

# Essays in Environmental and Development Economics

Allan Hsiao

MPhil, University of Oxford (2015)  
AB, Harvard College (2013)

Submitted to the Department of Economics  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy in Economics

MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
June 2021  
© 2021 Allan Hsiao. All rights reserved.

The author hereby grants to MIT permission to reproduce and to  
distribute publicly paper and electronic copies of this thesis document  
in whole or in part in any medium now known or hereafter created.

Author .....  
 Department of Economics  
 May 14, 2021

Certified by .....  
 Nikhil Agarwal  
 Associate Professor of Economics  
 Thesis Supervisor

Certified by .....  
 Esther Duflo  
 Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics  
 Thesis Supervisor

Certified by .....  
 Benjamin Olken  
 Jane Berkowitz Carlton and Dennis William Carlton Professor of Microeconomics  
 Thesis Supervisor

Accepted by .....  
 Amy Finkelstein  
 John & Jennie S. MacDonald Professor of Economics  
 Chairman, Departmental Committee on Graduate Studies

# Essays in Environmental and Development Economics

Allan Hsiao

Submitted to the Department of Economics  
on May 14, 2021, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Economics

## Abstract

**Chapter 1.** 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.

**Chapter 2.** Does electoral accountability discipline public spending? After the fall of Suharto, Indonesia held local elections for the first time in decades. I use a dynamic discrete choice framework to study how democratization affected the spatial allocation of public investment in healthcare infrastructure. On one hand, democratization limits distortions from Suharto-era biases toward certain areas, such as those within the patronage network. On the other hand, spillover effects are less internalized as districts become more focused on their own constituents.

**Chapter 3.** Many infrastructure investments have spatial effects that make optimal allocation a difficult, combinatorial problem. Schools are one such example: when graduates seek employment nationally and migrate, schools have effects that extend beyond local labor markets. But policy-makers often allocate infrastructure investments with simple rules like population cutoffs, ranked lists, and need-based formulas that do not account for spatial interdependencies. How effective are these simple rules compared to more sophisticated approaches? I use a spatial equilibrium model of individuals' education and migration decisions to study this question in the context of Indonesia's Sekolah Dasar INPRES program, the largest school construction program in history.

Thesis Supervisor: Nikhil Agarwal  
Title: Associate Professor of Economics

Thesis Supervisor: Esther Dufo  
Title: Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics

Thesis Supervisor: Benjamin Olken  
Title: Jane Berkowitz Carlton and Dennis William Carlton Professor of Microeconomics

## Acknowledgments

I owe much to my advisors, Nikhil Agarwal, Esther Duflo, Ben Olken, and Rob Townsend. Nikhil taught me how to tackle ambitious questions with intellectual discipline and rigor. Esther has shown me how research can improve lives and inspired me to pursue meaningful work. Ben has nurtured my interest in Indonesia and helped me navigate the richness of this important setting. Rob has been a source of guidance and encouragement from my earliest days at MIT. I hope to pass their kindness forward.

Before MIT, Melanie Morten and Gharad Bryan gave me a foundation for research. Working with them was more influential than I could have imagined, as one project in Indonesia turned into many more. Undergraduate classes with Nathan Nunn, Larry Katz, Claudia Goldin, and Kiran Gajwani introduced me to economics and set me on my current path.

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.

Finally, I am grateful to my family. Their support has made all the difference.

## Chapters

1	Coordination and Commitment in International Climate Action	5
2	Democratization and Infrastructure Investment	74
3	Allocation Rules for Infrastructure Investment	106

# **Chapter 1**

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

# 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<sub>2</sub> 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<sub>2</sub> 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<sub>2</sub> emissions from palm oil plantations

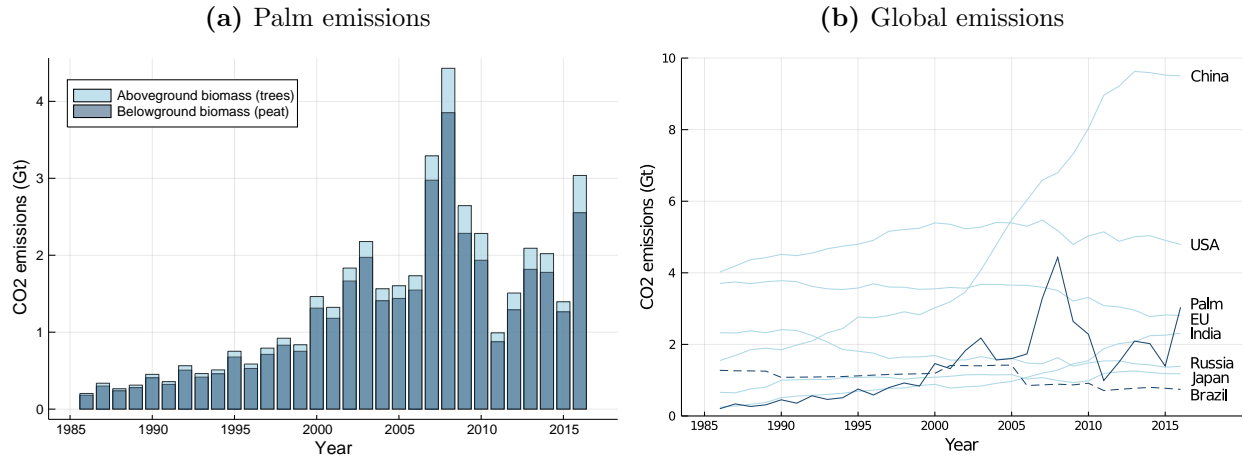


Figure 1a computes CO<sub>2</sub> emissions by combining data on palm oil plantations from Xu et al. (2020) and Song et al. (2018) with data on tree biomass from Zarin et al. (2016) and peat deposits from Gumbrecht et al. (2017). Figure 1b uses data on global CO<sub>2</sub> emissions by country, including 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<sub>2</sub> emissions during this period. I highlight Brazil, which also generates significant emissions through land-use change. These calculations focus on CO<sub>2</sub> emissions, which account for 73% of total greenhouse gas emissions during the study period. I focus on CO<sub>2</sub> 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.<sup>1</sup> 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,

---

<sup>1</sup> 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 an 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 Kalouptsi 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<sub>2</sub> 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  $\{Q_t^n, Q_{t+1}^n, \dots\}$ , and how the resulting production is allocated across markets. Given discrete time, discount factor  $\beta$ , new development  $Q_t^n = Q_t^{rn} + Q_t^{un}$ , and old development  $Q_{t+1}^o = Q_t^o + Q_t^n$ ,

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

#### Domestic regulation

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

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

where the tilde denotes net present value. There is no leakage problem because direct domestic regulation of supply achieves complete regulation. There is no commitment problem because the regulator can 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

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

where  $\varepsilon_t^S > 0$  and  $\varepsilon_{t+1}^{Du} < 0$  are elasticities of supply and unregulated demand, 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

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

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

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

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.<sup>2</sup> Figure 1b shows that palm emissions account for more CO<sub>2</sub> 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.

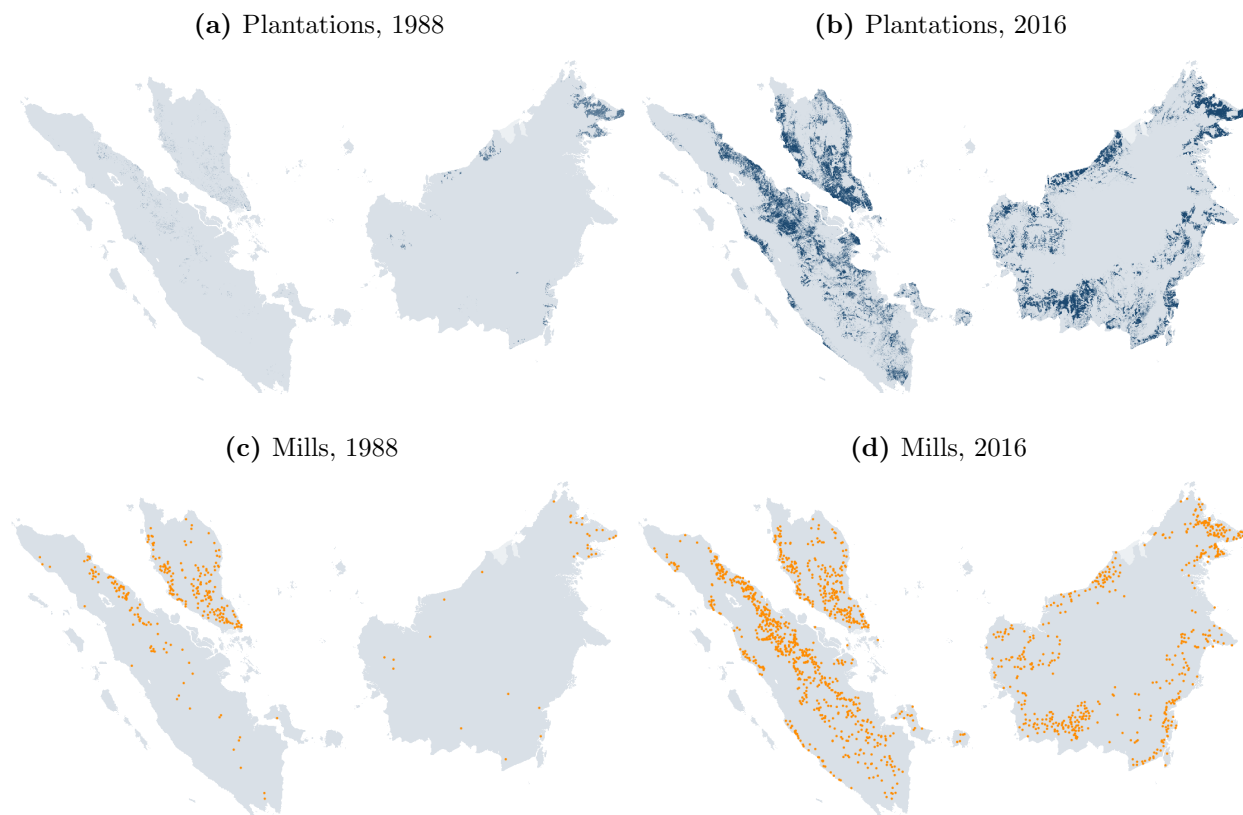
Third, foreign governments are actively discussing trade-policy interventions, particularly in Europe. French parliament debated a “Nutella” tax on palm oil in food products in 2016, although it failed to pass. Furthermore, the European Union initially provided green subsidies for palm-based biofuels, 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<sup>2</sup> – that measures the annual extent of plantations

<sup>2</sup> 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 Gumbrecht et al. (2017), respectively. I use conversion factors of 0.5 for biomass to carbon, 65.1 kg C/m<sup>3</sup> for peat to carbon, and 3.67 for carbon to CO<sub>2</sub>. 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).

