_{it})$ by definition, where $a^*_{it} \equiv a^*_{it}(0; \boldsymbol{w}_{it}, \varepsilon_{it})$ and $a^*_{it+1} \equiv a^*_{it+1}(a^*_{it}; \boldsymbol{w}_{it+1}, \varepsilon_{it+1})$. For the second sequence, $v^e(0, 1, a'_{it+1}; \boldsymbol{w}_{it})$ involves choices that may differ from the optimal choices implied by $v^e(0; \boldsymbol{w}_{it})$. The difference in payoffs is

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

where $a^*_{it+1} \equiv a^*_{it+1}(0; \boldsymbol{w}_{it+1}, \varepsilon_{it+1})$, as derived in appendix D.3. Substituting into equation 9, I obtain an Euler equation in which continuation values cancel. Applying the functional forms of revenues and costs, and noting $a^*_{it+1}(a^*_{it}; \boldsymbol{w}_{it+1}, \varepsilon_{it+1}) = a^*_{it+1}(0; \boldsymbol{w}_{it+1}, \varepsilon_{it+1})$ given linear revenues,

$$\ln\left(\frac{p^e(\boldsymbol{w}_{it})}{1 - p^e(\boldsymbol{w}_{it})}\right) - \beta \mathbb{E}^e_{it}\big[\ln p^e(\boldsymbol{w}_{it+1})\big] = \mathbb{E}^e_{it}[I_{it+1}] - (1-\beta) x_i \gamma^e - (1-\beta)\kappa^e_m - \alpha^e_m \tilde{t}\,,$$

for $\tilde{t} = t - \beta(t+1)$ and $I_{it+1} = [\beta Y_{it+1} P_{t+1} - (1-\beta) x_i \gamma - (1-\beta)\kappa_m - \alpha_m \tilde{t}\,] a^*_{it} + \delta[-\frac{1}{2} {a^*_{it}}^2 + \beta a^*_{it} a^*_{it+1}]$. Intuitively, developing earlier brings forward plantation revenues, but also investment costs.

I apply expectational errors η^e_{it} and substitute estimated values to obtain a regression equation.

$$\ln\left(\frac{\widehat{p^e}(\boldsymbol{w}_{it})}{1 - \widehat{p^e}(\boldsymbol{w}_{it})}\right) - \beta \ln \widehat{p^e}(\boldsymbol{w}_{it+1}) = \widehat{I}_{it+1} - (1-\beta) x_i \gamma^e - (1-\beta)\kappa^e_m - \alpha^e_m \tilde{t} + \eta^e_{it} \tag{10}$$

I estimate conditional choice probabilities $\widehat{p^e}(\boldsymbol{w}_{it})$ from the data. In particular, I use the predicted values from a logit regression of observed investment choices on a flexible set of basis terms: piecewise linear splines in Y_{it+1}, P_{t+1}, x_i, and $\tilde{t}$, as well as their interactions. I do so separately for each region and therefore account non-parametrically for regional heterogeneity. Consistent with the model, this procedure accommodates unobserved heterogeneity by region while allowing only observed heterogeneity by site. I estimate intensive-margin choices $\widehat{a}^*_{it}$ in the same way, but with OLS instead of a logit regression. Intensive-margin profits $\widehat{I}_{it+1}$ are a function of these intensive-margin choices, observables, and intensive-margin parameters estimated previously, and are denominated in dollars because $Y_{it+1}P_{t+1}$ measures revenues directly. This denomination provides a scale normalization, such that parameters γ^e, κ^e_m, and α^e_m are interpretable in dollar terms. Intercepts κ^e_m are only identified relative to the outside option, as is typical with discrete choice models.

Discussion

This Euler approach to estimation has several advantages. I can address endogeneity concerns using standard instrumental variable techniques because estimation reduces to linear regression. Furthermore, while I do need to assume that agents have rational expectations, for estimation I do not need to model exactly what these expectations are. This flexibility is a significant advantage over a conventional full-solution approach that would require explicit structure on expectations. The full-solution approach also requires solving the model repeatedly for estimation, with each iteration involving the time-consuming calculation of continuation values. 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 nonstationarity of the problem in my setting (Aguirregabiria and Mira 2007; Bajari et al. 2007; Pakes et al. 2007; Pesendorfer and Schmidt-Dengler 2008).

One disadvantage is that rational expectations can still be a strong assumption. Biased expectations load onto costs, with pessimism over future prices having the same effect as high costs. Regional effects κ_m capture cost heterogeneity across regions and therefore absorb expectational bias to the extent that it is fixed within regions. This approach is similar to Diamond et al. (2017), who difference out expectational bias by assuming that it is constant among individuals within a group. For counterfactuals, the assumption is that expectational bias remains uninfluenced by trade policy. A more careful treatment of expectations would require separate variation in actual and expected profits, as well as specifying how trade policy changes expectations.

Another disadvantage is that tractability relies on several assumptions that may also be strong. First, the Euler comparison between investing today or tomorrow implicitly assumes property rights. If delaying investment risks losing land claims, then sites will be biased toward investing today. Regional effects κ_m also help here: low costs make delayed investment less appealing, and so regions susceptible to land grabbing will appear to have low costs. Second, I assume

Table 2: Rainfall shocks as price instruments

	All	All	Palm	Other
Rain, deviation from optimal (100 mm)	0.139***	0.114***	0.117***	0.0828
	(0.0374)	(0.0328)	(0.0354)	(0.0656)
Oil FE	x	x		
Oil-year trend		x		
Year trend			x	x
Observations	74	74	37	37
F-statistic	13.92	12.11	10.91	1.594

The outcome variable is log price of a given oil product in a given year. Rainfall is constructed at the oil-product level by aggregating rainfall across producing regions, weighting by total production over the study period. For a given region, I measure the total absolute deviation from optimal monthly rainfall levels over the course of the growing season. The first two columns pool across all oil products, controlling for oil product fixed effects, while the last two columns show results for each oil product separately. Palm oil aggregates palm and palm kernel oil, while "other" oils include coconut, olive, rapeseed, soybean, and sunflower oil. Data are annual and cover 1980 to 2016. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

independent, atomistic sites because finite dependence does not hold otherwise. If a price-maker delays investment, then competitors will respond, thereby changing the state of the economy in all future periods. Independence also rules out spatial competition. In particular, although sites are all price-takers in the global palm oil market, even small sites can be price-makers in local input markets for land, labor, and capital. Furthermore, spatial interdependence introduces a dimensionality problem that makes estimation intractable. Third, I assume that plantation age does not affect profits, again such that finite dependence holds. If younger plantations are more productive, then delaying investment changes profits in all future periods.

6 Results

This section describes both demand and supply estimates. Demand estimates suggest inelastic Indonesian and Malaysian demand, while supply estimates quantify palm oil production costs.

6.1 Demand estimates

Table 2 shows the first stage for the price instrument. Rainfall shocks significantly increase world oil prices. The first two columns pool data across oil products and show that the rainfall coefficient is significant with and without year trends. The last two columns show estimates for palm and other oils separately. Smaller sample sizes mean less precision, but the point estimates are similar to those of the pooled specifications. Newey-West standard errors account for serial correlation, and the F-statistics reflect this correction. Appendix table C1 considers alternative instruments. Temperature is not a strong instrument because palm oil is grown in tropical climates

Table 3: Demand elasticities for vegetable oils

		Estimates		SEs	
		Palm	Other	Palm	Other
EU	Palm	-1.06***	1.56**	(0.28)	(0.73)
	Other	0.04	-0.22	(0.25)	(0.17)
China/India	Palm	-0.74***	0.96	(0.24)	(0.73)
	Other	0.01	-0.50***	(0.26)	(0.14)
Other importers	Palm	-0.66**	0.82	(0.26)	(0.58)
	Other	0.30	-0.42***	(0.23)	(0.16)
Indonesia	Palm	-0.11	0.04	(0.88)	(1.19)
	Other	0.31	-0.37	(0.90)	(0.39)
Malaysia	Palm	-0.02	-0.31	(0.31)	(0.43)
	Other	-0.93	1.51	(0.88)	(1.98)

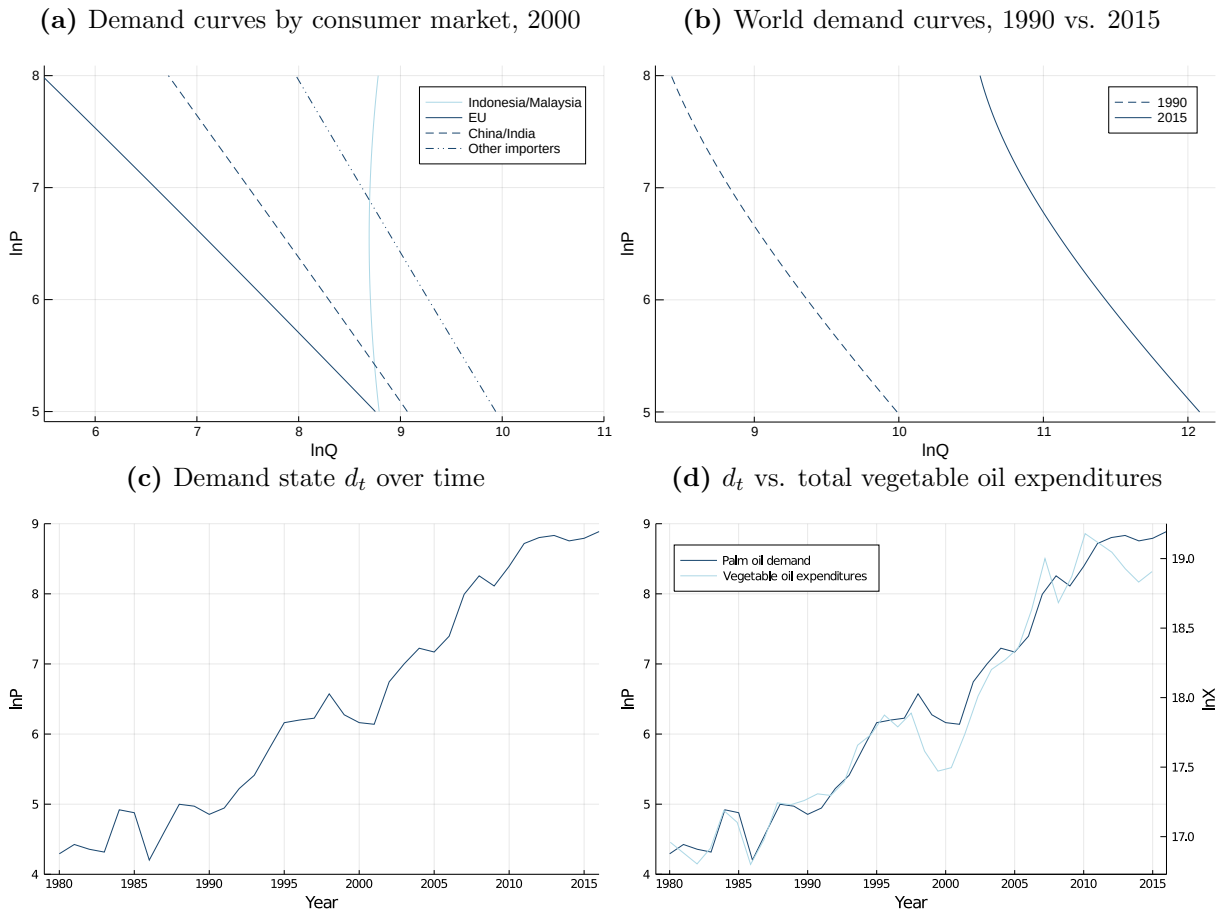
Uncompensated price elasticities are computed from estimated demand parameters using equation 3. Palm oil aggregates palm and palm kernel oil, while "other" oils include coconut, olive, rapeseed, soybean, and sunflower oil. I evaluate expenditure shares, prices, and the time trend at their averages over the study period. I instrument for prices with foreign rainfall shocks. Data are annual and cover 1980 to 2016. Standard errors are computed with the delta method, and I apply a Prais-Winsten transformation to account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

where year-to-year variation in temperatures is limited. Allowing for asymmetric effects, negative rainfall shocks have larger effects than positive shocks, but the difference is not significant.

Table 3 shows the estimated demand elasticities by consumer market, and figure 6 plots the resulting demand curves and implied demand shifters. Appendix tables C2 and C3 present the lower- and upper-level estimates. I obtain reasonable estimates with negative own-price elasticities and positive cross-price elasticities for elasticities that are statistically significant. For Malaysia, elasticities for other oils are noisy because other oils account for only 3% of consumption in the data. I find that palm oil demand among non-EU importers does respond to prices, suggesting unilateral EU action is susceptible to leakage. While non-EU demand elasticities are below one in magnitude, leakage concerns apply as long as unregulated demand is less than perfectly inelastic. Furthermore, non-EU importers account for a large proportion of global palm oil consumption – 68% in table 1. On the other hand, leakage concerns are limited for Indonesia and Malaysia, which have demand elasticities that are statistically indistinguishable from zero.

To address exclusion restriction concerns, I omit domestic rainfall shocks when constructing instruments for estimation. Appendix table C4 shows that, for each consumer market, this leave-out estimator produces estimates similar to those in table 2. I also assess the effects of foreign rainfall shocks on domestic incomes and consumption patterns. Appendix table C5 shows that foreign rainfall shocks have no effect on either GDP or total consumption expenditures, both nationally

Figure 6: Palm oil demand



For each consumer market, figure 6a plots the palm oil demand curves implied by the estimates in table 3. Figure 6b aggregates over market-specific demand to obtain world demand. Figure 6c plots the marginal willingness to pay for the quantity $\ln Q = 10$ in each period. These values capture the shifting of the demand curve over time. Figure 6d compares these values to total vegetable oil expenditures.

and for households. These results also provide reassurance that the instruments do not simply capture idiosyncratic fluctuations in macroeconomic conditions.