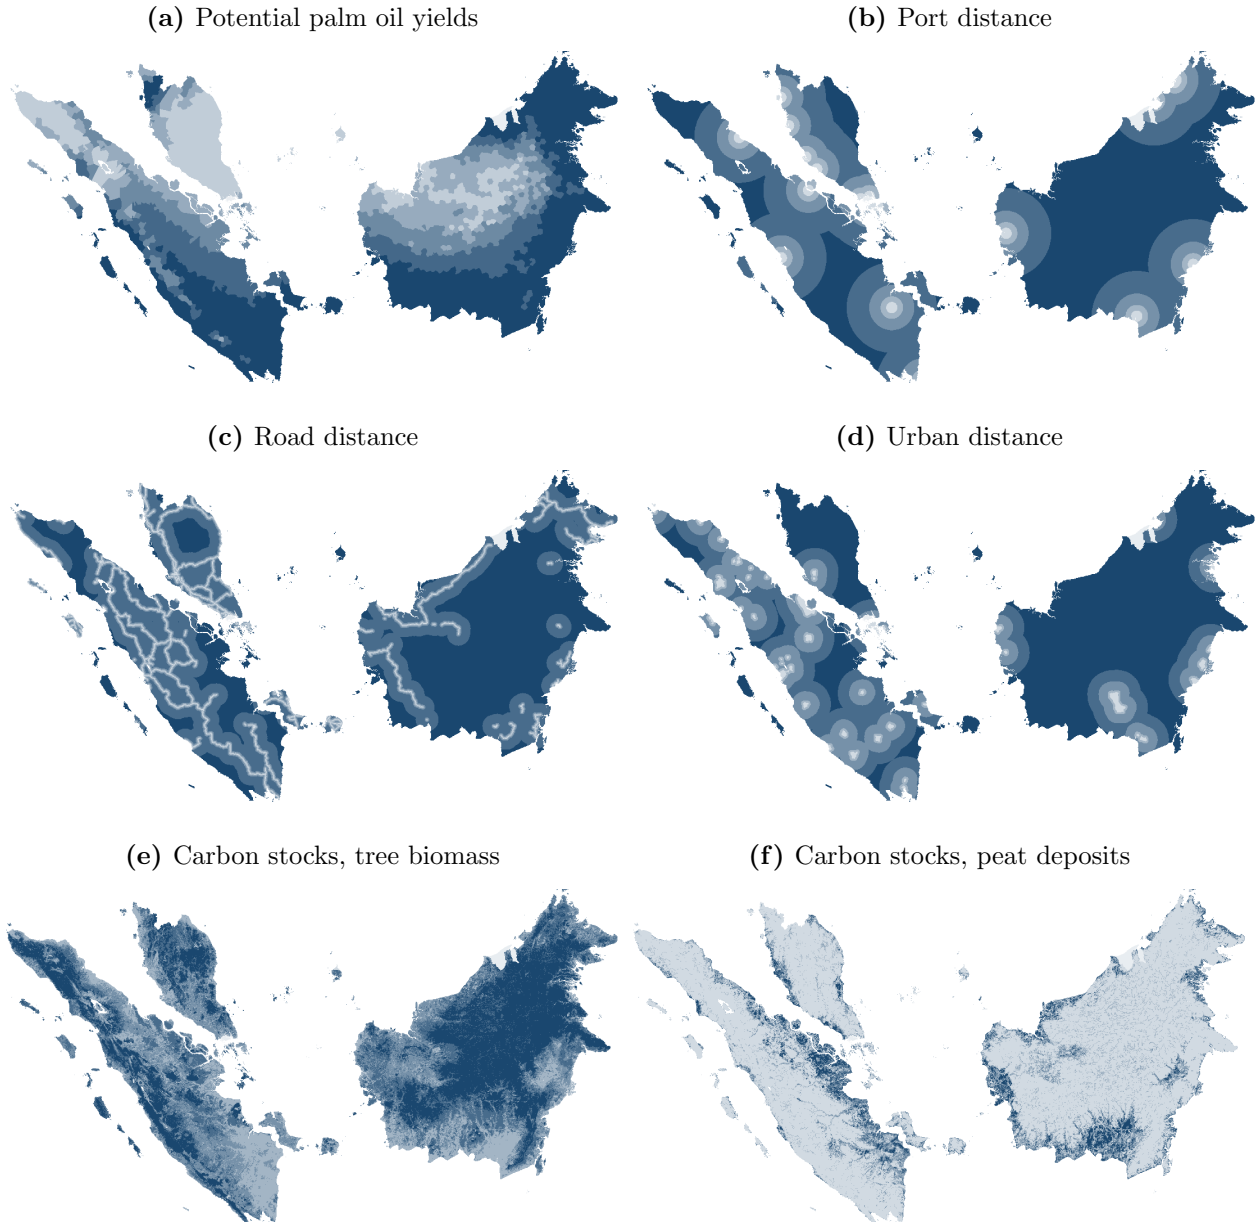
For consumption and world prices, I compile annual panel data from 1980 to 2016 on palm oil and its substitutes. Consumption data come from the USDA Foreign Agricultural Service and are measured at the country level, while world price data come from the International Monetary Fund and the World Bank. 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.<sup>3</sup> For each crop, year, and region, I compute rainfall shocks

---

<sup>3</sup> 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.

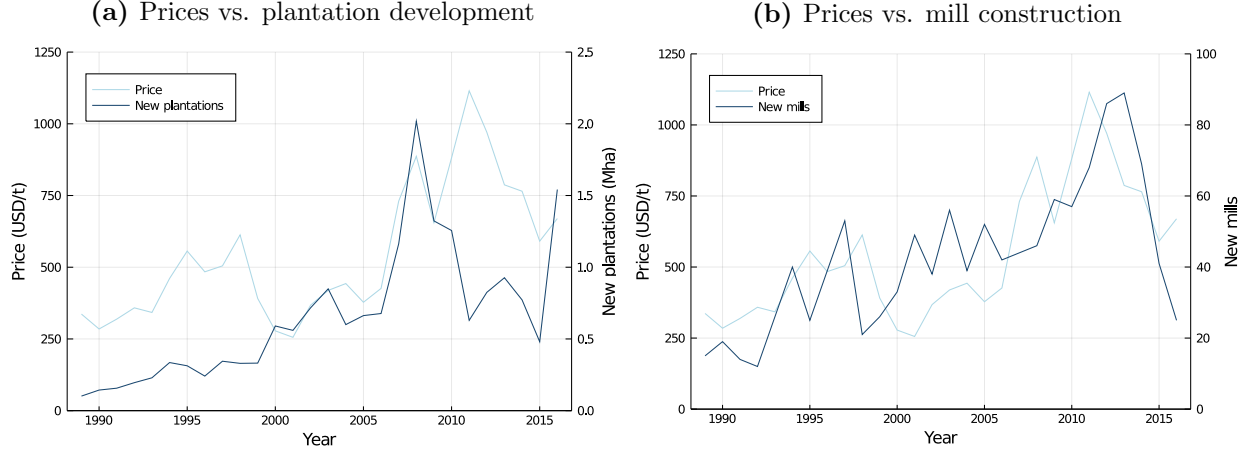
**Figure 3:** Land characteristics



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



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.<sup>4</sup> 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,

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

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

<sup>4</sup> 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  $\omega_{it}$ , palm and other oil prices  $p_{jt}$ , total vegetable oil expenditures  $X_t$ , and translog price index  $P_t$ . For the upper level,

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

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.

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

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  $\delta_{ij}$  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  $\gamma$  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

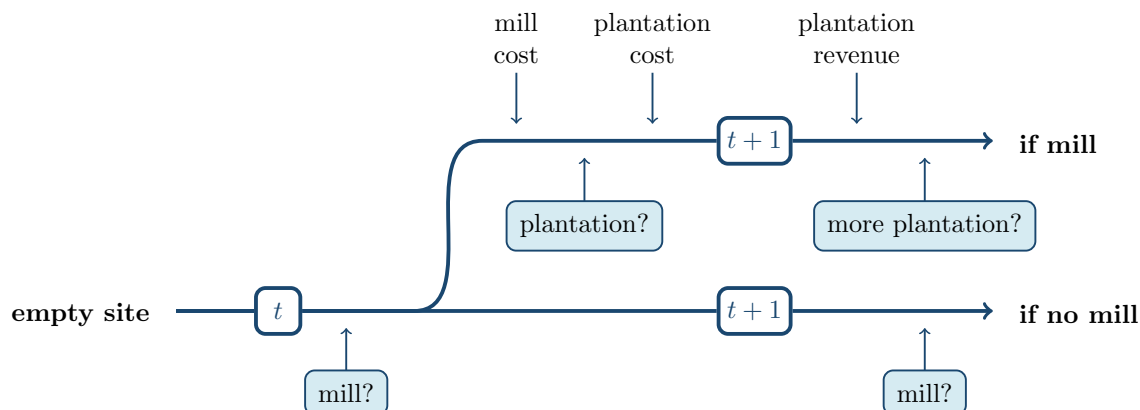
As is typical, prices are endogenous. Unobservables  $\varepsilon_{it}$  and  $\varepsilon_t$  shift demand and therefore affect equilibrium prices  $p_{jt}$ . I instrument with rainfall shocks in producing regions as a supply shifter. The exclusion restriction is that these shocks affect vegetable oil 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.<sup>5</sup> Figure 5 shows the timeline.

<sup>5</sup> 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  $\mathbf{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$  and privately observed state  $\varepsilon_{it}$ , 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  $\varepsilon_{it}$ , 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  $(\mathbf{w}_{it}, \varepsilon_{it})$  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

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

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

where investment increases future revenues while incurring some cost. Expectations are over next-period state  $(\mathbf{w}_{it+1}, \varepsilon_{it+1})$ , with shorthand  $\mathbb{E}_{it}[\cdot] \equiv \mathbb{E}[\cdot | s_{it}, \mathbf{w}_{it}, \varepsilon_{it}]$ . 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  $\kappa_m$  and regional time trends  $\alpha_m$ . Unobserved cost shocks  $\varepsilon_{it}$  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  $\varepsilon_{it}^Y$  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  $\mathbf{w}_{it} = \{Y_{it}, x_i, s_t, d_t\}$  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  $\sigma^e$ . For sites without mills, the timing of each period is as follows. First, sites realize state  $(\mathbf{w}_{it}, \varepsilon_{it}^e)$ . 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  $\varepsilon_{it}$ . 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

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

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

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

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

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

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

I suppress regional subscripts  $m$  in the notation, but these probabilities are more precisely  $p_m^e(\mathbf{w}_{it})$  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  $\mathbf{w}_{it+1}$ . 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  $\varepsilon_{it}$ . In equation 5d, cost factors  $x_i$  capture observed heterogeneity at the site level, and I accommodate regional unobserved heterogeneity  $\kappa_m^e$  and regional time trends  $\alpha_m^e$ . It is isomorphic to think of the difference in logit shocks  $\varepsilon_{it}^e$  as a cost shock. Finally, the outside option is never constructing a mill, with utility normalized to zero given mean-zero shocks  $\varepsilon_{it}^e$ .

### 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.<sup>6</sup>

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  $\varepsilon_{it}$ . 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  $Y_i^P$ . Demand shifters come

<sup>6</sup> 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.

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

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  $\varepsilon_{it}^e$  are assumed to be uncorrelated with plantation shocks  $\varepsilon_{it}$ .

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.<sup>7</sup> 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.

<sup>7</sup> The adding-up restrictions are  $\sum_i \alpha_i^0 = 1$ ,  $\sum_i \alpha_i^1 = 0$ ,  $\sum_i \beta_i = 0$ ,  $\sum_i \gamma_{ij} = 0 \forall j$  and are automatically satisfied since expenditure shares sum to one. Homogeneity imposes  $\sum_j \gamma_{ij} = 0 \forall i$ , such that proportional changes in prices and income have no impact on demand. Symmetry imposes  $\gamma_{ij} = \gamma_{ji} \forall i, j$ .

## 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<sup>2</sup>. 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.

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

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

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

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  $\eta_{it}$ , and I rearrange to obtain the regression equation

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

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

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

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.<sup>8</sup> [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  $\gamma$ ,  $\kappa_m$ , and  $\alpha_m$  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,

<sup>8</sup> 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 [Kalouptsi \(2014\)](#).

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.

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

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

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

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

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

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

where  $a_{it+1}^* \equiv a_{it+1}^*(0; \mathbf{w}_{it+1}, \varepsilon_{it+1})$ , 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  $a_{it+1}^*(a_{it}^*; \mathbf{w}_{it+1}, \varepsilon_{it+1}) = a_{it+1}^*(0; \mathbf{w}_{it+1}, \varepsilon_{it+1})$  given linear revenues,

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

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

I apply expectational errors  $\eta_{it}^e$  and substitute estimated values to obtain a regression equation.

$$\ln \left( \frac{\hat{p}^e(\mathbf{w}_{it})}{1 - \hat{p}^e(\mathbf{w}_{it})} \right) - \beta \ln \hat{p}^e(\mathbf{w}_{it+1}) = \hat{I}_{it+1} - (1 - \beta)x_i \gamma^e - (1 - \beta)\kappa_m^e - \alpha_m^e \tilde{t} + \eta_{it}^e \quad (10)$$

I estimate conditional choice probabilities  $\hat{p}^e(\mathbf{w}_{it})$  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  $\hat{a}_{it}^*$  in the same way, but with OLS instead of a logit regression. Intensive-margin profits  $\hat{I}_{it+1}$  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  $\gamma^e$ ,  $\kappa_m^e$ , and  $\alpha_m^e$  are interpretable in dollar terms. Intercepts  $\kappa_m^e$  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 non-stationarity 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  $\kappa_m$  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  $\kappa_m$  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*** (0.0374)	0.114*** (0.0328)	0.117*** (0.0354)	0.0828 (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)

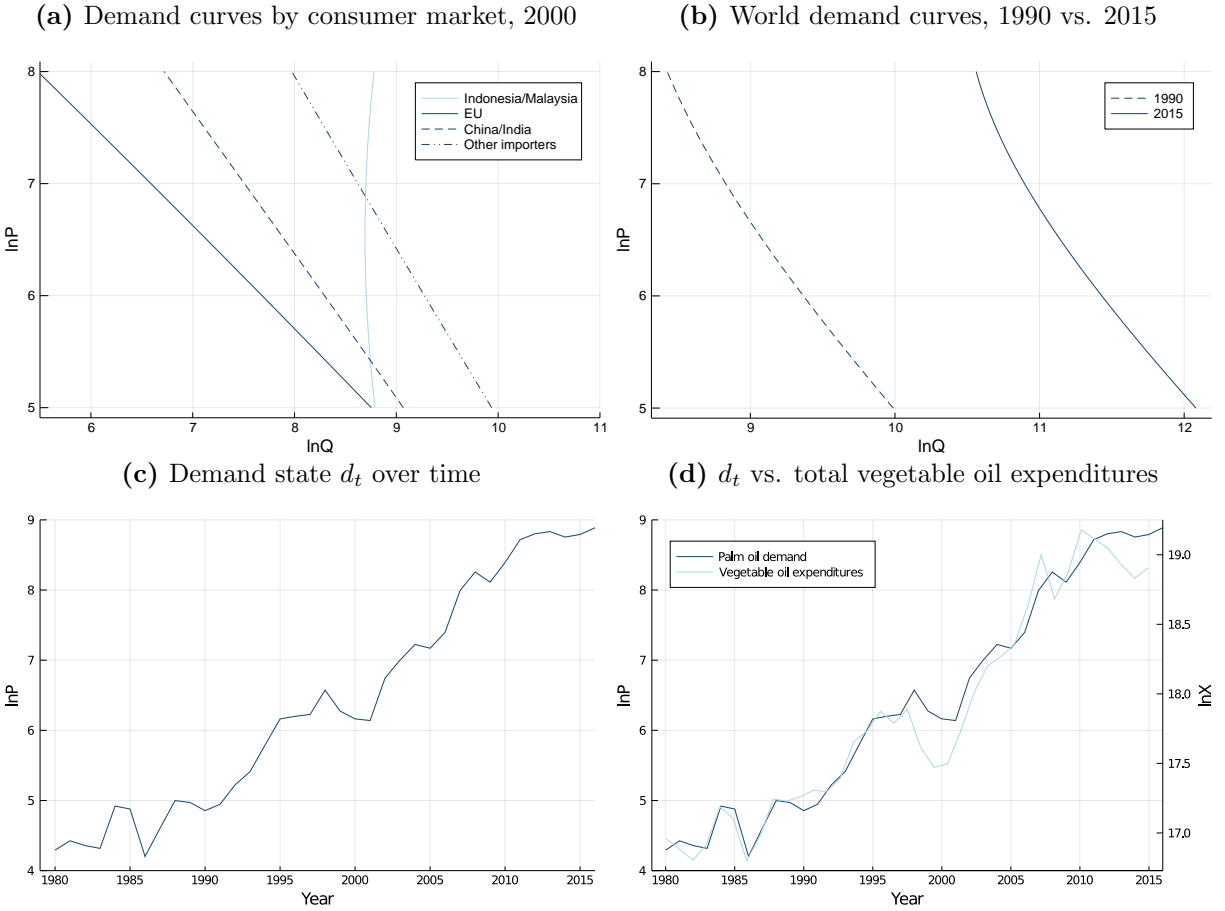
Uncompensated price elasticities are computed from estimated demand parameters using equation 3. Palm oil aggregates palm and palm kernel oil, while “other” oils include coconut, olive, rapeseed, soybean, and sunflower oil. I evaluate expenditure shares, prices, and the time trend at their averages over 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 $\times$ price ( $Y_{it+3}P_{t+3}$ )	0.113*** (0.00714)	0.200*** (0.0355)	
Potential yield $\times$ 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 ([Gumbrecht 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 ( $\kappa_m$ )	9,674***	(856)	6,398	14,655
Province-specific cost trends ( $\alpha_m$ )	-374***	(21)	-729	-99
Quadratic costs ( $\delta$ )	4.50***	(0.80)	–	–
Cost factors ( $\gamma$ )				
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 ( $\kappa_m^e$ )	22,881,886***	(391,964)	15,804,464	29,636,816
Province-specific cost trends ( $\alpha_m^e$ )	88,477***	(15,261)	-483,608	625,779
Logit scale ( $\sigma^e$ )	3,075,006***	(107,831)	–	–
Cost factors ( $\gamma^e$ )				
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  $\kappa_m$  and  $\kappa_m^e$  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  $\tau_t$  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  $\varepsilon_{it}$  and  $\varepsilon_{it}^e$ , 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	<u>\$/t NPV</u>	<u><math>\Delta\%</math></u>	<u><math>\Delta\%</math> surplus</u>			<u>\$/t CO<sub>2</sub></u>	
	Tax	CO <sub>2</sub>	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

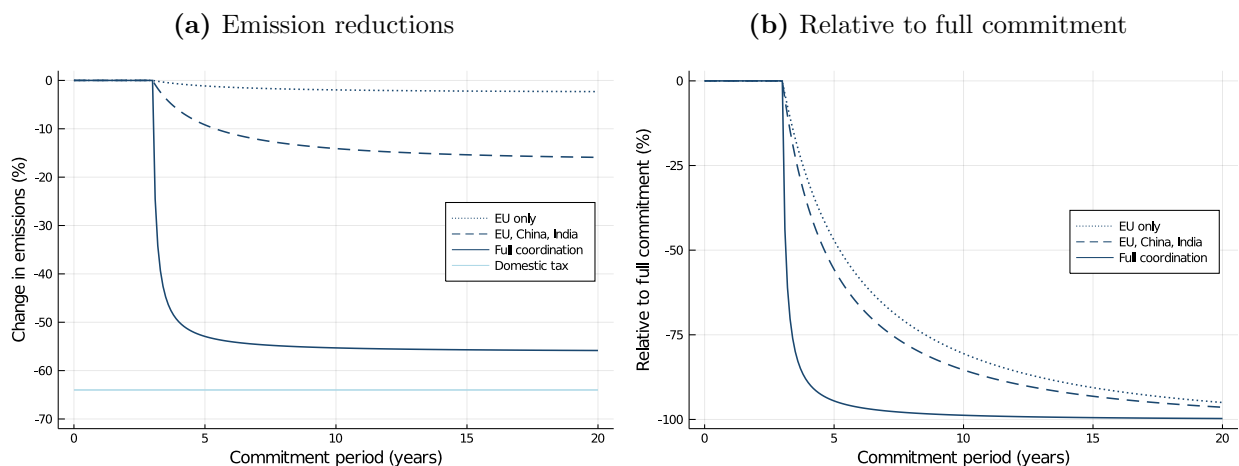


Figure 7a shows emission reductions under several scenarios. Starting at the top, the dotted line shows reductions under unilateral EU action for each of the commitment periods listed on the  $x$ -axis. Emission reductions are zero when the commitment period does not exceed time to build because otherwise 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<sub>2</sub>)

	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<sub>2</sub> 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.
- Fowle, Meredith. Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage. American Economic Journal: Economic Policy, 1(2):72–112, 2009.
- Fowle, Meredith, Mar Reguant, and Stephen Ryan. Market-Based Emissions Regulation and Industry Dynamics. Journal of Political Economy, 124(1):249–302, 2016.
- Gaveau, David, Bruno Locatelli, Mohammad Salim, Husna Yaen, Pablo Pacheco, and Douglas Sheil. Rise and Fall of Forest Loss and Industrial Plantations in Borneo (2000–2017). Conservation Letters, 12(3): e12622, 2019.
- Gumbrecht, Thomas, Rosa Maria Roman-Cuesta, Louis Verchot, Martin Herold, Florian Wittmann, Ethan Householder, Nadine Herold, and Daniel Murdiyarso. An Expert System Model for Mapping Tropical Wetlands and Peatlands Reveals South America as the Largest Contributor. Global Change Biology, 23: 3581–3599, 2017.
- Hall, Robert. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. Journal of Political Economy, 86(6):971–987, 1978.
- Hansen, Lars Peter and Kenneth Singleton. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. Econometrica, 50(5):1269–1286, 1982.
- Hansen, M. C., P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend. High-Resolution Global Maps of 21st-Century Forest Cover Change. Science, 342(6160): 850–853, 2013.
- Harstad, Bård. Buy Coal! A Case for Supply-Side Environmental Policy. Journal of Political Economy, 120(1), 2012.

- Harstad, Bård. The Market for Conservation and Other Hostages. Journal of Economic Theory, 166:124–151, 2016.
- Harstad, Bård. Technology and Time Inconsistency. Journal of Political Economy, 128(7):2653–2689, 2020.
- Harstad, Bård and Torben Mideksa. Conservation Contracts and Political Regimes. Review of Economic Studies, 84:1708–1734, 2017.
- Hausman, Jerry, Gregory Leonard, and J. Douglas Zona. Competitive Analysis with Differentiated Products. Annales d’Economie et de Statistique, 34:159–180, 1994.
- Helm, Dieter, Cameron Hepburn, and Richard Mash. Credible Carbon Policy. Oxford Review of Economic Policy, 19(3):438–450, 2003.
- 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.
- Kalouptsi, Myrto. Time to Build and Fluctuations in Bulk Shipping. American Economic Review, 104(2): 564–608, 2014.
- Kalouptsi, 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.
- Kopplitz, 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 Processes. Econometrica, 70 (2):801–816, 2002.

- Man, Elaine Lau Ying and Adam Baharum. A Qualitative Approach of Identifying Major Cost Influencing Factors in Palm Oil Mills and the Relations towards Production Cost of Crude Palm Oil. American Journal of Applied Sciences, 8(5):441–446, 2011.
- Markusen, James. International Externalities and Optimal Tax Structures. Journal of International Economics, 5:15–29, 1975.
- Marsiliani, Laura and Thomas Renström. Time Inconsistency in Environmental Policy: Tax Earmarking as a Commitment Solution. Economic Journal, 110(462):C123–C138, 2000.
- 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 Murdiyarto, 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  $\beta$ ,

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

where  $Q_{t+s}^o = Q_t^o + Q_t^n + Q_{t+1}^n + \dots + Q_{t+s-1}^n$ . Domestic regulation 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.

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

Assuming an interior solution  $Q_{t+1}^{o*}(\tilde{\tau}_t) > 0$ , 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

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

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

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

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

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

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

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  $\{\tau_{t+1}, \tau_{t+2}, \dots\}$ . 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

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

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

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

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

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

From the perspective of time  $t$ , tariffs  $\tau_t$  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.

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

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  $\tau_{t+s}$  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  $\frac{dQ_{t+s}^{ro}}{d\tau_{t+s}} = 0$  for  $s \in \{0, L, 2L, \dots\}$ . But  $\frac{dQ_{t+s}^{ro}}{d\tau_{t+s}}$  and  $\frac{dQ_{t+s}^{uo}}{d\tau_{t+s}}$  are time-invariant because the demand and supply curves are time-invariant, and thus  $\tau_{t+s} = \tau$  for all  $s \notin \{0, L, 2L, \dots\}$ . Furthermore, any response to announced tariffs will occur in the initial period. To see why, suppose not. Development in a later period must be profitable in that period, but if so then developing in the first period 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.<sup>9</sup>

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.