A potential shortcoming is that I ignore dynamics on the demand side. Dynamics are important if consumers stockpile vegetable oils when prices are low, or if consumption is sticky because of taste preferences. For stockpiling, I observe oil stocks held in storage facilities and find that they are small at 12.5% of consumption by volume. For taste preferences, stickiness is unlikely because palm, soybean, rapeseed, and sunflowerseed oils – 94% of consumption by volume – are very similar products. On the other hand, long-term contracts or firm relationships could induce stickiness. In addition, over the long term, unregulated markets could respond to palm oil tariffs by developing new industries to supply regulated markets with palm oil in other forms. For example, unregulated markets might produce palm-based biofuels for EU consumption. While this static demand system does capture short-term increases in unregulated input demand for palm oil, such as from existing

Table 4: 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 × price ($Y_{it+3}P_{t+3}$)	0.113***	0.200***	
	(0.00714)	(0.0355)	
Potential yield × demand ($Y_i^p d_t$)			30.58***
			(1.220)
Province FE	x	x	x
Province-year trend	x	x	x
Observations	17,181	17,181	17,181
F-statistic			628

Each column is one regression, and each observation is a site-year. The dependent variables are shown in the column headings. The first column is OLS, and the second column IV. The IV specification uses the interaction of potential yields and demand shifters to instrument for the interaction of yields and prices, with the third column showing the first stage. Potential yields are computed using the agronomic model of Hoffmann et al. (2014). Demand shifters are computed during demand estimation. Prices combine palm and palm kernel oil prices and are inflation-adjusted to year 2000 dollars. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

industries that expand production, it does not capture long-term increases from new industries. This concern is mitigated by the fact that most palm oil is exported in raw form, but nonetheless a simple response is for import tariffs to cover both palm oil and palm oil content.

Moving forward, I focus on demand for palm oil and ignore the carbon effects of substitution to other oils. I therefore do not account for substitution to South American soybean oil, which involves Amazonian deforestation. Two facts mitigate the resulting bias. First, South American soybean oil is just one of several close substitutes for palm oil: soybean oil is only 32% of total oil consumption, and South America supplies less than half of soybean crops globally. Second, South American soy does not destroy peatlands. Amazonian peatlands are concentrated deep in the Amazon, while deforestation is primarily at the Amazon's outskirts (Gumbricht et al. 2017; Song et al. 2018). Thus, the carbon consequences are smaller than those of palm oil, and indeed palm oil emissions would be five to ten times smaller without peatland destruction.

6.2 Supply estimates

Tables 4 and 5 present supply estimates. Table 4 shows that higher revenues – whether they be from higher yields or higher prices – indeed lead to more development, with a larger effect in the IV specification. Table 5 shows the estimated model parameters, all of which are interpretable in dollar terms. On the intensive margin, I estimate the average lifetime costs of plantation development to be \$10 thousand per hectare in net-present-value terms, ranging from \$6 thousand at the 10th percentile to \$15 thousand at the 90th percentile across provinces. By comparison, accounting estimates suggest costs of \$7 thousand per hectare: \$4 thousand upfront and \$3 thousand for

Table 5: Supply model parameter estimates

	Mean	SE	10th percentile	90th percentile
Province-specific costs (κ_m)	9,674***	(856)	6,398	14,655
Province-specific cost trends (α_m)	-374***	(21)	-729	-99
Quadratic costs (δ)	4.50***	(0.80)	–	–
Cost factors (γ)				
Log port distance, km	-711	(486)	–	–
Log road distance, km	-333*	(199)	–	–
Log urban distance, km	-278	(278)	–	–
Log carbon in tree biomass, t	206	(540)	–	–
Log carbon in peat deposits, t	-93	(68)	–	–
Province-specific costs (κ^e_m)	22,881,886***	(391,964)	15,804,464	29,636,816
Province-specific cost trends (α^e_m)	88,477***	(15,261)	-483,608	625,779
Logit scale (σ^e)	3,075,006***	(107,831)	–	–
Cost factors (γ^e)				
Log port distance, km	685,682***	(194,359)	–	–
Log road distance, km	506,299***	(88,269)	–	–
Log urban distance, km	267,636**	(129,626)	–	–
Log carbon in tree biomass, t	706,172***	(174,548)	–	–
Log carbon in peat deposits, t	835	(30,598)	–	–

The top panel shows intensive-margin parameters, and the bottom panel shows extensive-margin parameters. All estimates are interpretable in dollar terms (inflation-adjusted to year 2000 dollars). For region-specific parameters, I include the 10th and 90th percentiles for estimates across regions. I report province-specific costs κ_m and κ^e_m for a mean year and at mean values for cost factors. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

operations (Butler et al. 2009). I estimate costs to be decreasing at an annual rate of $400 on average, again with some heterogeneity across provinces. Within provinces, I find costs to be similar across sites with different characteristics (conditional on mill construction).

On the extensive margin, I estimate lifetime mill construction costs of $23 million on average, ranging from $16 million at the 10th percentile to $30 million at the 90th percentile. By comparison, accounting estimates suggest costs of $20 million: $5 million upfront and $15 million for operations (Man and Baharum 2011). I estimate costs to be increasing at an annual rate of $88 thousand on average, with large heterogeneity across provinces. I estimate the standard deviation of the logit shock to be $3 million, suggesting that changing producer behavior requires incentives measured in the millions of dollars. Within provinces, site characteristics have a significant impact on costs, which are increasing in distances from major ports, roads, and urban centers. This transportation-cost effect is smallest for distance from urban areas, with higher transportation costs partially offset by lower land and labor costs in remote regions. Furthermore, tree biomass does discourage mill construction, as entering heavily forested areas demands significant effort in land development and may face scrutiny from local governments and native populations. But peat deposits – the main

source of carbon emissions – have little effect on mill construction. Indeed, palm oil producers fail to internalize their carbon externalities.

7 Counterfactuals: Assessing Coordination and Commitment

This section evaluates the individual and combined roles of coordination and commitment in determining the impacts of import tariffs. I find that import tariffs can be an effective substitute for domestic regulation, but only when both coordination and commitment hold.

7.1 Setting tariffs

I set tariffs to maximize social welfare, penalizing emission damages while also weighing consumer surplus from palm oil use and producer surplus for Indonesia and Malaysia. The domestic tax, which serves as a benchmark, is also set to maximize social welfare. Unlike the domestic tax, however, tariffs sidestep domestic obstacles to regulation by directly targeting the prices producers receive in world markets. In particular, prices equalize across markets in each period t.

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

For example, new EU tariffs cause revenues from EU sales to decline relative to other sales, and so producers respond by shifting sales to other markets. I assume zero trade costs for simplicity, but adding exogenous trade costs would be inconsequential because they would be invariant across tariffs. Furthermore, the above equation connects the three components of the empirical model: tariffs, demand, and supply. Tariff τ_t changes world price P_t depending on demand curves $P_t^{Dr}(q)$ and $P_t^{Du}(q)$, and world price P_t in turn affects the investments that produce supply $Q_t^{ro} + Q_t^{uo}$.

I focus on uniform tariffs that treat all palm oil equally. The alternative is to condition on the emissions specific to each unit of palm oil. For example, if palm oil can be certified as green, then tariffs can differentiate by certification status. In practice, however, tracking production histories to this extent is difficult. Similarly, I focus on a uniform domestic tax because of its administrative convenience: it can be applied at the point of sale without the need to monitor production. Indeed, uniform taxes are common despite being "second-best" relative to a Pigouvian tax. For example, fuel taxes are uniform despite heterogeneity across vehicles in tailpipe emissions (Knittel and Sandler 2018). Nonetheless, an alternative is to condition on emissions with site-specific license fees or ex-post penalties, and my framework can readily accommodate such policies.

I quantify the effects of coordination and commitment by studying the following scenarios. For coordination, I study tariffs set under three tariff coalitions: all importers together, an EU-China-India partnership, and the EU alone. For commitment, I study full, no, and limited commitment. Full commitment assumes that, once set, tariffs are upheld in perpetuity. No commitment assumes

that tariffs are reset every period, with the result being sequential static optimization. Limited commitment assumes commitment to L-period tariff plans revised every L periods, similar to "five-year plans" in Indonesia and China or any policy based on decennial census results.

7.2 Solving the model

Counterfactuals require solving the supply model and thus involve an additional set of assumptions over how firms set expectations. I model the non-stationary evolution of demand d_t with an ARIMA process, and I assume firms' expectations are given by this process. Supply s_t is determined endogenously as the result of an entry game in which beliefs are correct in equilibrium. Intuitively, if firms believe all other firms will enter, then they will anticipate low prices and not enter, in which case their beliefs are not consistent with reality. I assume that yields Y_{it} evolve at a constant and exogenous rate per year. Finally, I assume that while firms know current-period cost shocks ε_{it} and ε^e_{it}, they only know the distribution of future shocks. Appendix E.1 discusses these expectations in further detail. Note that estimation does not rely on these assumptions because the Euler approach estimates the model without solving it. And while I do need to solve the model for counterfactuals, I still avoid the computational burden of solving it repeatedly for estimation.

For a given set of tariffs, I solve the model by backward inducting from the steady state. Suppose the steady state is reached in period S. At this point, all feasible lands are developed and there is no further development, but plantations continue to generate revenues over the infinite horizon. Finite lands ensure that such a period exists, but the challenge is that it takes many periods to exhaust all available land. Backward induction over such a long horizon is computationally intensive. I address this computation burden in two ways. First, I solve each subproblem using an iterative algorithm that approximates the solution with a fixed look-ahead horizon instead of always looking ahead to the end of the game tree. This algorithm breaks the usual curse of dimensionality in which the state space grows exponentially in the length of the look-ahead window. Second, I approximate period S by choosing an arbitrary period $T < S$ and solving as if it were the steady state. This approach is biased if substantial development occurs after period T, but I resolve taking periods $T+5$, $T+10$, and so on as the steady state until the solutions converge. Intuitively, entry today becomes less appealing when competitors have a longer window of opportunity to enter, but discounting means a diminishing marginal impact of extending this window. Appendix E.1 documents each step of the solution algorithm.

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. Assuming plantation development releases carbon stocks completely, these data provide a direct link to counterfactual emissions. I also assume that producers' outside option remains

unimpacted by palm oil tariffs. This assumption is strong, but it is typical of work that studies individual industries in detail at the cost of missing cross-industry effects in general equilibrium. For carbon emissions, this assumption imposes that non-palm deforestation does not expand in response to palm oil tariffs. To assess this restriction, I consider mining, selective logging, and acacia (paper pulp) plantations – the other primary drivers of deforestation in Indonesia and Malaysia. The first two are unlikely to generate large bias: mining depends on the exogenous distribution of mining deposits, and selective logging involves limited emissions because it does not destroy peatlands.

Acacia plantations, however, do destroy peatlands, and thus substitution from palm oil to acacia has significant carbon implications. I address this concern by compiling data on the acacia industry and estimating the reduced-form relationship between acacia and palm oil plantation development in appendix E.2. I find the magnitude of the relationship to be small, at least in partial equilibrium, with a one-percent reduction in palm oil plantation development corresponding to a 0.02% increase in acacia plantation development. Capturing general equilibrium effects would require a two-industry model in which producers first choose between palm oil and acacia, then proceed with the extensive- and intensive-margin investment decisions of the baseline model. But oil palm is more profitable than other crops – seven times more so than acacia (Sofiyuddin et al. 2012) – and thus acacia expansion is unlikely to fully offset palm reductions. Conceptually, substitution toward acacia plantations is a source of supply-side leakage that makes tariffs less effective, and the policy response is to levy acacia tariffs alongside palm oil tariffs.

7.4 Import tariffs can be an effective substitute for domestic regulation

Table 6 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. Finally, the magnitude of the emission externality leads to a domestic tax that is itself quite large at several times observed prices.

7.5 But only when both coordination and commitment hold

Emission reductions diminish as coordination and commitment weaken. Figure 7a plots emission reductions under each of the scenarios in table 6. First, weak coordination decreases the level of achievable emission reductions because importers have relatively elastic demand. Emissions fall by at most 56% under full coordination among all 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 consumption, respectively – because leakage

Table 6: Counterfactual experiments

Experiment	$/t NPV	Δ%	Δ% 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

The first column shows the net present value of taxes or tariffs in dollars per ton of palm oil. The second column shows percentage changes in carbon emissions relative to observed net present values, and the third, fourth, fifth, and sixth columns show percentage changes in surplus by market. Figures for Indonesia and Malaysia combine consumer and producer surplus, and all figures include government tax or tariff revenue where applicable. The last column shows average social surplus losses per ton of carbon averted. 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. Each shows several commitment scenarios: full commitment over all future periods, and limited commitment in which commitment is only for five, ten, or twenty years at a time. Under no commitment, tariffs have no effect because they do not affect new development given time to build. The discount factor is $\beta = 0.9$.

concerns lead to smaller tariffs. Second, weak commitment can significantly undermine emission reductions. This effect is especially stark when the commitment period does not exceed time to build, in which case tariffs and emission reductions are zero. In this case, tariffs have no effect on new development because new development does not generate taxable production until after the commitment period has passed. Third, coordination and commitment interact. Figure 7b shows how weak coordination increases the importance of commitment. A five-year commitment period achieves 95% of full-commitment outcomes when all importers coordinate, but does much less under an EU-China-India coalition or unilateral EU action. These scenarios instead require twenty-year commitment periods to approach full-commitment outcomes.

The division of surplus highlights why coordination and commitment are difficult to achieve when countries focus only on their individual outcomes. For coordination, importers gain by defect-

Figure 7: Counterfactual emissions

(a) Emission reductions **(b)** Relative to full commitment

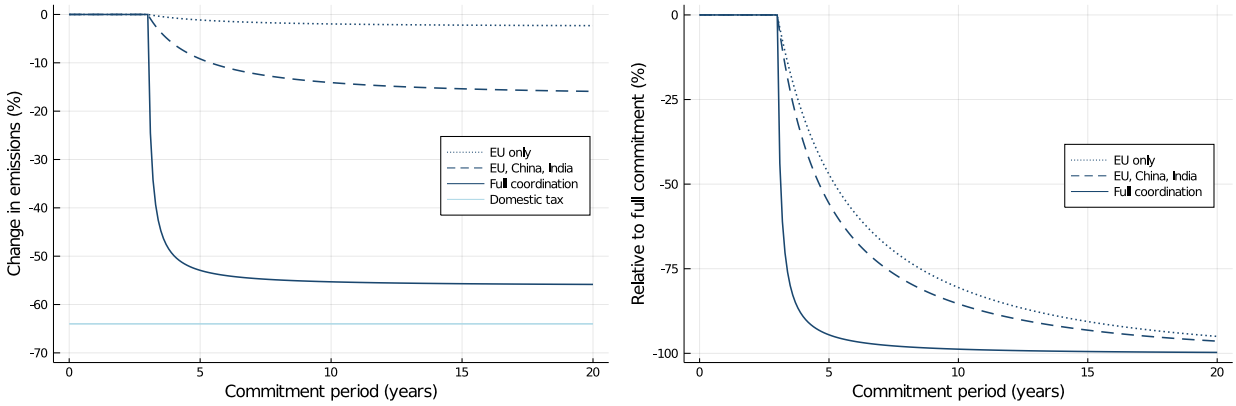


Figure 7a shows emission reductions under several scenarios. Starting at the top, the dotted line shows reductions under unilateral EU action for each of the commitment periods listed on the x-axis. Emission reductions are zero when the commitment period does not exceed time to build because otherwise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

ing from the tariff coalition because they can free-ride on the emission reductions that the coalition achieves. Furthermore, defectors benefit from leakage as the tariff coalition cuts its consumption and world prices fall in response. For example, focusing on full commitment, other importers have 25% lower consumer surplus when they join the EU, China, and India in imposing tariffs, but 45% higher consumer surplus when they unilaterally defect. For commitment, acting importers have higher surplus when commitment levels are low because low commitment leads to low tariffs and thus limited 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.