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

<sup>9</sup> 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  $(Q_{t+1}^{ro}, Q_{t+1}^{uo})$  when tariffs are in place and  $(Q_{t+L}^{ro}, Q_{t+L}^{uo})$  when they are not. In equilibrium,

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

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

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

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

Assuming an interior solution, the net present value of the stream of tariffs given by  $\tau$  is

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

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

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

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

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

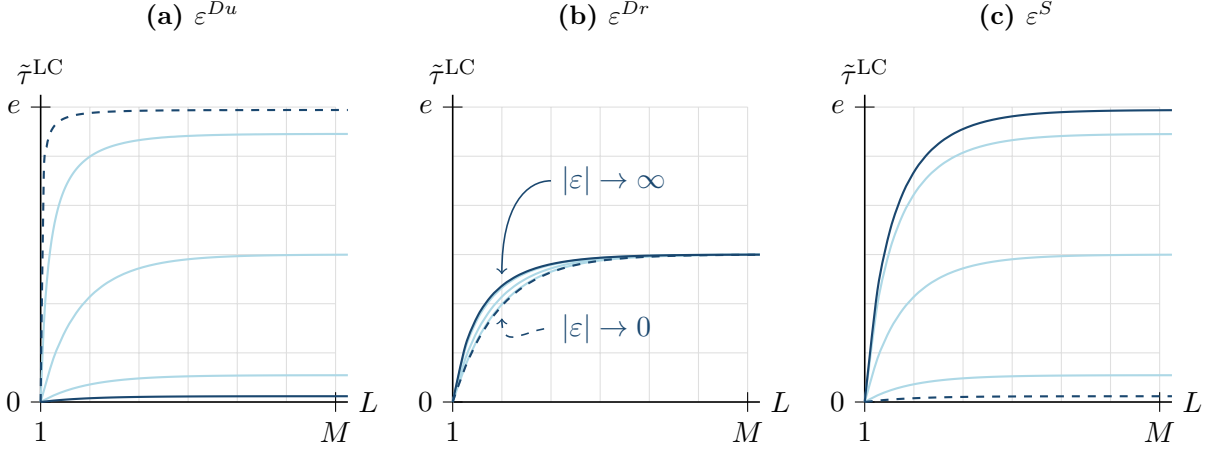
Lastly, the same mechanism also applies in the more general case if tariffs are declining over time. Indeed, importers that take a sequential static approach to setting tariffs will be governed by equations 11, which imply declining tariffs if the elasticity of supply is declining over time. Such will 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  $\varepsilon_{t+1}^{Du}$ ,  $\varepsilon_{t+1}^{Dr}$ , and  $\varepsilon_t^S$ ) and commitment (given  $L$ ) interact to determine total tariffs  $\tilde{\tau}_t^{\text{LC}}(L)$ . First,  $\tilde{\tau}_t^{\text{LC}}(L)$  increases more rapidly in  $L$  for smaller  $|\varepsilon_{t+1}^{Du}|$ .

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

**Figure A1:** Total tariffs by leakage and commitment



For various values of each leakage-relevant elasticity – namely elasticity of unregulated demand  $\varepsilon^{Du}$ , elasticity of regulated demand  $\varepsilon^{Dr}$ , and elasticity of supply  $\varepsilon^S$  – I plot the relationship between total tariffs  $\tilde{\tau}^{LC}$  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,  $\tilde{\tau}_t^{LC}(L)$  increases more rapidly in  $L$  for larger  $|\varepsilon_{t+1}^{Dr}|$ , although this effect is relatively small.

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

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

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

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.<sup>10</sup>

## 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

<sup>10</sup> 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.<sup>11</sup> 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.

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

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  $\tau^k$ ,

$$P^{Sk}(Q_1^{k*}(\tau^c, \tau^d)) = \frac{1}{1-\beta} \left( P^{Dr}(Q_1^{r*}(\tau^c, \tau^d)) - \tau^k \right) \quad \text{if } Q_1^{r*}(\tau^c, \tau^d) > 0 \text{ for } k \in \{c, d\},$$

$$P^{Sk}(Q_1^{k*}(\tau^c, \tau^d)) = \frac{1}{1-\beta} \left( P^{Du}(Q_1^{u*}(\tau^c, \tau^d)) \right) \quad \text{if } Q_1^{u*}(\tau^c, \tau^d) > 0 \text{ for } k \in \{c, d\},$$

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

$$\frac{dW_1}{d\tau} = \left( \frac{1}{1-\beta} \tau - e^c \right) \frac{dQ_1^{rc}}{d\tau} + \left( \frac{1}{1-\beta} \tau - e^d \right) \frac{dQ_1^{rd}}{d\tau} - e^c \frac{dQ_1^{uc}}{d\tau} - e^d \frac{dQ_1^{ud}}{d\tau} = 0,$$

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

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  $Q_1^u = 0$ .

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

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

<sup>11</sup> 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  $\tau$  under commitment, the objective function is

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

In equilibrium, marginal entry is indifferent between markets.

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

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

$$\begin{aligned} \frac{dW_1}{d\tau} &= -Q_1^r \frac{dP^{Dr}}{dQ_1^r} \frac{dQ_1^r}{d\tau} + \tau \frac{dQ_1^r}{d\tau} + Q_1^r - (1-\beta)e \frac{dQ_1}{d\tau} = 0, \\ \tau^C &= \underbrace{(1-\beta) \left( \frac{\varepsilon^S}{\varepsilon^S - \frac{Q_1^u}{Q_1} \varepsilon^{Du}} \right) e}_{\text{emissions}} + \underbrace{(1-\beta) \left( \frac{\frac{Q_1^r}{Q_1} P^S}{\varepsilon^S - \frac{Q_1^u}{Q_1} \varepsilon^{Du}} \right)}_{\text{terms of trade}}, \end{aligned}$$

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  $Q_1^u = 0$ .

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

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
<a href="#">Xu et al. (2020)</a>	2001-2016	Indonesia, Malaysia	Palm oil plantations over time, 100m resolution
<a href="#">Song et al. (2018)</a>	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
<a href="#">Gumbricht et al. (2017)</a>	2011	World	Peatlands and depth, 231m resolution
<a href="#">Zarin et al. (2016)</a>	2000	World	Aboveground biomass, 30m resolution
<a href="#">Hansen et al. (2013)</a>	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.<sup>12</sup> I aggregate the data to the 30-arc-second resolution (approximately 1 km<sup>2</sup>) 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.<sup>13</sup> 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  $t$ ,

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

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

<sup>12</sup> 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.

<sup>13</sup> 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 preceding 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<sup>2</sup> tiles). My core analysis respects this constraint: it centers on sites that are, on average, about 650 km<sup>2</sup>, 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 preceding 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.<sup>14</sup> 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.<sup>15</sup> 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.

---

<sup>14</sup> In the spatial merge, the closest mill within one kilometer is considered a match.


<sup>15</sup> 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

	$\Delta\text{Plantation}_t$	$\Delta\text{Plantation}_t$	$\Delta\text{Plantation}_t$
$\Delta\text{Tree cover}_t$	-0.00314*** (0.000156)	-0.00253*** (0.000155)	-0.00262*** (0.000153)
$\Delta\text{Tree cover}_{t-1}$	-0.00524*** (0.000192)	-0.00441*** (0.000191)	-0.00435*** (0.000190)
$\Delta\text{Tree cover}_{t-2}$	-0.00103*** (0.000194)	0.000199 (0.000193)	0.000408** (0.000194)
$\Delta\text{Tree cover}_{t-3}$	-0.000672*** (0.000162)	6.47e-05 (0.000161)	7.27e-05 (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.

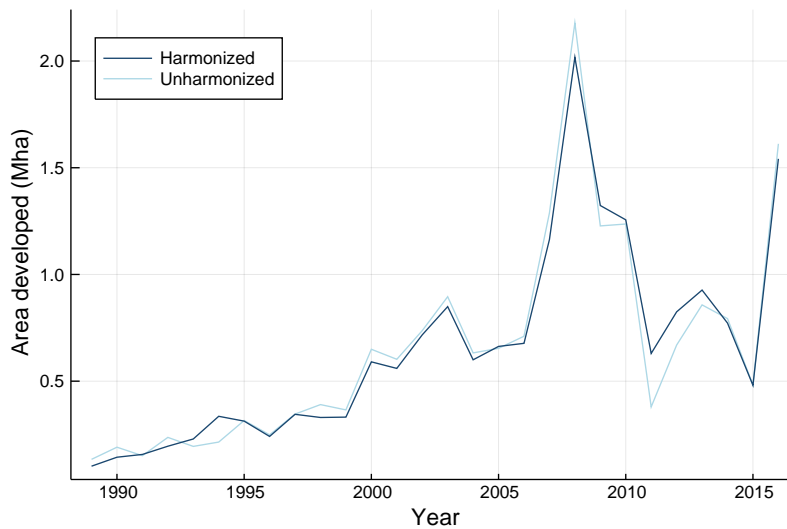
First, I impose that all plantation development have access to a mill within 50 kilometers.<sup>16</sup> In cases where no mill such mill is built by 2016, development is infeasible and I drop it entirely. In cases where such a mill does not exist in the current period but is eventually built, I delay 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).<sup>17</sup> 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.

<sup>16</sup> 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.

<sup>17</sup> 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

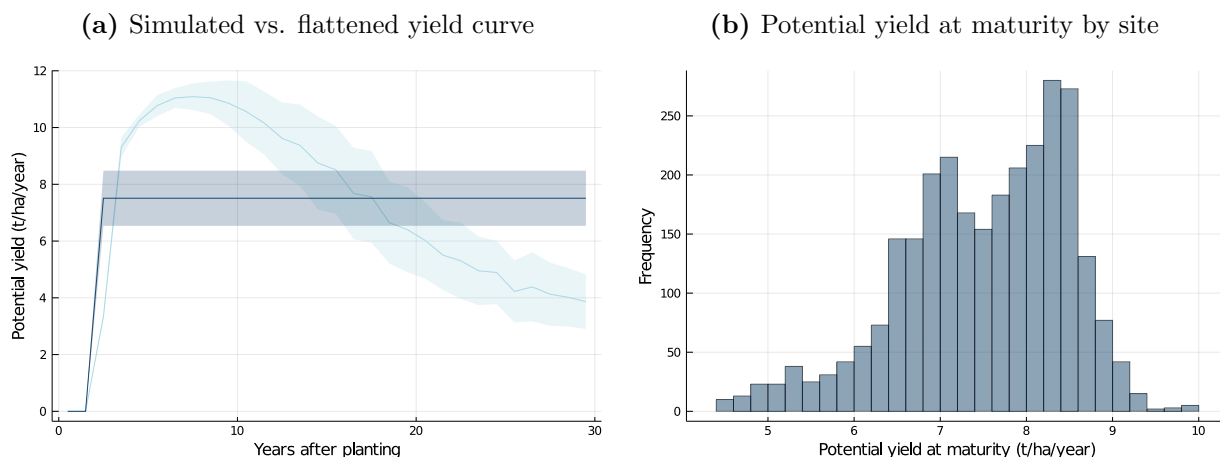
I construct data on palm oil yields by site over time by combining cross-sectional, site-level data on potential yields from the PALMSIM model of Hoffmann et al. (2014) with panel, province-level data on attained yields from government statistics. I proceed in the following steps:

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



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.<sup>18</sup> 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  $\gamma_{mt}$ .

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

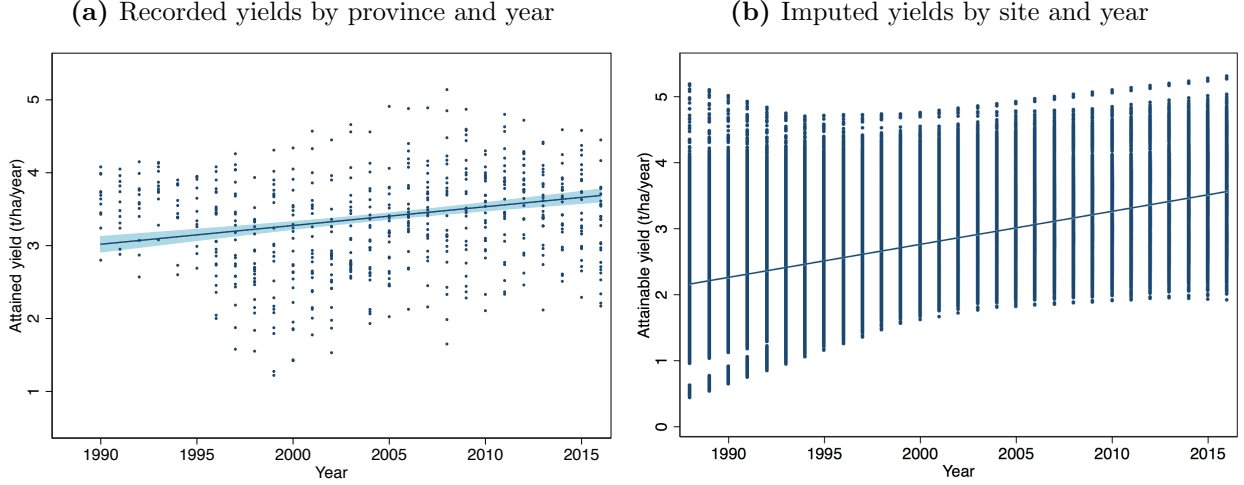
The underlying restriction is that, while potential yields are allowed to vary by site, yield gaps are fixed across sites in a given province-year. Site-year yields  $Y_{it}$  aggregate to the observed province-year yields  $Y_{mt}$  as

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

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

<sup>18</sup> 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



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  $\gamma_{mt}$  to obtain

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

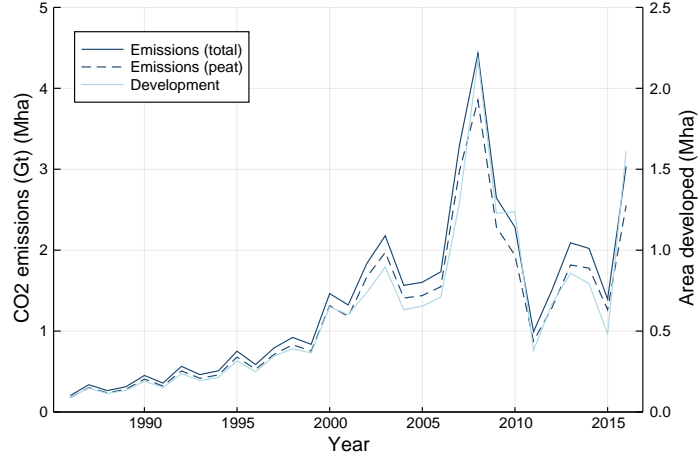
To isolate the underlying levels and trends of these yield gaps, I estimate the specification

$$\gamma_{mt} = \alpha_m + \beta_m t + \varepsilon_{mt}, \quad (16)$$

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 [Gumbrecht 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 [Gumbrecht 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 conversion factor of 0.5. To convert belowground biomass, I use the conversion factor of 65.1 kg C/m<sup>3</sup> 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 [Gumbrecht 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		Obs
	Estimate	SE	Estimate	SE	Estimate	SE	
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<sup>2</sup> (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<sup>2</sup> 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<sup>2</sup>. 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<sup>2</sup>. 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<sup>2</sup>, 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<sup>2</sup>. All three methods yield average site sizes in the neighborhood of 500 km<sup>2</sup>.

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  $C_1, C_2, \dots, C_k$ .
2. For the  $m$  mills observed in the data, move the  $m$  closest centers to the mill coordinates.
3. Assign points to the nearest cluster centers.
4. Update each cluster center by averaging over the points assigned to it.
5. Repeat (2) to (4) until convergence.

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**  $\eta_{it}$

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

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

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

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

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

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

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

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

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

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

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

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

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

## 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.

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

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  $\beta$ ) 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

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

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 $Y_{it}$ , $x_i$ , $\varepsilon_{it}$ , and $\varepsilon_{it}^e$

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  $\varepsilon_{it}$  and  $\varepsilon_{it}^e$ , they only know the distribution of future shocks.

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

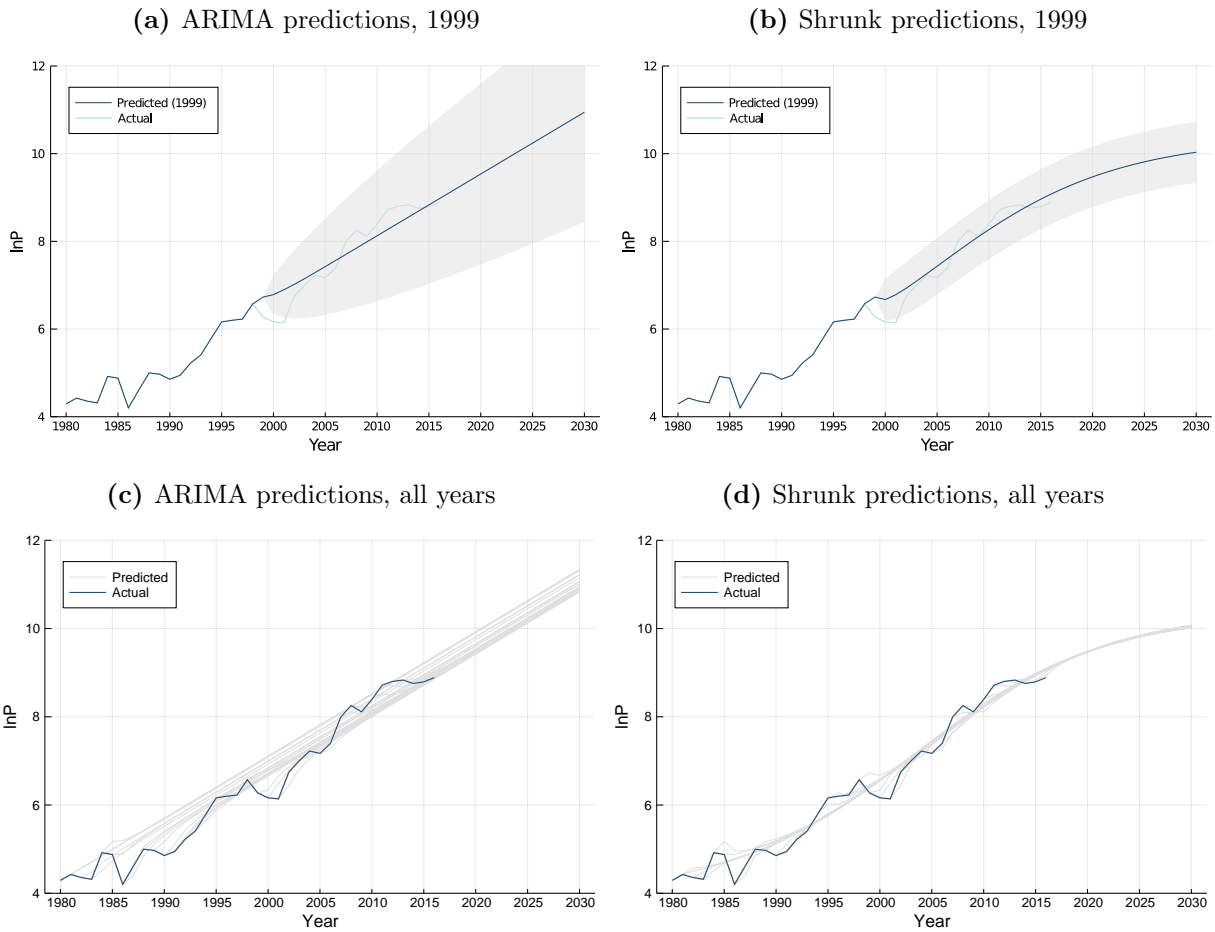
$$v_{it} = -\frac{1}{\delta}\varepsilon_{it} + \eta_{it}$$

**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  $\mathbb{E}_{it}[d_{t+s}]$



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

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

Thus, I can estimate cost shocks as a function of residuals  $v_{it}$  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  $\varepsilon_{kt}^e$ . 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  $s_S(a_{S-1}, a_{S-2}, \dots, a_1, s_1)$ .
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  $\{a_1, a_2, \dots, a_S\}$ .

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.

$$s_t = \sum_i Y_{it} s_{it}, \quad a_t = \sum_i \left( s_{it}^e a_{it} + (1 - s_{it}^e) p_{it}^e a_{it} \right), \quad (18)$$

where for sites without mills in period  $t$  ( $s_{it}^e = 0$ ), new development depends on both extensive-margin probability  $p_{it}^e$  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  $T$  ( $s_{iT}^e = 1$ ), the first order condition for  $a_{iT}$  determines development.

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

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

$$p_{iT}^e = \frac{\exp(-x_i \gamma^e - \kappa_m^e - \alpha_m^e T + \mathbb{E}_{iT}^e[V(0; \mathbf{w}_{iT}, \varepsilon_{iT})])}{1 + \exp(-x_i \gamma^e - \kappa_m^e - \alpha_m^e T + \mathbb{E}_{iT}^e[V(0; \mathbf{w}_{iT}, \varepsilon_{iT})])}, \quad (20)$$

where the one in the denominator arises from  $v^e(0; \mathbf{w}_{iT}) = 0$  since there is no further entry after period  $T$  (for an outside option normalized to zero).<sup>19</sup> 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  $\mathbf{a}_T = \{a_{iT}, a_{iT}^e\}$  as a function of previous development  $\mathbf{s}_T = \{s_{iT}, s_{iT}^e\}$ .

<sup>19</sup> To determine the probability of extensive-margin entry, I compute intensive-margin profits assuming  $\mathbb{E}_{iT}^e[\varepsilon_{iT}] = 0$  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  $\varepsilon_{it}$ . Furthermore, since intensive-margin profits  $\mathbb{E}_{iT}^e[V(0; \mathbf{w}_{iT})]$  are not linear in  $\varepsilon_{iT}$  (even though choices  $a_{iT}$  are), I cannot simply apply  $\mathbb{E}_{iT}^e[\varepsilon_{iT}] = 0$  and must instead compute expected intensive-margin profits based on the distribution of  $\varepsilon_{iT}$ , 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  $\mathbf{a}_T$  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  $\mathbf{a}_T$ . 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  $p_{iT}^e$ . 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  $\mathbf{a}_{T-1}$  in period  $T - 1$ . Taking previous development  $\mathbf{s}_{T-1}$  as given, I solve for new development  $\mathbf{a}_{T-1}$  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  $\mathbf{a}_{T-1}$ . Since the first order condition is monotonic in prices, only one such division exists.
3. With  $\mathbf{s}_{T-1}$  and  $\mathbf{a}_{T-1}$ , I obtain site-specific  $\mathbf{s}_T$  and therefore total  $s_T$ .
4. Given  $\mathbf{s}_T$ , I solve the subproblem for  $\mathbf{a}_T$  using the solution algorithm described above for entry in period  $T$ , after which there is no further entry. With  $\mathbf{s}_T$  and  $\mathbf{a}_T$ , I obtain site-specific  $\mathbf{s}_{T+1}$  and therefore total  $s_{T+1}$ .
5. Given totals  $s_T$  and  $s_{T+1}$ , I compute site-specific  $\mathbf{a}_{T-1}$  with analogues of equations [19](#) and [20](#).<sup>20</sup>
6. Finally, I check if site-specific new development  $\mathbf{a}_{T-1}$  sums to the guess for total new development  $a_{T-1}$ . If so, then  $\mathbf{a}_{T-1}$  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](#).

$$\text{acacia}_i = \beta_0 + \beta_1 \text{palm}_i + \beta_2 \text{mill\_distance}_i + \alpha_m + \varepsilon_i, \quad (21)$$

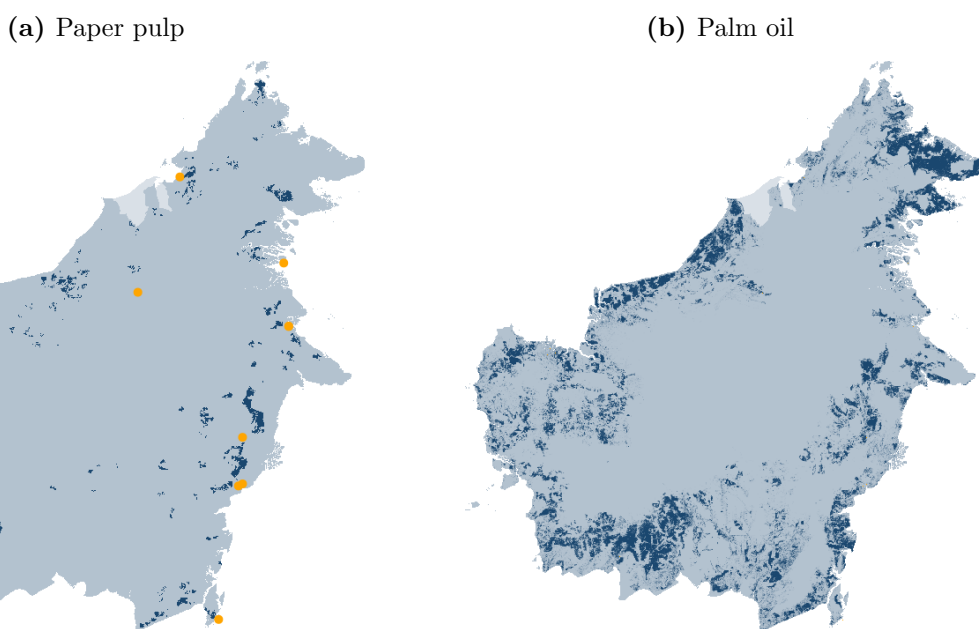
<sup>20</sup> For intensive-margin entry in equation [19](#), the analogue in period  $T - 1$  is similar except that prices depend on  $s_T$  in period  $T$  and  $s_{T+1}$  thereafter. A firm's expected development  $a_{iT}$  in period  $T$  does not enter. For extensive-margin entry probabilities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