More broadly, the same considerations underscore why Indonesia and Malaysia have not implemented the socially optimal domestic tax. If importers cannot coordinate, then the domestic tax greatly reduces producer surplus, leaving Indonesia and Malaysia better off accepting import tariffs. But if importers threaten coordinated tariffs, then the domestic tax becomes appealing. It has low marginal impact on producer surplus since coordinated tariffs are already high, and it generates government revenue that would otherwise go to foreign governments.

7.6 Robustness and extensions

Table 7 shows that the qualitative results hold across a series of robustness checks. First, the baseline model assumes a discount factor of $\beta = 0.9$, but effects are larger for lower discount factors, which imply larger benefits from delaying development. Second, the baseline result relies on

Table 7: Robustness and extensions, carbon emission reductions ($\Delta\%$ CO_2)

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						
$\varepsilon^{DI}, \varepsilon^{DM} = 0.22$	-50	-43	-13	-7	-2	-1
$\varepsilon^{DI}, \varepsilon^{DM} = 0.44$	-44	-34	-10	-5	-2	-1
$\varepsilon^{DI}, \varepsilon^{DM} = 0.66$	-39	-28	-8	-4	-1	-1
Limited commitment, early planning	-56	-55	-16	-14	-2	-2
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 experiment. The first panel corresponds to table 6. The second panel changes the discount factor. The third panel changes the elasticities of Indonesian and Malaysian demand, where 0.66 is the demand elasticity for other importers. The fourth panel allows planning for the next L-year plan under limited commitment to begin before the end of each plan. This early planning prevents tariffs from being set to zero at the beginning of each L-year tariff regime. The fifth panel assumes tariffs are set to maximize the surplus of the acting coalition, net of its own costs of carbon as computed by Ricke et al. (2018). The sixth panel allows import tariffs to condition on the emissions specific to each unit of palm oil. The last panel assumes a static supply model.

inelastic demand in Indonesia and Malaysia, but elastic demand increases leakage and lowers carbon reductions, although coordinated, committed tariffs continue to have large effects. Third, I allow importers under limited commitment to revise their L-year plans several years before the end of each plan. Early planning prevents tariffs from being set to zero at the start of each regime because of time to build, and thus lessens the difference between long and short commitment periods.

I also consider other extensions. First, I set tariffs to maximize the acting coalition's welfare rather than social welfare. I assume the coalition considers only its own proportion of the social costs of carbon: 1%, 17%, 80%, and 2% for the EU, China/India, other importers, and Indonesia/Malaysia, respectively, based on pooling the country-level estimates of Ricke et al. (2018). When importers coordinate, tariffs rise because they improve terms of trade – importers no longer value Indonesian and Malaysian producer surplus – and importers internalize nearly the full social cost of carbon. When importers do not coordinate, tariffs decline sharply because small coalitions internalize only a small part of the social cost of carbon. Second, baseline tariffs are uniform across all units of palm oil, but conditioning on unit-specific emissions leads to larger carbon reductions by more efficiently targeting peatland destruction. A non-uniform domestic tax achieves similar

gains: a carbon reduction of 91% relative to 64% in the baseline. Finally, a static supply model leads to low supply elasticities and much smaller effects for tariffs. 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 an empirical framework for quantifying these forces, and I apply it to the market for palm oil. The palm oil industry is of first-order importance: deforestation for palm oil plantations accounts for more CO_2 emissions over the last three decades than the entire economy of India. My framework quantifies the extent to which import tariffs could have reduced these emissions. It accounts for leakage to unregulated consumer markets, and it captures firms' dynamic considerations over sunk investment in palm oil plantations and mills. Using data from satellite imagery, it delivers predictions of plantation development – and therefore deforestation – at a fine level of spatial disaggregation.

I find coordinated, committed trade policy to be effective, reducing carbon emissions by 56% compared to 64% under domestic regulation. In the case of Europe, where recent legislation has targeted palm oil imports, EU import 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. 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.

I leave several directions open for future work. First, despite its environmental consequences, oil palm yields significantly more oil per hectare than any comparable oilseed. Future work might take a global view of oilseed production and account more explicitly for substitution to other oilseed crops, including soybeans from the Amazon. Second, given my estimates of the social welfare gains from coordination, future work might study the dynamic bargaining game that the EU, China, and India face in deciding whether to form a coalition. Lastly, spatial interaction across palm oil plantations might create path dependence, which tariffs can leverage to protect carbon-rich regions by conditioning on where a given unit of palm oil is produced.

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.

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.

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.

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.

Butler, Rhett, Lian Pin Koh, and Jaboury Ghazoul. REDD in the Red: Palm Oil Could Undermine Carbon Payment Schemes. Conservation Letters, 2:67–73, 2009.

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.

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.

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.

Ericson, Richard and Ariel Pakes. Markov-Perfect Industry Dynamics: A Framework for Empirical Work. Review of Economic Studies, 62(1):53–82, 1995.

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.

Gumbricht, 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 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 Differenciated 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.

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.

Holmes, Thomas. The Diffusion of Wal-Mart and Economies of Density. Econometrica, 79(1):253–302, 2011.

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.

Illanes, Gastón. Switching Costs in Pension Plan Choice. 2017.

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.

Kalouptsidi, Myrto. Time to Build and Fluctuations in Bulk Shipping. American Economic Review, 104(2): 564–608, 2014.

Kalouptsidi, Myrto, Paul Scott, and Eduardo Souza-Rodrigues. Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models. 2018.

Knittel, Christopher and Ryan Sandler. The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation. American Economic Journal: Economic Policy, 10(4):211–242, 2018.

Koplitz, Shannon, Loretta Mickley, Miriam Marlier, Jonathan Buonocore, Patrick Kim, Tianjia Liu, Melissa Sulprizio, Ruth DeFries, Daniel Jacob, Joel Schwartz, Montira Pongsiri, and Samuel Myers. Public Health Impacts of the Severe Haze in Equatorial Asia in September–October 2015: Demonstration of a New Framework for Informing Fire Management Strategies to Reduce Downwind Smoke Exposure. Environmental Research Letters, 11:094023, 2016.

Kortum, Samuel and David Weisbach. The Design of Border Adjustments for Carbon Prices. National Tax Journal, 70(2):421–446, 2017.

Magnac, Thierry and David Thesmar. Identifying Dynamic Discrete Decision Pro- cesses. 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.

Morales, Eduardo, Gloria Sheu, and Andrés Zahler. Extended Gravity. Review of Economic Studies, 86: 2668–2712, 2019.

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, A., J. Porter, Kate Ho, and Joy Ishii. Moment Inequalities and Their Application. Econometrica, 83 (1):315–334, 2015.

Pakes, Ariel. Alternative Models for Moment Inequalities. Econometrica, 78(6):1783–1822, 2010.

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.

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 Bank. The Cost of Fire: An Economic Analysis of Indonesia's 2015 Fire Crisis. Technical report, World Bank Group, Washington, DC, 2016.

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.

A Appendix: Illustrative Model of Emission-Based Trade Policy

I derive optimal tariffs, illustrating the leakage and commitment problems, and I consider extensions for heterogeneous emissions and terms-of-trade effects.

A.1 Import tariffs under incomplete regulation and sunk investment

Domestic regulation

In the absence of an unregulated market, I denote the total inverse demand curve by $P_t^D(q)$. 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,

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

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

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

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

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

where upfront tax $\tilde{\tau}_t$ only directly affects contemporaneous new development Q_t^n, and where I apply the envelope theorem in ignoring second-order effects on new development in future periods.

The leakage problem

To isolate the leakage problem, I first suppose that importers are able to impose tariff $\tilde{\tau}_t$ on development directly, as is possible under domestic regulation. The difference is that producers can choose between regulated market r and unregulated market u. Social welfare is

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

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

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

Development tariff $\tilde{\tau}_t^{\mathrm{L}}$ only directly affects new development Q_t^n. Assuming an interior solution, the first order condition and resulting tariff are

$$\frac{dW_t}{d\tilde{\tau}_t} = (\tilde{\tau}_t - e)\frac{dQ_t^{rn}}{d\tilde{\tau}_t} - e\frac{dQ_t^{un}}{d\tilde{\tau}_t} = 0\,, \quad \tilde{\tau}_t^{\mathrm{L}} = \left(\frac{\varepsilon_t^S}{\varepsilon_t^S - \frac{Q_{t+1}^{uo}}{Q_{t+1}^{o}}\varepsilon_{t+1}^{Du}}\right)e < \tilde{\tau}_t^{\mathrm{FB}}\,, \tag{11}$$

for elasticities $\varepsilon_t^S > 0$ and $\varepsilon_{t+1}^{Du} < 0$ evaluated at quantities $Q_{t+1}^o \equiv Q_{t+1}^{o*}(\tilde{\tau}_t^{\mathrm{L}})$ and $Q_{t+1}^{uo} \equiv Q_{t+1}^{uo*}(\tilde{\tau}_t^{\mathrm{L}})$, 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 $\tilde{\tau}_t^{\mathrm{L}} = \tilde{\tau}_t^{\mathrm{FB}}$.

The leakage problem is limited when supply is elastic or unregulated demand is inelastic. In the first case, tariffs have limited effects on world prices; in the second case, world prices do fall but unregulated consumption does not increase in response. In both cases, tariffs do not affect unregulated consumption, and so they approach the size of the first-best tax. The leakage problem is also limited when the unregulated share of consumption is small. Conversely, elastic unregulated demand leads to a severe leakage problem and pushes tariffs to zero. Tariffs also go to zero when supply is inelastic, in which case tariffs produce allocative inefficiency without reducing emissions.

The commitment problem

In reality, importers cannot impose an upfront tax $\tilde{\tau}_t$ directly on new development Q_t^n. Rather, they can only target individual units of consumption at each point in time. This constraint has two consequences. First, given time to build, this constraint means that tariffs today cannot target new development directly. Time to build implies that new development does not begin production until the next period, and so this new development is unaffected by tariffs on consumption today. New development is instead governed by the stream of future tariffs $\{\tau_{t+1}, \tau_{t+2}, \ldots\}$. Second, the allocation of consumption between markets can shift from period to period depending on the tariffs in place. This shifting occurs because producers reallocate sales toward higher-priced markets in each period until the prices they receive are equalized. Such reallocation does not occur with upfront tax $\tilde{\tau}_t$ because producers that have paid taxes upfront have no further cost of selling to the regulated market and therefore no incentive to reallocate sales.

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

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

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

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

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

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

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.

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

In the absence of leakage, there is no such problem: $\frac{dQ_t^{ro}}{d\tau_t} = 0$, and the first order condition is satisfied without setting tariffs to zero.

Under limited commitment, I assume that importers commit to L-period tariff plans that get revised every L periods. Indeed, this scenario is common in practice: 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.

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

The remaining tariffs are set anticipating these periodic breaks. With the goal of highlighting intuition and obtaining manageable closed-form expressions, I simplify the problem by assuming that the demand and supply curves are constant over time. I relax this simplifying assumption in the empirical implementation by solving numerically.

Under time-invariant demand and supply curves, the problem simplifies because the non-zero tariffs will also be time-invariant. To see why, note that the first order condition for a tariff τ_{t+s} is

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

nesting $\frac{dW_t}{d\tau_{t+s}} = \tau_{t+s}\frac{dQ_{t+s}^{ro}}{d\tau_{t+s}} = 0$ given $\frac{dQ_{t+s}^{o}}{d\tau_{t+s}} = 0$ for $s \in \{0, L, 2L, \ldots\}$. 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 $s \notin \{0, L, 2L, \ldots\}$. 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 and generating revenues for the interceding periods is more profitable: flow profits do not decrease over time because demand, supply, and tariffs are fixed. Thus, development in a later period is not profit-maximizing.[9]

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 can shift sales away from the regulated market where tariffs are in place, and toward the regulated market when they are not.

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

[9] A benefit of developing later is that it delays development costs. But if firms prefer to delay, then they will do so forever given constant supply and demand over time. In this case, developing later is not optimal to begin with.