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

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



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*** (0.00610)	-0.0235*** (0.00734)
Log paper pulp mill distance (km)	-0.0265*** (0.00447)	-0.0210*** (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: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

## Chapter 2

Democratization and Infrastructure Investment:  
Evidence from Healthcare in Indonesia

# 1 Introduction

Infrastructure investment is at the heart of economic development. The world invests more than \$2 trillion each year in infrastructure, and this figure continues to grow ([Oxford Economics 2017](#)). But while these investments have long-lasting benefits when allocated efficiently, they are often targets of corruption. Particularly in countries with weak institutions, funds may go missing ([Olken 2007](#)) or be distributed by favoritism ([Burgess et al. 2015](#)). More broadly, corruption accounts for low economic growth ([Mauro 1995](#)). This paper asks whether electoral accountability limits corruption in the context of public infrastructure spending.

I study how the Indonesian government allocated new healthcare facilities – hospitals, clinics, and subclinics – over the last three decades in its efforts to expand healthcare coverage. This expansion spans Indonesia’s democratization in 1999, allowing me to compare the allocation of new facilities before and after democratization. To do so, I quantify the welfare effects of new facilities by estimating a spatial model of demand for healthcare. I then model the government’s allocation decision as a dynamic discrete choice problem, and I use revealed-preference techniques to estimate how it weighs the benefits of new facilities for citizens against a range of favoritism motives.

The fall of Suharto ushered in democratization with Indonesia’s first free elections since 1955, as well as the decentralization of decision-making power from Jakarta to local governments. The result was *local* electoral accountability, and indeed this bundling of democratization and decentralization is common in other settings as well ([Gadenne and Singhal 2014](#); [Mookherjee 2015](#)). While local elections may reduce corruption by increasing electoral accountability, they may also introduce their own distortions as constituents take priority over non-constituents. In particular, there are welfare losses when investments have spillover effects on non-constituents that local governments fail to internalize.

To assess these competing effects, I begin by quantifying the consumer surplus generated by new facilities. I do so by modeling demand for healthcare facilities over space and estimating the model with geocoded panel data on facility access and usage. Conditional on being sick, individuals choose among visiting their closest public hospital, private hospital, clinic, or subclinic, or the outside option of not seeking treatment. They have disutility from the distance and congestion of any given facility, and the staggered rollout of new facilities over time generates panel variation in both. The estimated demand system allows me to compute the consumer surplus generated by any given spatial allocation of facilities.

Next, I evaluate misallocation for each district by comparing consumer surplus under the actual allocation to the maximum achievable with the same budget. Misallocation is zero when the actual and maximizing allocations coincide. I use stochastic optimization techniques to approximate the maximum achievable surplus because the high dimensionality of the solution space makes conventional algorithms intractable. Structural estimation often avoids stochastic algorithms: in

the non-convex settings that motivate stochastic algorithms to begin with, these algorithms only provide good approximations of optimal objective values, but not necessarily of the optimal parameters themselves. This discrepancy is a major problem when the goal is to assign an economic interpretation to the parameter estimates, and so stochastic algorithms are more common in applications like machine learning, where the goal is prediction. My application sidesteps this challenge because my measure of misallocation is based on the achieved objective value (consumer surplus) and not on the maximizing values (the allocation itself). I find relatively large levels of misallocation, particularly before democratization. That is, the actual allocation achieves only 60% of achievable surplus because new facilities do not go to the places that would benefit most.

To understand why, I model the facility placement decision as a dynamic discrete choice problem. The previous analysis relies only on the structure of the demand model, but it can only compare misallocation across districts. Additional structure on the supply side allows me to analyze within-district misallocation at the village level. By revealed preference, I estimate the surface of village-level preferences that rationalizes the observed deviations from surplus maximization. I then show how these preferences map onto observable village characteristics in line with the channels described above. For estimation, I use moment-inequality techniques that circumvent the high dimensionality of the problem by comparing the actual allocation with a subset of local deviations (Pakes 2010; Pakes et al. 2015; Holmes 2011). This approach simplifies the dynamics of the problem by holding long-term allocations fixed, achieving finite dependence as in Arcidiacono and Miller (2011). I allow placement decisions to depend on unobservables that I accommodate flexibly and allow to vary at a fine level of spatial disaggregation.

My main finding is that democratization decreases misallocation overall. The structural estimates show that Suharto-era biases toward certain areas, such as those within the patronage network, are substantially lower after democratization. At the same time, spillover effects are less internalized as districts become more focused on their own constituents, but this effect is smaller than the first. I also find reduced-form evidence that supports this narrative. For electoral accountability, I use variation in the appointment dates of Suharto-regime district mayors, who were allowed to complete their terms after the Suharto's fall (Martinez-Bravo et al. 2017). Suharto mayors were not subject to electoral accountability – they were all replaced by elected officials – and I find misallocation to be higher in Suharto-mayor districts. For uninternalized spillovers, I use variation in the timing of redistricting, by which districts split into smaller districts (Burgess et al. 2012; Bazzi and Gudgeon 2017). I find no significant effect on district-level misallocation, consistent with only muted distortions from uninternalized spillovers.

I contribute to the literature on misallocation by proposing a measure of misallocation in infrastructure investment and identifying mechanisms for the misallocation I observe in the data. Hsieh and Klenow (2009) consider factor misallocation across firms and its associated productivity losses, and Bryan and Morten (2019) perform a similar exercise to quantify the spatial misalloca-

tion of labor. Compared to these papers, I consider the impacts of a particular set of reforms that changed the allocation process within my study period. Similar to [Asker et al. \(2019\)](#), I define misallocation based on the difference between achieved and achievable results. For infrastructure investment in particular, [Fajgelbaum and Schaal \(2017\)](#) show how to compute optimal road networks, and [Balboni \(2019\)](#) estimates the consequences of coastal bias in road investment as sea levels rise. These papers make important progress in quantifying misallocation under an assumed social objective function, but they are not geared toward determining whether the misallocation they measure results from irrational behavior, or simply a social objective function that differs from the ones used in their models. In comparison to these papers, I provide a specific measure of misallocation, and I take seriously the idea that observed misallocation is not the result of irrationality. Rather, I estimate the government objective function in order to understand both why I see deviations from surplus maximization and why these deviations change over time.

A large literature in political economy highlights potential mechanisms driving misallocation in this context. [Seabright \(1996\)](#) emphasizes voter information in a theoretical model to argue that local elections improve accountability by allowing local issues take center stage. Empirically, [Ferraz and Finan \(2008\)](#) show that voters hold candidates accountable by responding to performance, and [Casey \(2015\)](#) finds that the shift to local elections increases political accountability by empowering voters. This work provides support for the broader finding that democracy facilitates economic growth ([Acemoglu et al. 2019](#)). At the same time, decentralized decision-making may be socially suboptimal in the presence of spillovers and economies of scale ([Oates 1972](#)). [Sigman \(2002\)](#), [Kahn et al. \(2015\)](#), and [Lipscomb and Mobarak \(2017\)](#) show empirically that uninternalized spillovers result in lower water quality where water flows from one jurisdiction into another. My contribution is to provide an empirical framework for quantifying these mechanisms in the context of one of recent history’s most aggressive efforts to expand access to healthcare.

## 2 Institutional Details

This section describes the democratization and decentralization reforms that followed the end of Suharto’s decades-long regime in Indonesia. It also describes the Indonesian healthcare system.

### 2.1 Democratization

In response to mounting pressure domestic and international pressure, interim president Habibie announced free elections in 1999 – the first since 1955. Previously, Suharto’s regime had suppressed opposition parties by forcing them to merge into two parties, one Islamic and one non-Islamic, controlling the opposition leadership, and implementing a recall system that enabled the removal of individual legislators. As a result, between 1973 and 1998 Suharto’s party, *Golkar*, won landslide victories in five legislative elections.

District governments are headed by mayors, and under the Suharto regime these mayors were appointed by the central government. Mayors became subject to elections under democratization, but only after the end of the Suharto mayors' five-year terms. Since nearly no Suharto mayor won reelection, these mayors largely operated without electoral concerns. Term end dates varied by district, and [Martinez-Bravo et al. \(2017\)](#) establishes that this variation is quasi-random. Later reforms in 2005 brought the direct election of district mayors.

## 2.2 Decentralization

Following Suharto's fall, the transitional government passed legislation calling for the transfer of power to local governments (Laws 22/1999 and 25/1999). Decentralization proceeded at a rapid pace and placed district governments at the center of Indonesia's governance structure. Within two years, local district governments received additional authority in the form of two million civil servants, 30% of government expenditures, and responsibility for the provision of a range of public services. Today, district governments perform the majority of administrative functions, particularly as they pertain to the provision of public goods, while the central retains power over issues of national importance, such as foreign affairs and defense.

Furthermore, after 2001 a number of new district governments were established as existing districts split into smaller districts. I refer to this reform as "redistricting." District governments were required to apply for central government approval to redistrict, and the central government placed a moratorium on redistricting from 2004 and 2006 and again from 2009 to 2012. [Burgess et al. \(2012\)](#) and [Bazzi and Gudgeon \(2017\)](#) argue that the timing of redistricting around the first moratorium is plausibly exogenous.

"Districts" are subdivisions of provinces (*provinsi*) and refer collectively to both regencies (*kabupaten*) and cities (*kota*). Districts are subdivided into subdistricts (*kecamatan*), which are further subdivided into rural villages (*desa*) and urban neighborhoods (*kelurahan*). These rural villages and urban neighborhoods form the smallest administrative entities in Indonesia, and in this paper I refer to both as "villages." In sum, the administrative hierarchy is as follows: nation, province, district, subdistrict, and village.

## 2.3 Healthcare

Indonesia relies on a multi-layered, referrals-based system to provide care to its 260 million citizens. Below, I summarize the institutional details that are relevant for this paper.

### **The public healthcare system is layered**

The public system consists of hospitals, clinics, and smaller facilities. Hospitals are themselves divided into classes: class A hospitals average 1,450 beds and cover a range of specialties, while class D hospitals are district-level facilities that average 70 beds and offer only general care.

Below hospitals are clinics (*puskesmas*), which are usually staffed by a physician and focus on providing primary care. Some clinics are equipped to provide basic inpatient services. Clinics are further supported by a network of subclinics (*pustu*) and village facilities, including village health posts (*poskesdes*), village maternity posts (*polindes*), and neighborhood health posts (*posyandu*). Subclinics are staffed with one to three nurses and visited weekly to monthly by a physician. Village facilities are often staffed by local volunteers trained by health workers and may operate on borrowed premises.

### **Access expanded with infrastructure**

The Indonesian government has expanded access to healthcare services by devoting significant resources to building infrastructure.<sup>1</sup> Since the origins of the clinic system in the 1970s, the government has worked toward its formal goal of one clinic per 30,000 people or subdistrict, and one subclinic per 10,000 people. In the 1990s, the government implemented the *Bidan di Desa* initiative, which sought to station a midwife in every village. At the same time, the hospital network has continued to grow. Figure 1 shows the visible expansion of hospitals in Java, Indonesia's most populous island, since 1990. Today, there are about 2,500 hospitals, 10,000 clinics, and 25,000 subclinics in Indonesia.