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

$$P^S\big(Q^{o*}_{t+1}(\tau)\big) = \left(\frac{\beta}{1-\beta} - \frac{\beta^L}{1-\beta^L}\right) P^{Du}\big(Q^{uo*}_{t+1}(\tau)\big) + \frac{\beta^L}{1-\beta^L} P^{Du}\big(Q^{uo*}_{t+L}(\tau)\big)\,,$$

$$P^{Dr}\big(Q^{ro*}_t(\tau)\big) - \tau_t = P^{Du}\big(Q^{uo*}_t(\tau)\big) \quad \forall\, t, \text{ given } \tau_{t+s} = \begin{cases} 0 & \text{for } s \in \{0, L, 2L, \ldots\} \\ \tau & \text{otherwise} \end{cases},$$

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

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

Assuming an interior solution, the net present value of the stream of tariffs given by τ is

$$\tilde{\tau}^{\mathrm{LC}}_t(L) = \left(\frac{\varepsilon^S_t}{\varepsilon^S_t - \frac{Q^{uo}_{t+1}}{Q^o_{t+1}}\varepsilon^{Du}_{t+1}\left[1 + \frac{(1-\beta)\beta^L}{\beta-\beta^L}\frac{P^{Du}_{t+L}}{P^{Du}_{t+1}}\left(1 - \frac{Q^o_{t+1}\varepsilon^S_t}{Q^{ro}_{t+L}\varepsilon^{Dr}_{t+1} + Q^{uo}_{t+L}\varepsilon^{Du}_{t+1}}\right)\right]}\right)e\,,$$

for elasticities $\varepsilon^S_t > 0$ and $\varepsilon^{Dr}_{t+1}, \varepsilon^{Dr}_{t+L}, \varepsilon^{Du}_{t+1}, \varepsilon^{Du}_{t+L} < 0$, and quantities and prices evaluated at τ^{LC}. For simplicity I assume constant elasticities of demand. Per-period tariff τ^{LC} is

$$\tau^{\mathrm{LC}}_t(L) = \left(\frac{\beta}{1-\beta} - \frac{\beta^L}{1-\beta^L}\right)^{-1}\tilde{\tau}^{\mathrm{LC}}_t(L)\,.$$

Total tariffs $\tilde{\tau}^{\mathrm{LC}}_t(L)$ are increasing in L, with $L \to \infty$ corresponding to full commitment and $L = 2$ to the minimum binding level of commitment.

$$\tilde{\tau}^{\mathrm{LC}}_t(L) < \left(\frac{\varepsilon^S_t}{\varepsilon^S_t - \frac{Q^{uo}_{t+1}}{Q^o_{t+1}}\varepsilon^{Du}_{t+1}}\right)e = \lim_{L\to\infty}\tilde{\tau}^{\mathrm{LC}}_t(L) = \tilde{\tau}^{\mathrm{C}}_t = \tilde{\tau}^{\mathrm{L}}_t\,.$$

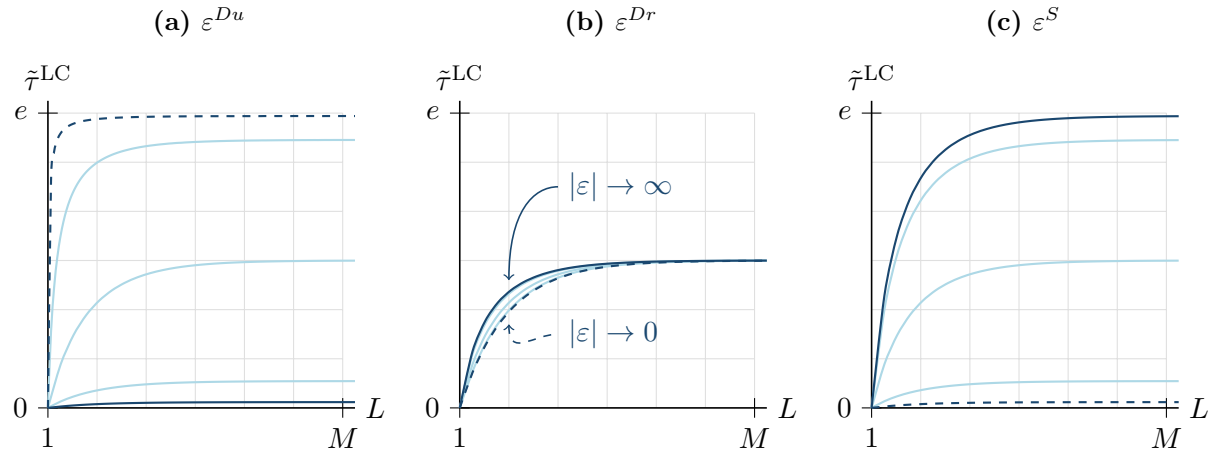
Lastly, the same mechanism also applies in the more general case if tariffs are declining over time. Indeed, importers that take a sequential static approach to setting tariffs will be governed by equations 11, which imply declining tariffs if the elasticity of supply is declining over time. Such will 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 feasible lands are exhausted: at this point, tariffs cannot reduce emissions because prior development is sunk, and no new development is possible. Thus, as tariffs decline, producers will be able to reallocate sales toward the regulated market as shown above.

How leakage and commitment interact

I study how leakage (given ε^{Du}_{t+1}, ε^{Dr}_{t+1}, and ε^S_t) and commitment (given L) interact to determine total tariffs $\tilde{\tau}^{\mathrm{LC}}_t(L)$. First, $\tilde{\tau}^{\mathrm{LC}}_t(L)$ increases more rapidly in L for smaller $|\varepsilon^{Du}_{t+1}|$.

$$\lim_{\varepsilon^{Du}_{t+1}\to 0}\tilde{\tau}^{\mathrm{LC}}_t(L) = e > 0 = \lim_{\varepsilon^{Du}_{t+1}\to -\infty}\tilde{\tau}^{\mathrm{LC}}_t(L)$$

Figure A1: Total tariffs by leakage and commitment



For various values of each leakage-relevant elasticity – namely elasticity of unregulated demand ε^{Du}, elasticity of regulated demand ε^{Dr}, and elasticity of supply ε^{S} – I plot the relationship between total tariffs $\tilde{\tau}^{\text{LC}}$ and the length of commitment L. The solid navy lines show the relationship for large values of the elasticities, the dashed navy lines for small values, and the light blue lines for intermediate values. Each of the values differs by an order of magnitude. Emissions e represents the externality, and M is an arbitrarily large number.

Second, $\tilde{\tau}_t^{\text{LC}}(L)$ increases more rapidly in L for larger $|\varepsilon_{t+1}^{Dr}|$, although this effect is relatively small.

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

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

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

Figure A1 graphs these relationships. As above, the leakage problem is most severe when unregulated demand is elastic or supply is inelastic. The elasticity of regulated demand plays a more limited role.[10]

A.2 Heterogeneous emissions

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

[10] It affects the scope for shifting but not the mechanism itself. In particular, commitment is more important when regulated demand is inelastic, in which case the need to shift toward the unregulated market is small: the tariff displaces only a small quantity, and regulated consumers bear the brunt of the tariff.

difficult and imperfect.[11] I therefore focus on a uniform tariff that treats all goods equally.

Consider incomplete regulation under commitment. The regulator considers consumption utility, for which clean and dirty products are perfect substitutes, and production costs, which vary both privately and socially. I again focus on the simple case of an initial period with no prior development and time-invariant demand and supply curves. Social welfare depends on the consumption of each good in each market.

$$W_1(Q_1^{rc}, Q_1^{rd}, Q_1^{uc}, Q_1^{ud}) = \frac{1}{1-\beta}\int_0^{Q_1^r} P^{Dr}(q)dq + \frac{1}{1-\beta}\int_0^{Q_1^u} P^{Du}(q)dq - \int_0^{Q_1^c}\Big(P^{Sc}(q)+e^c\Big)dq - \int_0^{Q_1^d}\Big(P^{Sd}(q)+e^d\Big)dq\,,$$

where $0 < e^c < e^d$. In equilibrium, new development – clean and dirty – equalizes marginal cost and marginal revenue. The equilibrium conditions bind when sales of a given product to a given market are positive, otherwise marginal cost exceeds marginal revenue. For per-period tariffs τ^k,

$$P^{Sk}\big(Q_1^{k*}(\tau^c,\tau^d)\big) = \frac{1}{1-\beta}\Big(P^{Dr}\big(Q_1^{r*}(\tau^c,\tau^d)\big) - \tau^k\Big) \quad \text{if } Q_1^{rk*}(\tau^c,\tau^d) > 0 \text{ for } k \in \{c,d\}\,,$$
$$P^{Sk}\big(Q_1^{k*}(\tau^c,\tau^d)\big) = \frac{1}{1-\beta}\Big(P^{Du}\big(Q_1^{u*}(\tau^c,\tau^d)\big)\Big) \quad \text{if } Q_1^{uk*}(\tau^c,\tau^d) > 0 \text{ for } k \in \{c,d\}\,,$$

If clean and dirty consumption must face equal tariffs ($\tau^c = \tau^d = \tau$), then all four equilibrium conditions bind simultaneously. The first order condition and optimal tariff are

$$\frac{dW_1}{d\tau} = \left(\frac{1}{1-\beta}\tau - e^c\right)\frac{dQ_1^{rc}}{d\tau} + \left(\frac{1}{1-\beta}\tau - e^d\right)\frac{dQ_1^{rd}}{d\tau} - e^c\frac{dQ_1^{uc}}{d\tau} - e^d\frac{dQ_1^{ud}}{d\tau} = 0\,,$$
$$\tau^{\mathrm{C}} = (1-\beta)\left(\frac{\frac{Q_1^c}{Q_1}\varepsilon^{Sc}}{\frac{Q_1^c}{Q_1}\varepsilon^{Sc} + \frac{Q_1^d}{Q_1}\varepsilon^{Sd} - \frac{Q_1^u}{Q_1}\varepsilon^{Du}}\right)e^c + (1-\beta)\left(\frac{\frac{Q_1^d}{Q_1}\varepsilon^{Sd}}{\frac{Q_1^c}{Q_1}\varepsilon^{Sc} + \frac{Q_1^d}{Q_1}\varepsilon^{Sd} - \frac{Q_1^u}{Q_1}\varepsilon^{Du}}\right)e^d,$$

for $\varepsilon^{Sc}, \varepsilon^{Sd} > 0$, $\varepsilon^{Du} < 0$, and $Q_1 = Q_1^c + Q_1^d = Q_1^r + Q_1^u$. The first best is special case $Q_1^u = 0$.

$$\tau^{\mathrm{FB}} = (1-\beta)\left(\frac{\frac{Q_1^c}{Q_1}\varepsilon^{Sc}}{\frac{Q_1^c}{Q_1}\varepsilon^{Sc} + \frac{Q_1^d}{Q_1}\varepsilon^{Sd}}\right)e^c + (1-\beta)\left(\frac{\frac{Q_1^d}{Q_1}\varepsilon^{Sd}}{\frac{Q_1^c}{Q_1}\varepsilon^{Sc} + \frac{Q_1^d}{Q_1}\varepsilon^{Sd}}\right)e^d > \tau^{\mathrm{C}}$$

In both cases, these "second-best" uniform tariffs are weighted averages of emission levels as in Diamond (1973), with weights given by level-specific supply elasticities.

A.3 Terms-of-trade effects

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

[11] Several certification schemes exist for palm oil, with the Roundtable on Sustainable Palm Oil being most prominent. Two tiers of differentiation – certified or not – is common and insufficient for a Pigouvian tax that differentiates across emission levels. Furthermore, these schemes have their own commitment problems. A common criticism is that they certify palm oil from previously deforested lands on the grounds that it involves no new emissions.

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. For simplicity, I analyze an initial period with no prior development and time-invariant demand and supply curves. For per-period tariff τ under commitment, the objective function is

$$W_1(Q_1^r, Q_1^u) = \frac{1}{1-\beta}\int_0^{Q_1^r}\Big(P^{Dr}(q) - P^{Dr}(Q_1^r) + \tau\Big)dq - \int_0^{Q_1} e dq\,.$$

In equilibrium, marginal entry is indifferent between markets.

$$P^S\big(Q_1^*(\tau)\big) = \frac{1}{1-\beta}\Big(P^{Dr}\big(Q_1^{r*}(\tau)\big) - \tau\Big) = \frac{1}{1-\beta}\Big(P^{Du}\big(Q_1^{u*}(\tau)\big)\Big)$$

Assuming $Q_1^{r*}(\tau), Q_1^{u*}(\tau) > 0$, the first order condition and optimal per-period tariff are

$$\frac{dW_1}{d\tau} = -Q_1^r\frac{dP^{Dr}}{dQ_1^r}\frac{dQ_1^r}{d\tau} + \tau\frac{dQ_1^r}{d\tau} + Q_1^r - (1-\beta)e\frac{dQ_1}{d\tau} = 0\,,$$

$$\tau^{\mathrm{C}} = \underbrace{(1-\beta)\left(\frac{\varepsilon^S}{\varepsilon^S - \frac{Q_1^u}{Q_1}\varepsilon^{Du}}\right)e}_{\text{emissions}} + \underbrace{(1-\beta)\left(\frac{\frac{Q_1^r}{Q_1}P^S}{\varepsilon^S - \frac{Q_1^u}{Q_1}\varepsilon^{Du}}\right)}_{\text{terms of trade}},$$

for quantities $Q_1^k \equiv Q_1^{k*}(\tau)$, prices $P^S \equiv P^{S*}(\tau)$, and elasticities $\varepsilon^S > 0$ and $\varepsilon^{Du} < 0$. The first-best tariff is the special case with $Q_1^u = 0$.

$$\tau^{\mathrm{FB}} = (1-\beta)\left(e + \frac{P^S}{\varepsilon^S}\right) > \tau^{\mathrm{C}}$$

In both cases, I obtain an additional terms-of-trade term, although this term is dominated when the emissions externality is large.

B Appendix: Data

This section lists data sources and discusses the construction of data on palm oil plantations, mills, yields, and carbon stocks.

B.1 Data 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 Plantation development

Data on the expansion of 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.[12] I aggregate the data to the 30-arc-second resolution (approximately 1 km^2) 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, and Malaysia, and I do the same in my analysis. These regions account for 96% of palm oil area harvested in Indonesia and Malaysia and 98% of production in 2016.[13] In particular, I ignore palm oil production in Papua, Sulawesi, and Java. Although Papua and Sulawesi are important frontiers for future expansion, they are small contributors in the period of study.

I extend the plantations data back to 1985 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,

$$\Delta\text{Plantation}_{it} = \sum_{s=0}^{3} \beta_s \Delta\text{Tree cover}_{it-s} + \varepsilon_{it}\,, \tag{12}$$

where $\Delta\text{Plantation}_{it}$ captures new plantation development and the $\Delta\text{Tree cover}_{it-s}$ terms capture tree cover loss in the preceeding periods. The Song et al. (2018) data are measured at a resolution of 5.6km ($0.05° \times 0.05°$), 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

[12] I use the midpoints of the upper and lower bounds in years where bounds are provided, and I use the point estimates in years where bounds are not provided.

[13] I calculate these figures from national data from the USDA Foreign Agricultural Service, combined with regional data from the Indonesian Ministry of Agriculture.

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. I also check for monotonicity over time, and I find this property to hold in the imputed values. The downscaling of the coarser Song et al. (2018) is one point of concern, and it implies that the imputed data should not be analyzed below a resolution of 5.6km (30-km^2 tiles). My core analysis respects this constraint: it centers on sites that are, on average, about 650 km^2, and not on individual tiles.

Figure B1 shows the average proportions of plantation development over time for tiles with nonzero development by 2016 – about 39% of all tiles. The second half of the data is as observed in the Xu et al. (2020) data, and the first half is imputed using the Song et al. (2018) data. For tiles with zero development by 2016, note that I do not need to impute values given that I impose uni-directional development: if a tile has no plantations in 2016, then it must have no plantations in every preceeding year. In figure B2, I verify the quality of the satellite data, both observed and imputed, by comparing them to aggregate data derived from government statistics on palm oil croplands in Indonesia and Malaysia. The data match well, although the satellite data imply more plantations in later years because the assumption of uni-directional development fails to account for field loss, such as from forest fires. Uni-directional development is therefore a simplifying assumption, but it is arguably appropriate in my context because plantation development releases carbon emissions irreversibly. Furthermore, my supply model can handle the risk of future field loss, which enters as an unobserved, region-specific cost.

B.3 Mill construction

Data on 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. Mill locations are recorded by latitude and longitude, and coordinates are manually verified using satellite imagery. I supplement these data with mill locations 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 using satellite imagery. I combine the datasets by merging them spatially, and I validate each mill using Landsat and DigitalGlobe satellite images from Google Earth by looking for nearby oil palm plantations, storage tanks, and effluent ponds.[14] I correct the coordinates where necessary, and I consult 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.

I identify 1,526 palm mills as of 2017. I omit mills in Java, where there is little palm oil cultivation; instead, Java primarily houses palm oil refineries and administrative offices. As a validation check, I compare these mill data with official government figures, namely the 2016 Indonesian economic census and 2016 figures from the Malaysia Palm Oil Board.[15] 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.

[14] In the spatial merge, the closest mill within one kilometer is considered a match.

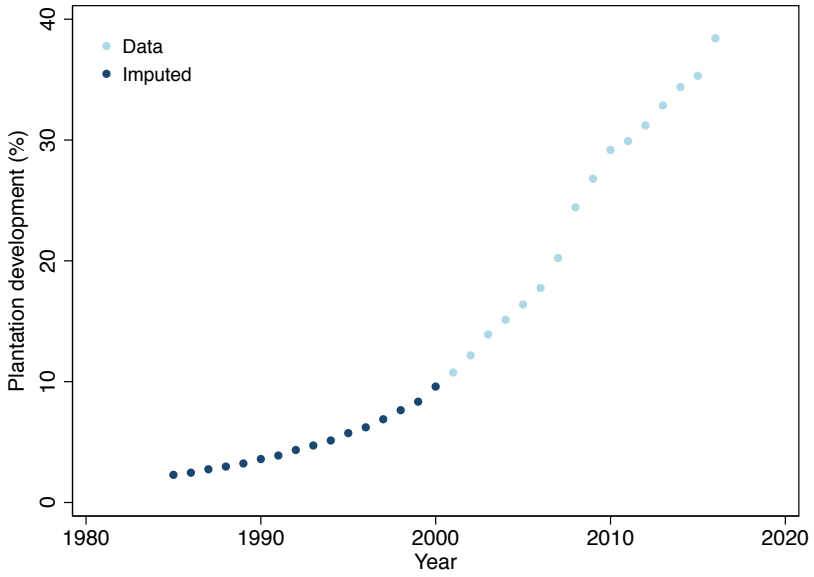
[15] The 2016 Indonesian economic census contains 1,248 palm-oil establishments, of which 1,154 are located outside of Java. Some of these firms are refineries, so I further restrict the dataset to firms involved in extracting crude oil from crops. I am left with 1,070 firms, covering both producers of crude palm oil (1,017 firms) and crude palm kernel oil (53 firms), as indicated by KBLI codes 10431 and 10432, respectively.

Table B5: Xu et al. (2020) plantation vs. Song et al. (2018) tree cover data, 2001-2016

	ΔPlantation$_t$	ΔPlantation$_t$	ΔPlantation$_t$
ΔTree cover$_t$	-0.00314***	-0.00253***	-0.00262***
	(0.000156)	(0.000155)	(0.000153)
ΔTree cover$_{t-1}$	-0.00524***	-0.00441***	-0.00435***
	(0.000192)	(0.000191)	(0.000190)
ΔTree cover$_{t-2}$	-0.00103***	0.000199	0.000408**
	(0.000194)	(0.000193)	(0.000194)
ΔTree cover$_{t-3}$	-0.000672***	6.47e-05	7.27e-05
	(0.000162)	(0.000161)	(0.000160)
Year FE	x	x	x
District FE		x	
Tile FE			x
Observations	9,095,175	9,095,175	9,095,175

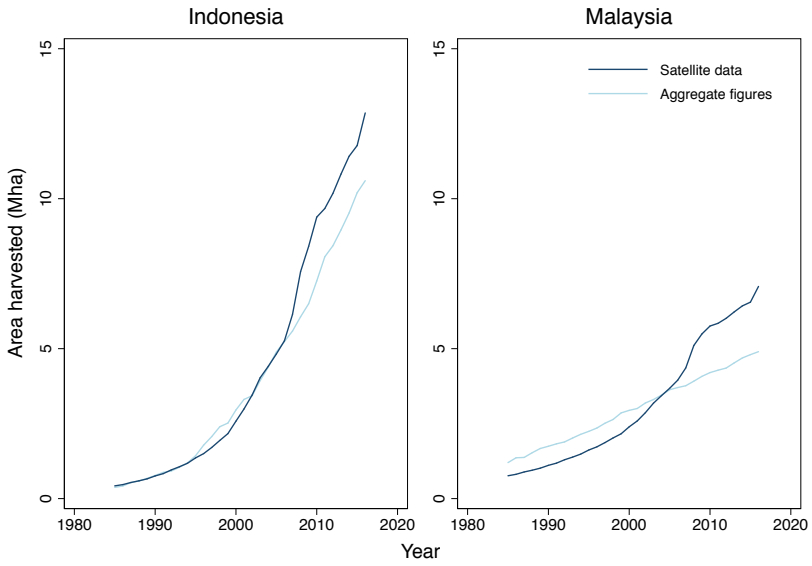
Each observation is a 30-arc-second tile in a given year, and each column is a regression. The dependent variable is from Xu et al. (2020), which measures the ratio of each tile that has been developed into palm oil plantations over time. The independent variables come from Song et al. (2018), which 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: Plantation development conditional on having nonzero development by 2016



Points show the average ratio of 30-arc-second tiles in the study area (Sumatra, Kalimantan, and Malaysia) that have been developed into palm oil plantations over time, conditional on having nonzero development by 2016. About 39% of all tiles have nonzero development by 2016. Points in light blue come directly from Xu et al. (2020), while points in navy blue are imputed using data from Song et al. (2018). These values are based on estimates from the third column of table B5.

Figure B2: Satellite data vs. aggregate figures



Satellite data on palm oil plantation development come from Xu et al. (2020) and Song et al. (2018), and aggregate figures come from the USDA Foreign Agricultural Service. The correlations are 0.9938 for Indonesia and 0.9757 for Malaysia.

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

	Mill data	Government figures
Indonesia	1054	1070
Kalimantan	329	260
Central Sumatra	264	358
North Sumatra	226	237
South Sumatra	206	178
Sulawesi	21	30
Papua	8	7
Malaysia	472	453
Peninsular Malaysia	266	247
Sabah	132	129
Sarawak	74	77
Total	1526	1523

The mill data contains mills built by 2017, and the government figures are from 2016. Mill data come from the Universal Mill List and CIFOR. Indonesia government data come from the 2016 economic census, and Malaysian government data come from the Malaysian Palm Oil Board. Regions are presented in descending order by number of mills. Kalimantan includes the provinces of North, South, East, West, and Central Kalimantan; 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; Sulawesi includes North, South, Southeast, West, and Central Sulawesi, and Gorontalo; Papua includes Papua and West Papua. The states of Sabah and Sarawak comprise East Malaysia, while Peninsular Malaysia includes all other states.

Table B7: Proportion of 2016 plantations impacted by harmonization

	Within 50km	Within 50km, in province	Within 50km, in district
Drop (%)	1.80	2.01	6.35
Delay (%)	5.10	5.29	7.16
Changed (%)	6.90	7.30	13.51

Each column is one method of harmonization, with the second column being the one I adopt in my analysis. The harmonization criterion applied in the first column imposes that all plantation development be within 50 kilometers of an existing mill. The second and third columns further impose that plantation development be associated with mills within the same provinces and districts, respectively. Weighting by the amount of plantation area in each year, the first row shows the percentage of plantation area dropped by the harmonization procedure. The second row shows the percentage delayed to achieve alignment with the mill data. The third row sums the first and second rows.

B.4 Harmonizing the plantation and mill data

Plantation development and mill construction are interdependent: plantations cannot generate revenues without nearby mills. Plantations produce fresh fruit bunches that rapidly deteriorate unless milled into palm oil, which can be sold in world markets. Because this deterioration is increasing in travel distance, nearby mills are preferred to faraway mills, and, by industry standard, plantations without a mill within 50 kilometers are considered infeasible. In this section, I use these rules to construct a correspondence between between the plantation and mill data described previously (since I do not observed these linkages directly). I focus on the period from 1988 to 2016 because the mill data cover 1988 to 2017 and the plantation data cover 1985 to 2016.

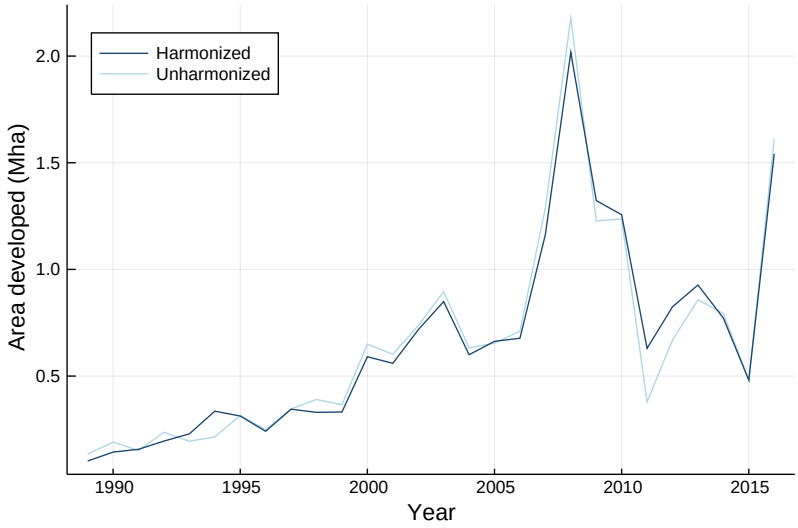
First, I impose that all plantation development have access to a mill within 50 kilometers.[16] In cases where no mill such mill is built by 2016, development is infeasible and I drop it entirely. In cases where such a mill does not exist in the current period but is eventually built, I delay the new development to align it with mill construction. Table B7 shows that less than 2% of plantation development is dropped because no mill is ever built within 50 kilometers, and about 5% is delayed to some degree. I do not drop mills.

Second, I further impose that plantations be linked to mills within the same province (Indonesia) or state (Malaysia).[17] This assumption simplifies computation in defining potential sites in section D.1 because it allows me to define sites separately for each region. Furthermore, there is anecdotal support for plantations' staying within these borders to avoid licensing with multiple regional governments. Table B7 shows that this criterion has little marginal effect on the results of harmonization and therefore does not introduce significant bias. I also experiment with restricting plantations to in-district mills, but this criterion introduces substantially more bias than the in-province one. In line with the low proportion of changes in table B7, figure B3 shows that the harmonized and unharmonized data align well with each other.

[16] I assume new plantations can be linked to new mills. The plantation data record when young palm trees have been established, and the mill data record when mill construction begins. I impose that these events align with each other. Young palm trees do not bear fruit, but proximity to an under-construction mill ensures that an operational mill will be available by the time these young trees reach maturity and begin to bear fruit.

[17] Since they are small and contain no mills, I combine Kuala Lumpur, Labuan, Perlis, and Putrajaya with neighboring provinces (Selangor, Sabah, Kedah, and Selangor, respectively).

Figure B3: Harmonized vs. unharmonized plantation data



The light-blue line plots new development over time in the unharmonized plantation data produced in section B.2. The navy-blue line plots new development in the harmonized plantation data, which impose consistency with the mill data produced in section B.3.

B.5 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. I proceed in the following steps:

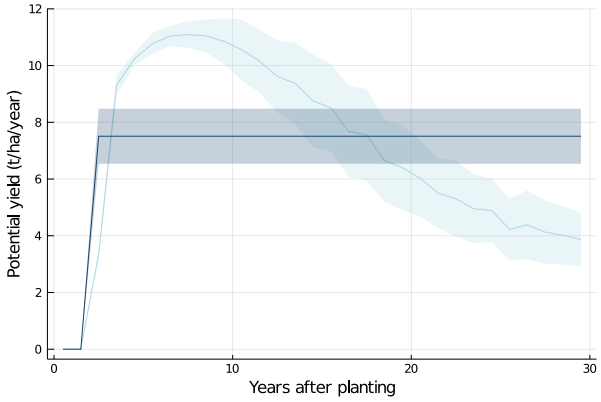
1. Compute potential yields by site from PALMSIM
2. Compile attained yields by province and year from government statistics
3. Combine to produce site-level panel data on attainable yields

First, I compute potential yields by site using the PALMSIM model of Hoffmann et al. (2014), who predict yields under optimal growing conditions by modeling the physiological growth process of the oil palm plant. Relative to other such models, this model requires only a simple set of input variables while still performing well on validation measures. In particular, I use average monthly solar radiation and precipitation from WorldClim, which measures these variables at a resolution of 30 arc-seconds, to compute solar radiation and precipitation by site (where sites are as defined in appendix section D.1). I then run the PALMSIM model for each site to compute how palm oil yields evolve in the 30 years after planting. Figure B4a shows the resulting yield curve, which starts at zero for several years before increasing steeply then declining gradually. Because the data on attained yields distinguish only between "immature" and "mature" palm oil crops, I flatten the yield curve as shown in the figure while holding fixed the average yield over time. The flattened yields at maturity, which vary across sites as shown in figure B4b, are the output I use in subsequent analysis. Note that these data are time-invariant because yields under optimal conditions are an inherent characteristic of the oil palm plant and therefore do not change over time.