### **Under decentralization, district governments place new facilities**

Before decentralization, the central government funded facility construction and possessed broad authority over the placement of new facilities. The clinic system, for example, was originally funded by the same INPRES program that funded the large-scale construction of more than 60,000 schools in the 1970s. Later, central funding continued through the Ministry of Health.

Since decentralization in 2001, district governments have been responsible for the direct implementation of healthcare services. Funding continues to come from the central government, including through disbursements from the Special Allocation Fund (*Dana Alokasi Khusus*) that are earmarked for facility construction. District governments negotiate with the central government for funds, including with proposals for new facilities. But afterwards, the central government cannot enforce agreed-upon proposals, and in some cases has limited information on the completion status of funded projects. It is therefore district governments that choose the placement of budgeted facilities.

### **The private system primarily serves the wealthy**

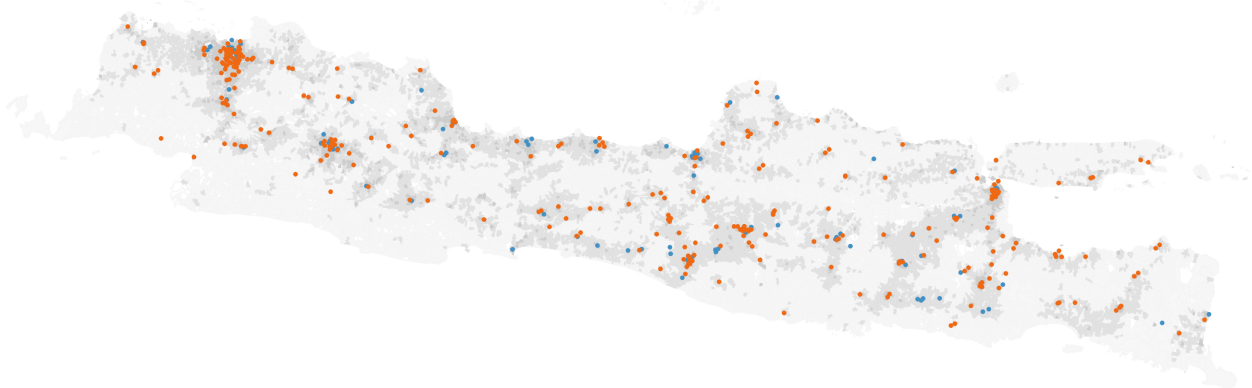
Private hospitals cater to the wealthy and operate outside of the public system. Growth in the number of private hospitals has largely involved the establishment of smaller, single-specialty hospitals – particularly in dentistry. Private doctor practices (*praktek dokter*) and polyclinics

---


<sup>1</sup> The Indonesian government has also expanded insurance coverage, launching universal healthcare – *Jaminan Kesehatan Nasional* (JKN) – in 2014 with plans to achieve full coverage by 2019. This program builds on the *Askeskin* (2004) and *Jamkesmas* (2008) programs, which provided coverage to the poor and near poor.

**Figure 1:** Hospitals versus population density in Java, 1990-2014

(a) Java 1990



(b) Java 2014



Orange dots are public hospitals, blue dots are private hospitals, and gray shading conveys population density. There are 390 hospitals in 1990 and 1,258 in 2014. Expansion seems strongest in areas of initial concentration and high population density – the cluster to the northwest is Jakarta. Data are from PODES.

(*poliklinik*) are the private counterparts to public clinics and often result from public doctors who open secondary practices. These facilities also serve relatively wealthy clientele, although to a lesser extent than private hospitals do.

### 3 Data

Village-level data on health infrastructure come from the Village Potential Statistics (PODES), a census of Indonesian villages conducted every few years. I use data from 1990 to 2014, and I merge the data over time using village locations. The core data cover hospitals, clinics, subclinics, and village-level facilities, and record the number of facilities for each type by village. In 2011, the data contain information on facility quality for clinics, subclinics, and village-level facilities. The PODES data also contain village-level voting results in the 1999 and 2004 legislative elections. The PODES data do not distinguish between public and private hospitals, so I also draw on Rumah Sakit Online (RSO), an online database of hospitals maintained by the Indonesian Ministry of

**Table 1:** Summary statistics by year (PODES)

Year	1990	1993	1996	2000	2003	2006	2008	2011	2014
Public hospitals	664	750	798	863	942	1,084	1,279	1,526	1,840
Private hospitals	231	260	282	307	351	395	465	544	654
Clinics	5,202	6,021	6,435	6,868	7,199	7,719	8,533	9,398	10,788
Subclinics	12,412	15,660	17,140	19,154	20,196	21,480	23,217	24,767	27,744
Distance, public hospital	30.58	28.71	28.36	27.43	26.21	23.91	21.51	19.72	18.32
Distance, private hospital	66.22	65.78	64.63	63.17	61.69	59.36	56.36	53.60	50.83
Distance, clinic	6.95	6.32	6.07	5.75	5.49	5.17	4.68	4.34	4.07
Distance, subclinic	4.21	3.46	3.18	2.79	2.63	2.47	2.25	2.07	1.81
Congestion, public hospital	442.83	424.52	418.10	374.01	363.22	328.79	296.96	258.44	225.54
Congestion, private hospital	1,274.83	1,280.61	1,261.05	1,127.77	1,087.34	1,035.04	984.82	929.34	811.11
Congestion, clinic	38.13	33.90	32.41	29.67	29.89	29.11	28.54	27.04	25.28
Congestion, subclinic	25.97	20.44	19.30	16.77	16.67	16.41	16.42	16.41	15.21
Rural (dummy)	0.89	0.89	0.89	0.89	0.82	0.81	0.82	0.81	0.81
Population	2,901	3,026	3,090	3,022	3,199	3,325	3,529	3,682	3,776
Area (km <sup>2</sup> )	26.69	26.69	26.69	26.69	26.69	26.69	26.69	26.69	26.69
Observations	62,194	62,194	62,194	62,194	62,194	62,194	62,194	62,194	62,194

Each observation is a village. The first four rows are totals, and all other rows are averages. Distance is to the closest facility of a given type and is measured in kilometers. Congestion is of the closest facility of a given type and is measured as the number of people (in thousands) for whom this facility is the closest of its type.

Health. This database lists approximately 2,500 hospitals and contains information on address, type (public or private), number of beds, number of personnel, and some measures of hospital quality. I use the RSO distinction between public and private hospitals in 2016 to classify hospitals in the PODES data.

The National Socioeconomic Survey (SUSENAS) dataset contains annually collected, individual-level data on healthcare usage by facility type. Unlike the PODES data, these data do distinguish between public and private hospitals. The data also contain demographic information and a limited set of health outcomes. Data with village-level locations are available from 1993 to 2010. Village locations allow me to link the data over time, as well as to the PODES data. Furthermore, villages are sufficiently small that I can geocode the data and calculate distances between individuals and facilities with a relatively high degree of accuracy.<sup>2</sup>

Table 1 summarizes the village-level data by year. For the more than 62,000 villages for which I was able to construct a balanced panel, the number of facilities has grown over time for all facility types. The panel data cover a tripling of hospitals from 1990 and a doubling of clinics and subclinics. Consistent with this growth, facility distance and congestion has declined.

<sup>2</sup> An empirical concern is that, within a district, most hospital construction occurs in the district capital. With a dataset coded at the city level, one would therefore struggle to find any location effects if hospitals were always built in the same city. Within a city, however, there are many neighborhoods, and data geocoded at the village level are capable of detecting shifts toward certain neighborhoods over others.

## 4 Welfare

This section quantifies welfare effects with a model of spatial demand for healthcare facilities. I specify the individual’s choice problem and discuss how I estimate the model.

### 4.1 Model

An individual seeks care at a facility type  $f \in \mathcal{F}$ , where the choice set  $\mathcal{F}$  contains public hospitals, private hospitals, clinics, and subclinics. I assume that individuals consider the closest facility of each type because the data record usage by facility type, but not by specific facility within a type. The utility of facility type  $f$  for individuals living in village  $v$  at time  $t$  is

$$\text{utility}_{fvt} = \underbrace{x_{fvt}\beta_f + p_{fvt}\alpha + \xi_{ft} + \delta_v + \delta_t}_{\equiv V_{fvt}} + \varepsilon_{fvt}.$$

Individuals consider facility characteristics  $x_{fvt} = [\text{distance}_{fvt}, \text{congestion}_{fvt}]$ , and I allow for preferences over these characteristics to vary freely by facility type. Individuals also consider facility prices  $p_{fvt}$ , which I observe, and facility quality  $\xi_{ft}$ , which I do not observe. This specification restricts facilities to be homogeneous within a given facility type and year. Parameters  $\delta_v$  and  $\delta_t$  are village and time fixed effects, and  $\varepsilon_{fvt}$  are logit errors. I normalize the utility of the outside option to zero. The logit inversion implies

$$\ln(s_{fvt}) - \ln(s_{0vt}) = x_{fvt}\beta_f + p_{fvt}\alpha + \xi_{ft} + \delta_v + \delta_t + \epsilon_{fvt}, \quad (1)$$

where  $s_{fvt}$  denotes market shares by facility type, village, and year. I also consider a specification with a population-density interaction term  $(x_{fvt} \cdot \text{popden}_{vt})\beta'_f$  that allows preferences to vary between rural and urban villages. The interacted specification also allows the facility quality  $\xi_{ft}$  to vary freely between villages with above- and below-mean density.

### 4.2 Estimation

I calculate market shares from individual-level data on the number of visits to each facility type in the last month. I focus on the period from 1993 to 2002 because the SUSENAS data do not distinguish between clinic and subclinic visits in other years. For each sick individual, which I define as those reporting at least one health concern, I classify the individual as having visited either a private hospital, public hospital, clinic, or subclinic. I do not distinguish between a single visit and multiple visits to a given facility type. For individuals who visited multiple facility types, I code them based on the most expensive facility type they visited (in order, private hospital, public hospital, clinic, and subclinic). The alternative is to compute market shares by visits instead of by individuals as I do here, but this alternative approach would treat visits as independent. Since

I focus on sick individuals, the outside option is choosing not to visit a facility despite being sick. Lastly, a practical concern is that the logit inversion is infeasible when market shares are zero or one, so I use inverse-distance weighting to smooth the market shares that I estimate from the data.

For facility characteristics, I measure distances as Euclidean distances between village centroids. Congestion is a function of how many individuals use a given facility. I proxy for this measure with the number of individuals for whom a given facility is the closest facility of its type. In the language of [Donaldson and Hornbeck \(2016\)](#), distance $_{fvt}$  captures the direct effects of facility construction, while congestion $_{fvt}$  captures the indirect effects. That is, a new facility directly increases usage in nearby villages by decreasing travel distance, and also indirectly increases usage in faraway villages as movement to the new facility decongests other facilities.

I construct prices from household-level data on health spending. For each village and year, the following regression of spending on visits yields the average amount of money spent on each facility type (or the outside option).

$$\text{spending}_{hvt} = \mu_{0vt} + \mu_{fvt} \sum_{f \in \mathcal{F}} \text{visits}_{fhvt} + u_{hvt}$$

To ensure prices are smooth over space, I run the regression for each village using data from all villages, using inverse-distance weights  $(1 + \text{distance}(v, v'))^{-2}$  that weigh nearby villages more heavily. I obtain variation in prices over time by repeating this procedure for each year in the data.

### 4.3 Estimates

I estimate specification [1](#) by OLS and present the resulting estimates in [table 2](#). Facility distance and congestion correspond to lower usage, and demand for clinics and subclinics is more elastic than demand for public hospitals. Demand for private hospitals is least elastic. In the appendix, [figure A1](#) plots the raw correlation between usage and distance, and [figure A2](#) shows the lack of pretrends. [Table A1](#) shows relatively little heterogeneity by population density. Urban areas are more elastic in terms of distance and less elastic in terms of congestion, although only the congestion differences are statistically significant. Thus, in the next stage of analysis I focus on the non-interacted demand estimates of [table 2](#).