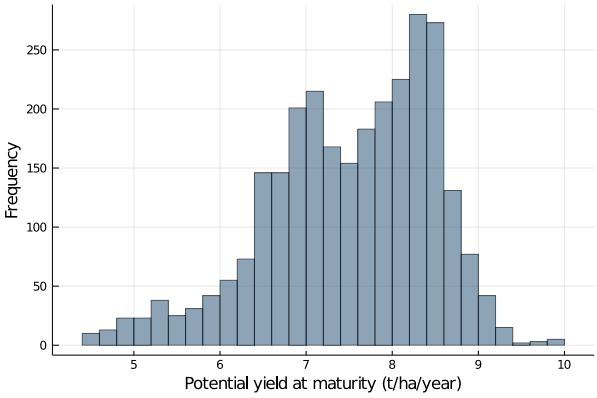
Second, I compile data on attained yields by province and year from government statistics, namely the Indonesian Ministry of Agriculture, the World Bank INDO-DAPOER database (via the Indonesian MoA), and the Malaysian Palm Oil Board. While the Indonesian data are also

Figure B4: Potential palm oil yields

(a) Simulated vs. flattened yield curve



(b) Potential yield at maturity by site



Yield curves are computed from the PALMSIM model (Hoffmann et al. 2014) using two climate inputs: average monthly solar radiation and precipitation from WorldClim. These inputs are measured at the field level; I aggregate inputs by site and run the PALMSIM model at the site level. Sites are those defined in appendix section D.1. On the left, the light blue curve shows the output of the PALMSIM model. The solid line is the average across all sites, and the shaded area shows the standard deviation. The navy blue line represents the flattened yield curve that I use for subsequent analysis, where the flattened curve is restricted to only two yield levels – "immature" (zero-yield) and "mature" – and has the same average over time as the simulated curve. On the right, I show the dispersion of (flattened) mature yields across sites.

available at the district level, I find that the province-level data evolve more stably over time. As well, the Malaysian data are available only at the state level, with Malaysian states analogous to Indonesian provinces. Both sources of data report yields for "mature" oil palm crops, omitting newly planted "immature" crops that do not produce fruit. Figure B5a shows that, on average, these yields are increasing over time as technological progress helps farmers approach the maximum potential yields, although attained yields fall far short of these potential levels in all provinces and years.[18] 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 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}.

$$Y_{it} = (1 - \gamma_{mt})Y_i^p \tag{13}$$

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. Site-year yields Y_{it} aggregate to the observed province-year yields Y_{mt} as

$$\frac{\sum_{i \in \mathcal{I}_m} Y_{it} d_{it}}{\sum_{i \in \mathcal{I}_m} d_{it}} = Y_{mt}\,, \tag{14}$$

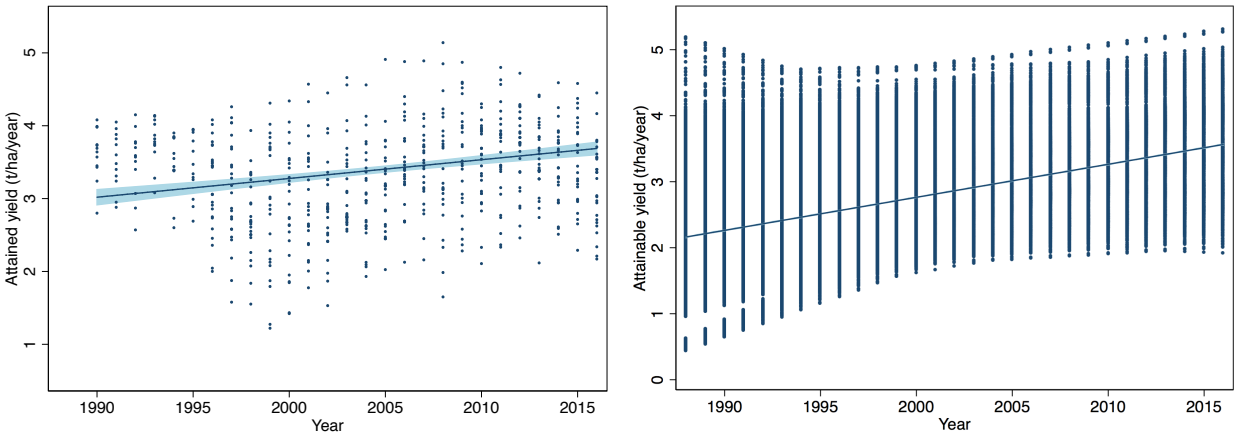
where d_{it} is the amount of land in site i has been developed as of year t. That is, the observed provincial yields are based only on developed lands, and not on the yields of all lands. Combining

[18] Compositional changes in the age mix of palm oil crops can also account for changes in realized yields over time. On one hand, newly planted crops will increase average yields as they reach their peak yields. On the other hand, aging crops will decrease average yields as their yields decline with age. Because these two effects offset each other, I rule out this channel and attribute the observed yield increases to technological progress.

Figure B5: Attained and attainable palm oil yields

(a) Recorded yields by province and year

(b) Imputed yields by site and 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. The fitted line shows a common time trend, accounting for province/state fixed effects. On the right, each observation is the annual attainable yield for a given site as imputed by combining potential variation across sites from PALMSIM with attained levels and time trends across provinces from government statistics. The fitted line shows a common time trend, accounting for site fixed effects. For both, shaded bands show 95% confidence intervals.

these relationships, I can solve for yield gaps γ_{mt} to obtain

$$\gamma_{mt} = 1 - Y_{mt}\left(\frac{\sum_{i\in\mathcal{I}_m} Y_i^p d_{it}}{\sum_{i\in\mathcal{I}_m} d_{it}}\right)^{-1}. \tag{15}$$

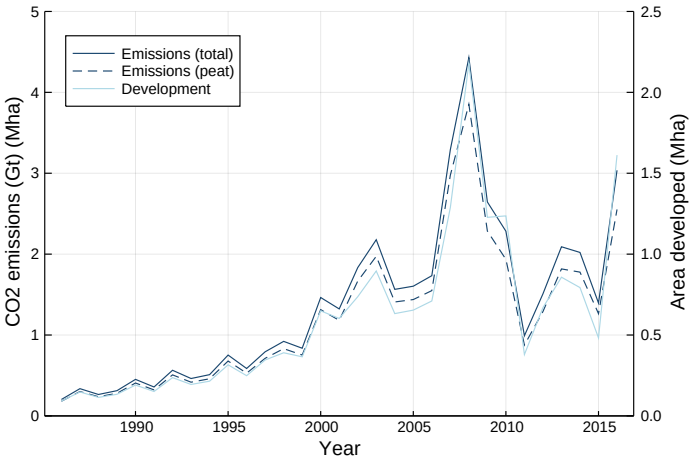
To isolate the underlying levels and trends of these yield gaps, I estimate the specification

$$\gamma_{mt} = \alpha_m + \beta_m t + \varepsilon_{mt}\,, \tag{16}$$

and I use the fitted values to estimate site-year yields Y_{it} using equation 13. Intuitively, I combine potential variation across sites from PALMSIM with attained levels and time trends across provinces from government statistics to estimate attainable yields by site and year. I do not have on attained yields for Malaysia before 1990 and for Indonesia before 1996, and so in both cases I extrapolate yield gaps back to 1986.

Figure B5b shows the estimated attainable yields, which maintain the uptrend observed in figure B5a while incorporating the site-level dispersion shown in figure B4b. The slope of the time trend is lower in figure B5a because these yields are among sites selected for development: if the best sites are developed first, then the improvement from technological progress over time is partially offset by the fact that future sites are negatively selected. In contrast, the attainable yields in figure B5b are for the full population of sites in every year. Furthermore, note that the fanning to the left in figure B5b reflects underlying province-specific trends that would also appear in figure B5a if I were to extrapolate the observed province measurements values back to 1986.

Figure B6: Plantation development vs. emissions



The light blue line shows changes in the extent of palm oil plantations, as measured using data from Xu et al. (2020) and Song et al. (2018). The navy blue lines show emissions, as measured using data on aboveground biomass from Zarin et al. (2016) and data on peatlands from Gumbricht et al. (2017).

B.6 Carbon stocks

I compute carbon stocks over space using two datasets: Zarin et al. (2016) measures aboveground biomass, capturing carbon stored in trees, at a resolution of 30m, and Gumbricht et al. (2017) measures belowground biomass in the form of peat depth at a resolution of 231m. I aggregate both datasets to a resolution of 30 arc-seconds. 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/m^3 peat in Warren et al. (2017). The implied carbon emissions from plantation development, which destroys both above- and belowground biomass, sums the above- and belowground quantities. I can convert carbon to carbon dioxide emissions using a molecular-weight conversion factor of 3.67.

I treat carbon stocks as predetermined, but one concern is that they are measured during the study period – not before. 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 mitigated by the way in which Gumbricht et al. (2017) construct the data. The authors rely primarily on precipitation and topography – predetermined features – in order to identify areas where precipitation exceeds evapotranspiration, and where water is likely to pool. Once wetlands are identified in this way, the authors use MODIS satellite imagery from 2011 to distinguish between different kinds of wetlands. Indeed, figure B6 shows that the relationship between plantation development and the resulting emissions is consistent over time. If peatlands destroyed by plantation development before 2011 were not captured in the peatland data, then estimated peatland emissions should be much smaller for plantation development before 2011.

C Appendix: Demand

Table C1: Rainfall as price instruments, alternative specifications

	Including temperature	Asymmetric effects
Rain, deviation from optimal (100 mm)	0.111***	
	(0.0323)	
Temperature, deviation from optimal (°C)	0.131	
	(0.123)	
Rain, below optimal (100 mm)		0.177***
		(0.0626)
Rain, above optimal (100 mm)		0.129***
		(0.0320)
Oil FE	x	x
Year-oil trend	x	x
Observations	74	74
F-statistic	6.599	8.196

The outcome variable is log price of a given oil product in a given year. In the first column, rainfall is constructed at the oil-product level by aggregating rainfall across producing regions, weighting by total production over the study period. For a given region, rainfall is measured as the total absolute deviation from optimal monthly rainfall levels over the course of the growing season. Temperature is calculated similarly, except that I assess deviations from optimal conditions at a daily frequency. In the second column, I separate positive and negative deviations from optimal monthly rainfall levels. Data are annual and cover 1980 to 2016. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C2: Lower-level demand elasticities

		Estimates		Standard errors	
		Palm	Other	Palm	Other
EU	Palm	-1.09***	0.48	(0.20)	(0.47)
	Other	0.02	-1.08***	(0.03)	(0.08)
China/India	Palm	-0.82***	0.14	(0.31)	(0.45)
	Other	-0.04	-1.03***	(0.07)	(0.10)
Other importers	Palm	-1.05***	0.01	(0.19)	(0.33)
	Other	0.01	-1.00***	(0.04)	(0.08)
Indonesia	Palm	-0.86***	-0.15	(0.08)	(0.36)
	Other	-0.41	-0.55	(0.26)	(1.09)
Malaysia	Palm	-0.93***	-0.15	(0.05)	(0.12)
	Other	-1.24	1.56	(0.80)	(2.08)

Uncompensated price elasticities are computed from estimated demand parameters using equation 3, omitting the final category-consumption term. Palm oil aggregates palm and palm kernel oil, while "other" oils include coconut, olive, rapeseed, soybean, and sunflower oil. I evaluate expenditure shares, prices, and the time trend at their averages over the study period. I instrument for prices with foreign rainfall shocks. Data are annual and cover 1980 to 2016. Standard errors are computed with the delta method, and I apply a Prais-Winsten transformation to account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C3: Upper-level demand elasticities

	Estimate	SE	Obs
European Union	-0.0547	(0.129)	37
China/India	-0.347	(0.720)	37
Other importers	-0.0583**	(0.0255)	37
Indonesia	-0.0638	(0.169)	37
Malaysia	-0.287*	(0.174)	37

Each row is a regression showing the effects of log oil prices on log oil consumption. I control for non-oil consumption prices, log GDP, and a time trend. Oil prices are measured as a translog price index based on lower-level demand estimates, and I instrument for them with foreign rainfall shocks. Data are annual and cover 1980 to 2016. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C4: Rainfall as price instruments by consumer market

	Estimate	SE	Obs	F-statistic
European Union	0.113***	(0.0328)	74	11.81
China/India	0.112***	(0.0324)	74	11.97
Other importers	0.114***	(0.0323)	74	12.45
Indonesia	0.0638***	(0.0227)	74	7.864
Malaysia	0.0804***	(0.0257)	74	9.817

Each row is one regression showing the effects of rainfall on log prices, controlling for oil-specific time trends. This table replicates column 2 of table 2 excluding rainfall shocks within each consumer market. Rainfall is constructed at the oil-product level by aggregating rainfall across producing regions, weighting by total production over the study period. For a given region, rainfall is measured as the total absolute deviation from optimal monthly rainfall levels over the course of the growing season. When aggregating across regions, the first row omits regions in EU countries, the second regions in China and India, the third regions in other importers, the fourth regions in Indonesia, and the fifth regions in Malaysia. Data are annual and cover 1980 to 2016. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C5: Effects of rainfall on GDP and oil expenditures

	Log GDP		Log final consumption		Log final HH consumption		
	Estimate	SE	Estimate	SE	Estimate	SE	Obs
China/India	0.0542	(0.0359)	0.0458	(0.0360)	0.0442	(0.0370)	74
European Union	-0.00531	(0.0212)	-0.00334	(0.0214)	-0.00482	(0.0210)	74
Indonesia	0.00889	(0.0280)	0.00565	(0.0257)	0.00607	(0.0240)	74
Malaysia	0.0275	(0.0168)	0.0350	(0.0241)	0.0325	(0.0232)	74
Other importers	0.00168	(0.0105)	0.000839	(0.00968)	-0.000115	(0.00890)	74

Each row is three separate regressions showing the effects of foreign rainfall shocks on CPI-adjusted log GDP, final consumption expenditures, and final household consumption expenditures, respectively, controlling for oil-specific time trends. Data are annual and cover 1980 to 2016. Newey-West standard errors account for serial correlation. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D Appendix: Supply

This section contains details on the defining of potential sites, the estimation of both intensive- and extensive-margin supply models, and the computation of supply elasticities.

D.1 Defining sites

The model analyzes investment decisions within sites, each managed by long-lived landowners that operate as independent firms. The benefit of this approach is tractability: the alternative is modeling the spatial entry problem and grappling with the resulting curse of dimensionality. On the other hand, a drawback is that I must define the boundaries of sites. Below, I describe the procedure I use to aggregate the field-level data, which are measured at a resolution of 30 arc-seconds, into potential palm oil sites. I conduct this procedure separately by province (Indonesia) or state (Malaysia). The assumed separability across provinces facilitates computation while introducing relative little bias (table B7), and it is anecdotally consistent with plantations' remaining within regional borders to avoid licensing with multiple regional governments. The result is a set of contiguous sites that is consistent with the mills observed in the data during the study period.

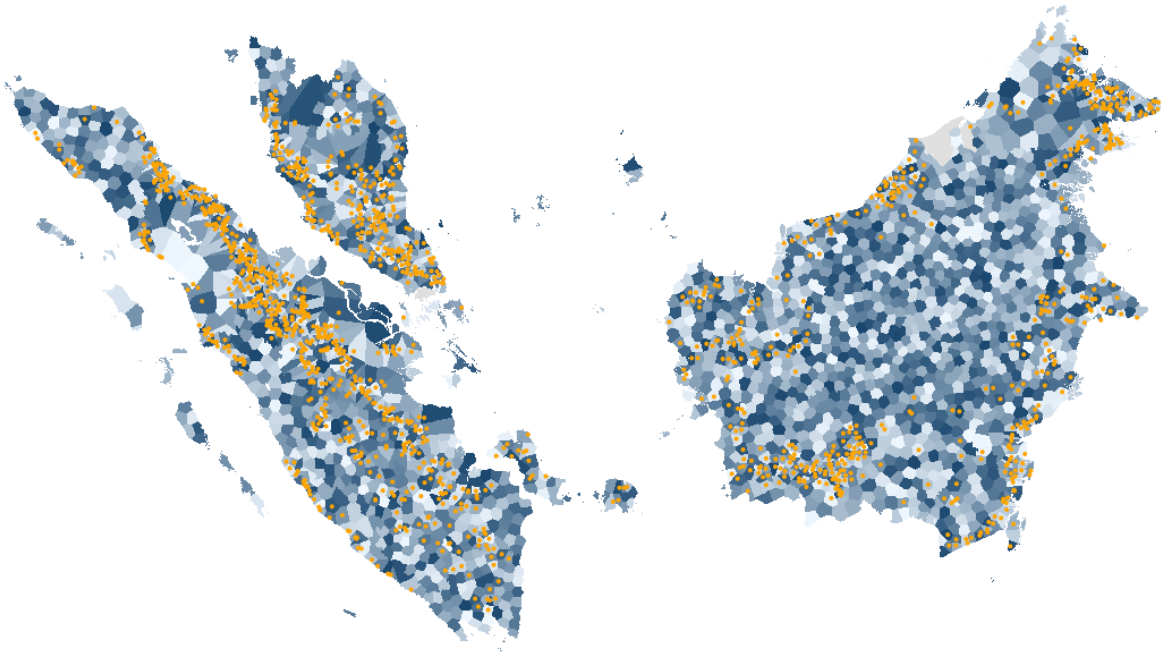
First, I compute the maximum number of sites k for each province. I do so by floor dividing total land area by 500 km^2 (585 30-arc-second tiles); if the actual number of mills by 2016 is higher, then I use the actual number instead. I arrive at this average site size of 500 km^2 in several ways. First, I consider the density of mills in regions where the palm oil industry is most developed. Computing the density of observed mills by province, I find that at the 75th percentile for density there is one mill per 453 km^2. Second, I consider provinces with existing mills but no recent mill construction – another way of identifying regions where development has plateaued. The median province with no mill construction in the last five years of the study period has one site per 500 km^2. Both the first and second methods imagine provinces' reaching the site density of the most developed regions, which both methods identify as consisting primarily of a set of Malaysian provinces. A third alternative is to consider the radius of plantations that a single mill can serve and to size sites accordingly. Plantations must be within 50 km of a mill because of production constraints specific to palm oil, but a 50-km radius implies a site size of 7,850 km^2, yielding far fewer sites than observed (171 vs. 1,492). This discrepancy suggests that the 50-km constraint is often not binding in practice. Instead, I consider the plantation-mill distances observed in the data, the upper extreme (75th percentile) of which is 11.7 km, implying a site size of 503 km^2. All three methods yield average site sizes in the neighborhood of 500 km^2.

Second, given the maximum number of sites k, I define site by k-means clustering. I cluster on geographic coordinates in order to produce contiguous sites, and I impose that each site contain no more than one of the mills observed in the data by 2016. I do so by applying a version of the constrained k-means clustering algorithm described in Wagstaff et al. (2001).

1. Choose initial cluster centers $C_1, C_2, \ldots, C_k$.
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.

Step (2) produces clusters that contain no more than one mill per cluster. Although convergence is not guaranteed, the algorithm converges in my case and yields 2,805 sites. I use multiple starts because convergence is to local optima. Figure D1 plots the potential sites.

Figure D1: Potential sites



Blue shading indicates different potential sites. Gray shading shows omitted regions, including sites with zero plantation development observed during the study period. Oranges dots are palm oil mills observed by 2016. Excluding omitted regions, I obtain 2,805 potential sites.

I do not use data on observed plantations in clustering, but I nonetheless obtain clusters consistent with these data. On average, 89% of observed plantations have access to an on-site mill. For the remaining 11%, I delay development if an on-site mill is eventually constructed, or I drop it if not. The alternative is to incorporate observed plantations in clustering, for example with a penalty for assigning plantations to sites without mills. The drawback, however, is that the resulting clusters may not be contiguous when clustering is not solely on geographic coordinates. Since 11% is relatively low, I do not pursue this alternative in the baseline analysis.

D.2 Estimating the intensive-margin model (plantation development)

Expectational errors η_{it}

$$a_{it} - \beta \mathbb{E}_{it}[a_{it+1}] = \frac{\beta}{\delta}\mathbb{E}_{it}[Y_{it+1}P_{t+1}] - \frac{1-\beta}{\delta}x_i\gamma - \frac{1-\beta}{\delta}\kappa_i - \frac{1}{\delta}\alpha_r[t - \beta(t+1)] - \frac{1}{\delta}\varepsilon_{it}$$

The Euler equation forms a telescoping series, which I iterate to obtain

$$a_{it} = \sum_{t'=1}^{\infty} \frac{\beta^{t'}}{\delta}\mathbb{E}_{it}[Y_{it+t'}P_{t+t'}] - \frac{1}{\delta}x_i\gamma - \frac{1}{\delta}\kappa_i - \frac{1}{\delta}\alpha_r\left(\frac{t - \beta(t+1) + \beta}{1-\beta}\right) - \frac{1}{\delta}\varepsilon_{it}\,,$$

noting that $\mathbb{E}_{it}[\varepsilon_{it+t'}] = 0$ for $t' > 1$ because cost shocks are mean-zero and IID. It follows that

$$\beta\mathbb{E}_{it}[a_{it+1}] - \beta a_{it+1} = \sum_{t'=2}^{\infty} \frac{\beta^{t'}}{\delta}\left(\mathbb{E}_{it}[Y_{it+t'}P_{t+t'}] - \mathbb{E}_{it+1}[Y_{it+t'}P_{t+t'}]\right) + \frac{\beta}{\delta}\varepsilon_{it+1}\,,$$

which by definition of the expectational errors implies

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

D.3 Estimating the extensive-margin model (mill construction)

Evaluating sequence $(0, 1, a'_{it+1})$

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

Proof. Comparing choice-specific conditional value functions $v^e(0; \boldsymbol{w}_{it})$ and $v^e(0, 1; \boldsymbol{w}_{it})$,

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

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

Proof. Comparing choice-specific conditional value functions $v^e(1; \boldsymbol{w}_{it})$ and $v^e(1, a_{it}; \boldsymbol{w}_{it})$,

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

where $a^*_{it} \equiv a^*_{it}(0; \boldsymbol{w}_{it}, \varepsilon_{it})$. The first equality is definitional. The second equality applies that costs are quadratic and revenues linear. The third equality applies the first order condition that holds at a^*_{it} and the linearity of revenues. 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 $v^e(1; \boldsymbol{w}_{it}) \geq v^e(1, a_{it}; \boldsymbol{w}_{it})$.

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

Proof. Comparing choice-specific conditional value functions $v^e(0; \boldsymbol{w}^e_{it})$ and $v^e(0, 1, a'_{it+1}; \boldsymbol{w}^e_{it})$,

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

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

E Appendix: Counterfactuals

This section describes how I solve the model and quantify carbon emissions.

E.1 Solving the model

I impose additional assumptions on expectations over the evolution of the state variables, and I solve by backward induction.

Expectations over aggregate states d_t and s_t

Expectations over the evolution of demand d_t and supply s_t together determine the expected path of prices $P(s_t, d_t)$. I make explicit assumptions about expectations for demand d_t, which I describe below. Supply s_t is determined endogenously as the result of an entry game in which beliefs are correct in equilibrium.

I model the non-stationary evolution of demand d_t with an ARIMA process, and I assume expectations for all firms are given by this process. Table E1 evaluates log likelihoods over a range of ARIMA specifications and finds that an ARIMA$(2, 1, 2)$ process produces the best fit to the data. In this specification, changes $d_t - d_{t-1}$ in demand follow an ARMA(2,2) process.

$$d_t - d_{t-1} = c + \upsilon_t + \sum_{t'=1}^{2} \left(\varphi_{t'}(d_{t-t'} - d_{t-t'-1}) - \theta_{t'} \upsilon_{t-t'} \right)$$

Since the demand curve is specified in logs, this ARIMA process can sometimes predict infinite exponential growth in demand. Such unbounded growth leads to unrealistically stark predictions: exponentially rising demand (at a rate that dominates discounting β) implies infinite returns to development and therefore immediate development of all undeveloped lands. Thus, I shrink the ARIMA estimates toward a sigmoid function fit to observed demand. Expectations are therefore

$$\mathbb{E}_{it}[d_{t+t'}] = \left(\frac{\widehat{V}^{\text{SIG}}}{\widehat{V}^{\text{ARIMA}}_{t+t'} + \widehat{V}^{\text{SIG}}} \right) \widehat{d}^{\text{ARIMA}}_{t+t'} + \left(\frac{\widehat{V}^{\text{ARIMA}}_{t+t'}}{\widehat{V}^{\text{ARIMA}}_{t+t'} + \widehat{V}^{\text{SIG}}} \right) \widehat{d}^{\text{SIG}}_{t+t'} \quad \text{for } t' \geq 1\,, \tag{17}$$

where I weight by inverse variances, with the variance of the sigmoid predictions given by the mean squared error. The ARIMA predictions have increasing variance for expectations taken farther into the future, implying greater reliance on the fitted sigmoid function in these periods. Figure E1 plots both ARIMA and shrunk demand expectations. Indeed, shrinking toward the sigmoid function helps in bounding demand expectations.

Expectations over site-specific states Y_{it}, x_i, ε_{it}, and ε^e_{it}

I assume that yields Y_{it} evolve at a constant and exogenous rate per year. Thus, no expectational error arises from changes in yields. There is no need to define expectations over cost factors x_i because they are constant. I assume that while firms know current-period cost shocks ε_{it} and ε^e_{it}, they only know the distribution of future shocks.

I obtain estimates of intensive-margin cost shocks ε_{it} from the residuals of equation 8. The complication is that these residuals combine cost shocks and expectational errors.

$$\upsilon_{it} = -\frac{1}{\delta}\varepsilon_{it} + \eta_{it}$$

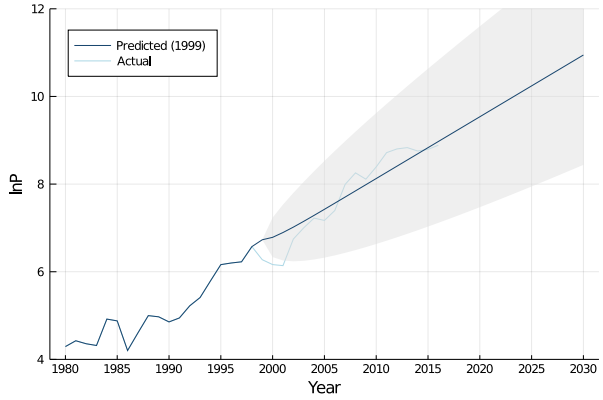
Table E1: ARIMA(p, d, q) log likelihoods for demand d_t

		ARMA(p, q)		
		(0,0)	(1,1)	(2,2)
Differencing (d)	0	-69.17	-6.91	-6.37
	1	-2.41	-0.53	1.44
	2	-14.82	-4.53	-0.37

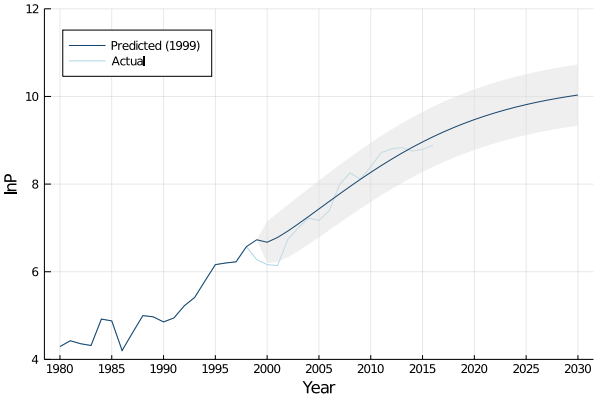
An ARIMA process with $d = 0$, the random variable is itself modeled as an ARMA process. For $d = 1$ it is the difference $x_t - x_{t-1}$, and for $d = 2$ it is the change in differences $(x_t - x_{t-1}) - (x_{t-1} - x_{t-2})$. I take $(p, d, q) = (2, 1, 2)$, which has the highest log likelihood, as my baseline specification.

Figure E1: Demand expectations $\mathbb{E}_{it}[d_{t+s}]$

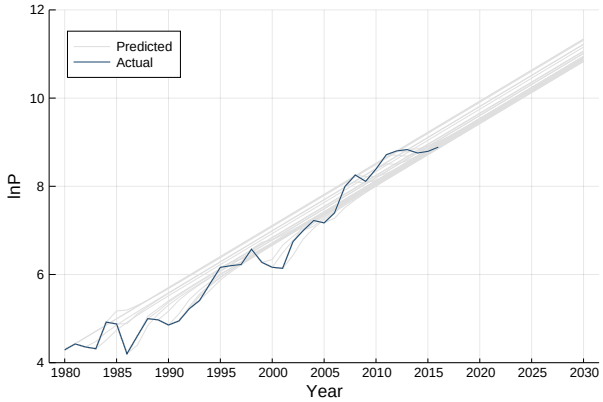
(a) ARIMA predictions, 1999



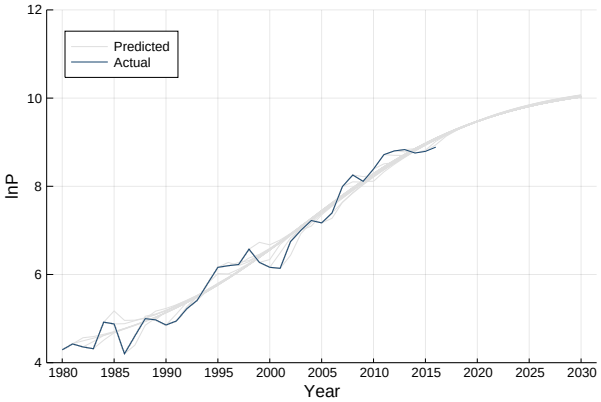
(b) Shrunk predictions, 1999



(c) ARIMA predictions, all years



(d) Shrunk predictions, all years



All figures show expectations for the evolution of demand state d_t. I estimate these demand states in section 5.1, and I plot the realized values as "actual." These realized values coincide with figure 6c. The top row shows predictions and the 95% confidence band from the perspective of a single year, while the bottom row shows such predictions for all years. The left column shows predictions arising from an ARIMA(2,1,2) process that I fit on observed values preceding each prediction year. This specification has the highest log likelihood among those tested in table E1. The right column shows the results of shrinking the ARIMA predictions toward a sigmoid function fit to realized values.

Substituting the expression for expectational errors and applying the above assumptions on expectations, I obtain

$$\varepsilon_{it} - \beta\varepsilon_{it+1} = -\delta\upsilon_{it} + \sum_{t'=1}^{\infty} \beta^{t'} Y_{it+t'} \Big(\mathbb{E}_t[P_{t+t'}] - \mathbb{E}_{t+1}[P_{t+t'}] \Big).$$

Thus, I can estimate cost shocks as a function of residuals υ_{it} and price expectations. The demand expectations of equation 17 translate into price expectations as a function of supply elasticities. Figure E1 shows that expectational errors for demand are relatively small in each period, so I approximate price expectations with the partial-equilibrium supply elasticities of table 4.

I do not obtain estimates of extensive-margin cost shocks ε^e_{kt}. Instead, counterfactuals evaluate the ex-ante value function and yield predicted probabilities of extensive-margin investment.

Backward induction from steady state

I solve the model by backward inducting from the steady state – period S – at which point all feasible lands have been developed. After period S, there is no further entry, but firms continue to generate revenues over the infinite horizon based on past entry. The existence of such a period is asymptotically guaranteed in my model: the total amount of development is non-decreasing given no exit, there are new cost shocks in each period, and there is a finite amount of land that can be developed. The challenge is that it may take many years for every hectare of available land to be developed.

I address this computation burden in two ways. First, I solve each subproblem using an iterative algorithm that uses a fixed look-ahead horizon instead of always looking ahead to the end of the game tree. Given initial state of development s_1, I backward induct from period S as follows.

1. Initialize the algorithm by solving for a_1 given s_1 assuming no further entry after period 1, then for a_2 given $s_2(a_1)$ assuming no further entry after period 2, and so on until a_S. With a_S and $s_S(a_{S-1})$, compute s_{S+1}. Note that $s_S(a_{S-1})$ is shorthand for $s_S(a_{S-1}, a_{S-2}, \ldots, a_1, s_1)$.
2. Taking s_{S+1} as fixed, work backward from period S. First, solve for a_{S-1} given s_{S-1} as a starting state and $\{s_S(a_{S-1}), s_{S+1}\}$ as the future states (with $s_{S+t'} = s_{S+1}$ for all $t' > 1$ given no future entry). Revise s_S given the previous solution to a_{S-1}. Second, solve for a_{S-2} given s_{S-2} as a starting state and $\{s_{S-1}(a_{S-2}), s_S, s_{S+1}\}$ as the future states. Revise s_{S-1} given the previous solution to a_{S-2}. Continue until a_1, noting that all states get revised except for initial state s_1, which must be taken as given.
3. To restart the chain of revisions, solve for a_S given s_S as the starting state and $s_{S+1}(a_S)$ as the future state given no further entry.
4. Repeat steps 2 and 3 until convergence in $\{a_1, a_2, \ldots, a_S\}$.

This algorithm breaks the usual curse of dimensionality in which the state space grows exponentially in the length of the look-ahead window.

Second, I approximate period S by choosing an arbitrary period $T < S$ and solving as if it were the steady state. In setting an earlier period T, computation is faster because the backward induction window is shorter, but there is more bias in ignoring post-T entry because there are more periods after T. My solution is to resolve taking periods $T+1$, $T+2$, and so on as the steady state until the solutions converge. Intuitively, entry today becomes less appealing when competitors have a longer window of opportunity to enter, but discounting means a diminishing marginal impact of extending this window.

Defining notation, world supply and entry in period t are functions of previous and new development, respectively.

$$s_t = \sum_i Y_{it} s_{it}\,, \quad a_t = \sum_i \Big(s^e_{it} a_{it} + (1 - s^e_{it}) p^e_{it} a_{it} \Big), \tag{18}$$

where for sites without mills in period t ($s^e_{it} = 0$), new development depends on both extensive-margin probability p^e_{it} of mill construction and intensive-margin choice a_{it} of plantation development. "Entry" involves plantation development in my context, so I refer to entry and development interchangeably. Entry determines future supply

$$s_{t+1} = s_t + a_t\,,$$

and therefore future world prices

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

To proceed, consider period T and suppose there is no further entry after this period. For sites with a mill in period T ($s^e_{iT} = 1$), the first order condition for a_{iT} determines development.

$$a_{iT} = \frac{1}{\delta} \sum_{t'=1}^{\infty} \beta^{t'} \mathbb{E}_{iT} \Big[Y_{iT+t'} P(s_{T+1}, d_{T+t'}, \tau_{T+t'}) - x_i\gamma - \kappa_m - \alpha_m(T + t') - \varepsilon_{iT+t'} \Big], \tag{19}$$

subject to constraint $0 \leq a_{iT} \leq \bar{s}_i - s_{iT}$. For sites without a mill in period T ($s^e_{iT} = 0$), development also depends on mill construction, which occurs with probability

$$p^e_{iT} = \frac{\exp\big(- x_i\gamma^e - \kappa^e_m - \alpha^e_m T + \mathbb{E}^e_{iT}[V(0; \boldsymbol{w}_{iT}, \varepsilon_{iT})]\big)}{1 + \exp\big(- x_i\gamma^e - \kappa^e_m - \alpha^e_m T + \mathbb{E}^e_{iT}[V(0; \boldsymbol{w}_{iT}, \varepsilon_{iT})]\big)}\,, \tag{20}$$

where the one in the denominator arises from $v^e(0; \boldsymbol{w}_{iT}) = 0$ since there is no further entry after period T (for an outside option normalized to zero).[19] In both cases, entry depends on world prices, which in turn depend on world supply.

The result is an entry game in which the returns to entry for a given firm depends on how many other firms enter. Intuitively, developing a given site has low returns when other sites develop extensively because high supply means low prices. In equilibrium, each firm's entry decision must be consistent with total entry. 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. I solve by selecting an arbitrary level of total development a_T, computing the site-specific development choices by equations 19 and 20, and calculating the implied total a'_T by equation 18. If the implied total is higher (lower) than the initial total, then for the next iteration I start with a higher (lower) initial total. In this way, I obtain site-specific period-T development $\boldsymbol{a}_T = \{a_{iT}, a^e_{iT}\}$ as a function of previous development $\boldsymbol{s}_T = \{s_{iT}, s^e_{iT}\}$.

[19] To determine the probability of extensive-margin entry, I compute intensive-margin profits assuming $\mathbb{E}^e_{iT}[\varepsilon_{iT}] = 0$ because I assume that firms make extensive-margin decisions before observing intensive-margin shocks. When computing actual intensive-margin entry, however, I use realized intensive-margin shocks ε_{it}. Furthermore, since intensive-margin profits $\mathbb{E}^e_{iT}[V(0; \boldsymbol{w}_{iT})]$ are not linear in ε_{iT} (even though choices a_{iT} are), I cannot simply apply $\mathbb{E}^e_{iT}[\varepsilon_{iT}] = 0$ and must instead compute expected intensive-margin profits based on the distribution of ε_{iT}, which I assume firms know.

The problem is computationally fast to solve. First, prices are monotonically decreasing in total entry a_T, so the solution is unique and standard root-finding algorithms work well. Second, I can iterate on total development a_T instead of site-specific development $\boldsymbol{a}_T$ because world prices are influenced only by total supply and not the spatial distribution of supply. This simplification rules out spatial competition concerns, which would otherwise generate a severe curse of dimensionality by requiring iteration over the I-dimensional space $\boldsymbol{a}_T$. Third, as in Hopenhayn (1992), I invoke that firms are small enough to approximate a continuum: by the law of large numbers, the implied total is simply the expected value resulting from extensive-margin entry probabilities p^e_{iT}. By contrast, with a small number of large firms, the extensive-margin entry probabilities induce a binomial distribution over total entry. In dealing with a scalar instead of a distribution, I avoid the computational burden of computing outcomes over each point of the distribution.

Working backward, consider development $\boldsymbol{a}_{T-1}$ in period $T-1$. Taking previous development $\boldsymbol{s}_{T-1}$ as given, I solve for new development $\boldsymbol{a}_{T-1}$ as follows.

1. I make an initial guess for total new development a_{T-1}.
2. I divide this total new development a_{T-1} into site-specific new development $\boldsymbol{a}_{T-1}$. Since the first order condition is monotonic in prices, only one such division exists.
3. With $\boldsymbol{s}_{T-1}$ and $\boldsymbol{a}_{T-1}$, I obtain site-specific $\boldsymbol{s}_T$ and therefore total s_T.
4. Given $\boldsymbol{s}_T$, I solve the subproblem for $\boldsymbol{a}_T$ using the solution algorithm described above for entry in period T, after which there is no further entry. With $\boldsymbol{s}_T$ and $\boldsymbol{a}_T$, I obtain site-specific $\boldsymbol{s}_{T+1}$ and therefore total s_{T+1}.
5. Given totals s_T and s_{T+1}, I compute site-specific $\boldsymbol{a}_{T-1}$ with analogues of equations 19 and 20.[20]
6. Finally, I check if site-specific new development $\boldsymbol{a}_{T-1}$ sums to the guess for total new development a_{T-1}. If so, then $\boldsymbol{a}_{T-1}$ is the solution. If not, then I repeat the above steps with a different guess for a_{T-1}.

Solving for entry in period $T-2$ and in earlier periods follows similarly, where I can solve the subproblems in step four by recursively applying the same algorithm.

E.2 Quantifying carbon emissions

I account for substitution to paper pulp (*acacia*) plantations by estimating the observed relationship between paper pulp and palm oil plantation development. I estimate this relationship using data on paper pulp plantation development as of 2016 on the island of Borneo (Gaveau et al. 2019), as mapped in figure E2.

$$\text{acacia}_i = \beta_0 + \beta_1 \text{palm}_i + \beta_2 \text{mill_distance}_i + \alpha_m + \varepsilon_i \,, \tag{21}$$

[20] For intensive-margin entry in equation 19, the analogue in period $T-1$ is similar except that prices depend on s_T in period T and s_{T+1} thereafter. A firm's expected development a_{iT} in period T does not enter. For extensive-margin entry probabilities in equation 20, the expression is simplified in period T because $v^e(0; \boldsymbol{w}_{iT}) = 0$ given no further entry. In earlier periods t, $v^e(0; \boldsymbol{w}_{it})$ is instead given by the logit log-sum formula

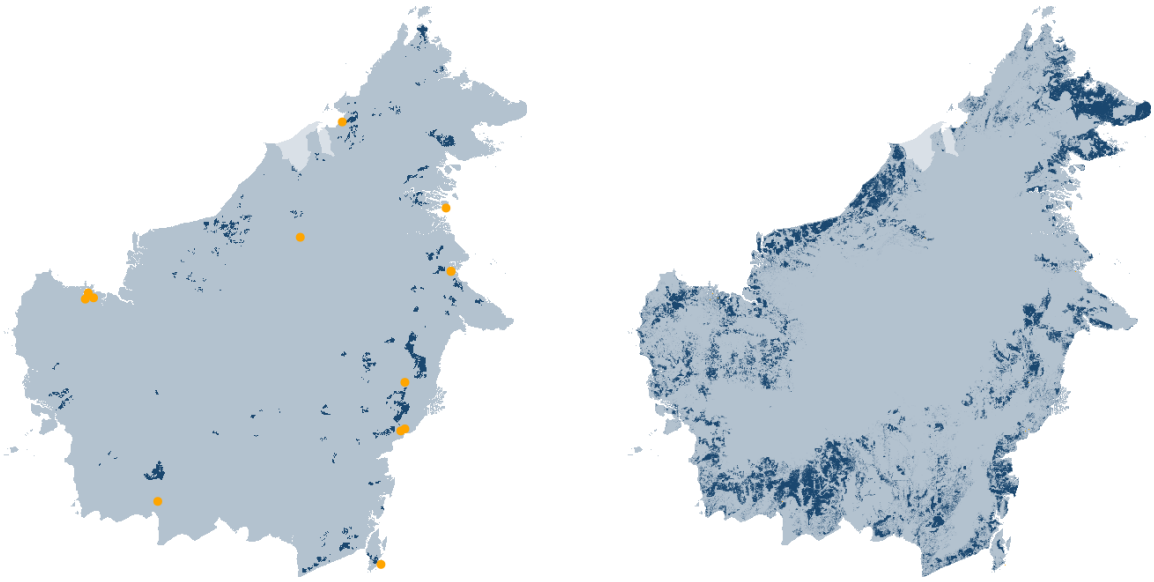
$$\begin{aligned} v^e(0; \boldsymbol{w}_{it}) &= \beta \mathbb{E}^e_{it}[V^e(\boldsymbol{w}_{it+1})] \\ &= \ln(e^{\mathbb{E}^e_{it}[v^e(1;\boldsymbol{w}_{t+1})]} + (e^{\mathbb{E}^e_{it}[v^e(1;\boldsymbol{w}_{t+2})]} + \cdots + (e^{\mathbb{E}^e_{it}[v^e(1;\boldsymbol{w}_T)]})^\beta)^\beta)^\beta . \end{aligned}$$

I account explicitly for the distribution of future intensive-margin cost shocks ε_{it+s}, which do not fall out because intensive-margin profits $V(0; \boldsymbol{w}_{it}, \varepsilon_{it})$ are not linear in ε_{it}, although development a_{it} is.

for sites i and regions m (provinces for Indonesia and states for Malaysia), and where I control for distance to the nearest paper pulp mill. Table E2 shows that lower levels of palm development are indeed associated with higher levels of paper pulp development, although the magnitude of the relationship does not seem to be large.

Figure E2: Plantation development, 2016

(a) Paper pulp **(b)** Palm oil



The figures map plantations as of 2016 for the island of Borneo, which is shared by Indonesia, Malaysia, and Brunei. The shaded out region is Brunei. On the left, data on paper pulp plantations come from Gaveau et al. (2019), and orange dots mark paper pulp mill locations based on information from the Indonesian Pulp and Paper Association. On the right, data on palm oil plantations come from Xu et al. (2020).

Table E2: Paper pulp vs. palm oil plantation development

Palm plantation development (%)	-0.0195***	-0.0235***
	(0.00610)	(0.00734)
Log paper pulp mill distance (km)	-0.0265***	-0.0210***
	(0.00447)	(0.00452)
Province FE		x
Observations	1,060	1,060

Each column is one cross-sectional regression using 2016 data, and each observation is a site. The sample is restricted to the island of Borneo, where data on paper pulp plantations are available (Gaveau et al. 2019). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.