One concern is that the price coefficient is biased because the price data are constructed and therefore subject to potentially substantial measurement error. In practice, in the analysis that follows I only use the price coefficient to denominate the demand system’s utility predictions in dollar terms. Alternatively, I can use the distance coefficients to denominate utility in terms of kilometers saved. Indeed, comparing the magnitudes of the price and distance estimates suggests that patients value one kilometer saved at approximately \$4, a figure that is perhaps high (as consistent with attenuation bias in the price coefficient) but still within range of a sensible prior.

**Table 2:** Usage by facility distance and congestion

	Estimate	Standard Error
Distance, public hospital	-2.295***	(0.0561)
Distance, private hospital	-1.127***	(0.0323)
Distance, clinic	-2.456***	(0.233)
Distance, subclinic	-2.404***	(0.442)
Congestion, public hospital	-0.0247***	(0.00298)
Congestion, private hospital	-0.0172***	(0.00127)
Congestion, clinic	-0.381***	(0.0389)
Congestion, subclinic	-0.589***	(0.0246)
Price	-0.537***	(0.152)
Village FE		x
Facility type-year FE		x
Observations		202,668

Each column is a single conditional multinomial logit regression with village and facility type-year fixed effects. The unit of observation is a village-year-facility type, where the set of facility types represents a village's choice set in a given year. The outcome is usage by facility type, as recorded in the SUSENAS data. Distance is to the closest facility of each type and is measured in units of 100 km. Congestion of the closest facility is the number of people for whom this facility is the closest of its type. This variable is measured in units of 100,000 people. Price is measured in units of \$100 (in year 2000 USD). Additional controls include population and ruralness. Standard errors are clustered by village. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

#### 4.4 Extensions

Future work will focus on the following extensions to this simple model of demand. First, I can avoid the issue of price endogeneity by using travel time in place of distance and price, and using a valuation of time to monetize the system. Since the SUSENAS data record incomes, the valuation of time could be allowed to vary by individual. Second, I follow the spatial demand literature by taking distance as exogenous, but a more sophisticated approach could instrument for distances using the distance-minimizing allocation, which is similar to the least-cost-path approach commonly adopted in work on transportation networks. Third, I can allow distance and congestion elasticities to vary with individual characteristics in a random-coefficients framework. This added flexibility would allow welfare effect to differ, for example, between rich and poor villages. Fourth, I restrict heterogeneity among facilities by imposing homogeneity within a facility type and year, but I can relax this restriction by allowing for heterogeneity by region or with grouped fixed effects. The data also contain some observables on hospitals and clinics that I can control for directly. Fifth, I restrict village heterogeneity to a village fixed effect and a time fixed effect, but I can also allow for differential time trends by village. Sixth, clinics provide referrals to hospitals for serious conditions, and I can capture this interaction by allowing complementarities across facility types.

## 5 Misallocation

The estimated demand system allows me to quantify the surplus generated by any given allocation of healthcare facilities. This section proposes a measure of misallocation that compares the surplus generated by observed allocations to the maximum achievable under the same budget constraint.

### 5.1 Consumer surplus

Consumer surplus is a function of the compensating variation associated with a given facility placement. Defining notation, let policy  $a$  specify where facilities are built over time. Placement  $a^t$  is the allocation of facilities resulting from policy  $a$  as of time  $t$ . As in [McFadden \(1981\)](#), I use the estimated price elasticity  $\hat{\alpha}$  to calculate the change in prices needed to compensate for some change in facility characteristics. Defined in relation to benchmark  $\underline{a}^t$ , the compensating variation associated with moving to placement  $a^t$  is

$$CV_{vt}(a^t, \underline{a}^t) = \frac{1}{\hat{\alpha}} \left[ \ln \left( \sum_{f \in \mathcal{F}} \exp(V_{fvt}(a^t)) \right) - \ln \left( \sum_{f \in \mathcal{F}} \exp(\underline{V}_{fvt}(\underline{a}^t)) \right) \right]. \quad (2)$$

$V_{fvt}$  are fitted values from specification [1](#), and  $\underline{V}_{fvt}$  can be calculated as

$$\underline{V}_{fvt} = V_{fvt} + (\underline{x}_{fvt} - x_{fvt})\beta_f,$$

where  $x_{fvt}$  and  $\underline{x}_{fvt}$  are distance and congestion under placements  $a^t$  and  $\underline{a}^t$ , respectively, assuming that a change in facility placement does not impact pricing.<sup>3</sup> The consumer surplus gains arising from placement  $a^t$  are therefore

$$\Delta \text{surplus}_t(a^t, \underline{a}^t) = \sum_{v \in \mathcal{V}} \text{population}_{vt} \cdot CV_{vt}(a^t, \underline{a}^t),$$

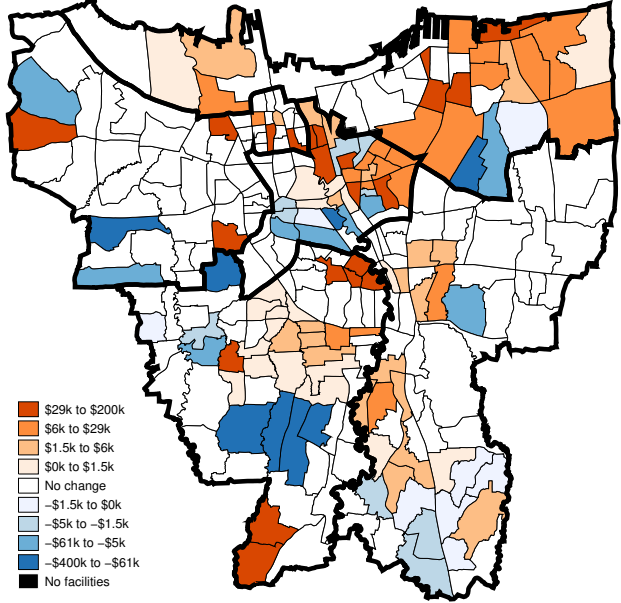
and gains from a policy over time are

$$\Delta \text{surplus}(a, \underline{a}) = \sum_t \beta^{t-1} \Delta \text{surplus}_t(a^t, \underline{a}^t).$$

---

<sup>3</sup> To obtain values for  $V_{fvt}$ , I extrapolate from the usage sample, which covers a subset of villages from 1993 to 2002, to other villages and other years. As such, calculating the fitted values requires some imputation of village fixed effects  $\delta_v$  and year fixed effects  $\delta_t$ . For out-of-sample villages in districts with at least 10 in-sample villages, I impute the village fixed effect as the distance-weighted average of the same-district, in-sample fixed effects. I use the same weighting scheme as when I smooth the choice probabilities, namely  $(1 + \text{distance}(v, v'))^{-2}$ . For the small proportion – about 0.5% – of out-of-sample villages in districts without at least 10 in-sample villages, I calculate the distance-weighted average of all in-sample fixed effects. For year fixed effects, I use the year 1993 fixed effect for pre-1993 years and the year 2000 fixed effect for post-2000 years.

**Figure 2:** Consumer surplus under surplus-maximizing vs. actual placement, Jakarta



This figure plots compares villages’ consumer surpluses under the surplus-maximizing and actual placements. In Jakarta, the administrative unit of a “village” can be thought of as a neighborhood. Villages in orange gain under the surplus-maximizing allocation, while villages in blue lose. That there are more orange villages reflects that total surplus under the surplus-maximizing allocation is larger than that under the actual allocation.

Note that the use of benchmark  $\underline{a}$  is necessary because, as is typical of discrete choice models, consumer surplus is identified in changes but not in levels.

### 5.2 Surplus-maximizing allocation

Consider the allocation  $\bar{a}$  that maximizes consumer surplus gains for society subject to the same construction budget over time as in observed placement  $a$ .

$$\bar{a} \equiv \arg \max_b \Delta \text{surplus}(b, \underline{a}) \quad \text{s.t.} \quad \sum_{v \in \mathcal{V}} n_v(b^t) = \sum_{v \in \mathcal{V}} n_v(a^t) \quad \forall t, \tag{3}$$

where  $n_v(\cdot)$  is the number of facilities in village  $v$  resulting from a given placement. Figure 2 compares the surplus-maximizing and actual allocations for Jakarta. Here, the problem is of relatively low dimension, and I can solve for the surplus-maximizing allocation exactly. More generally, however, obtaining this allocation requires solving a high-dimensional combinatorial optimization problem that is NP-hard. The solution space expands exponentially in the number of villages, the number of facilities, and the number of time periods. At the national level, the problem is of intractably high dimension.

I make progress in two ways. First, I solve the problem locally at the district level, and I aggregate these sub-solutions to approximate the full solution. This local approach will differ from

the full solution when cross-district spillover effects are large, so to mitigate this bias I account for these out-of-district spillover effects by specifying the local subproblems over a given district and a surrounding buffer zone. I choose the buffer zone to include any out-of-district village that would potentially be impacted by construction in the district of interest, whether it be through distance or congestion reductions. In this way, I account for any interaction between new construction in a district and the existing facilities in neighboring districts. However, the interaction between new in-district construction and new neighboring-district construction may still generate some bias.

Second, I apply simulated annealing to solve the problem heuristically. Simulated annealing is a stochastic, global optimization algorithm that is a variant of Metropolis-Hastings. I specify a starting temperature high enough to accept at least 95% of proposal solutions. The algorithm stops when the best candidate solution has not been surpassed for 500 iterations, and when the acceptance ratio is no higher than 5%. For robustness, I repeat the algorithm from 5 random starts and take the best solution. This algorithm breaks the curse of dimensionality, and I find that it performs consistently over the multiple starts in my setting.

This class of stochastic algorithms is not commonly used in estimating structural parameters because they do not deliver the exact solutions that gradient-based local optimizers do. The concern is that candidate solutions can differ substantially from the true solution but nonetheless be identified as optimal because they achieve similar objective values. For example, suppose a function is maximized by true parameters  $\theta = (1, 10)$  with objective value  $f(\theta) = 100$ , and a stochastic algorithm delivers an estimate  $\hat{\theta} = (10, 1)$  with objective value  $f(\hat{\theta}) = 99.9$ . This discrepancy can be a problem if the goal is to interpret the parameter estimates themselves, but it is not a problem if the quantity of interest is the objective value itself. Indeed, stochastic algorithms are more common in machine learning applications where the focus is on prediction and not on the interpretation of coefficients. Similarly, my measure of misallocation is only a function of the (approximately) maximized objective value  $\Delta\text{surplus}(\bar{a}, \underline{a})$  and not the maximizing allocation.

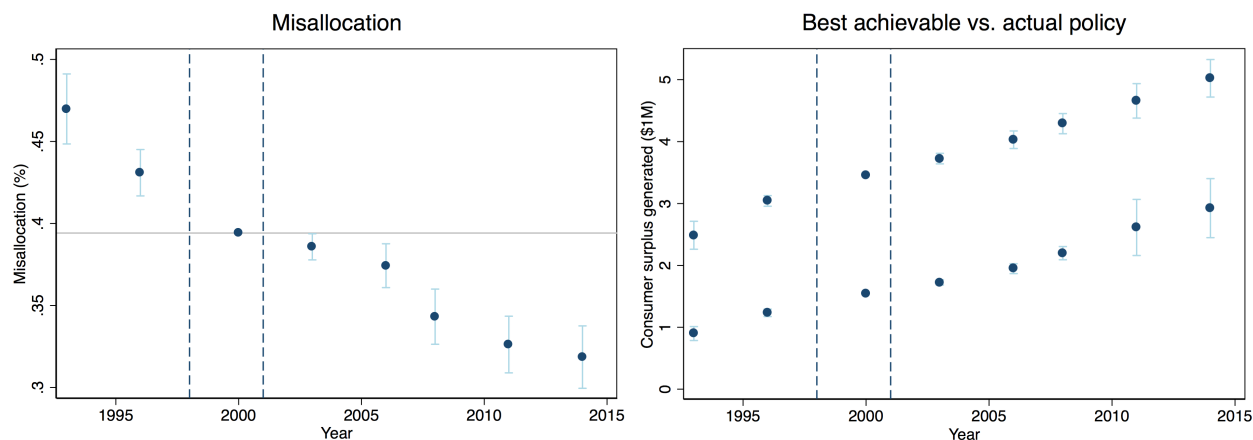
### 5.3 Measuring misallocation

In particular, my measure of misallocation compares the consumer surplus generated by the observed allocation to that of the surplus-maximizing allocation.

$$\text{misallocation}(a, \underline{a}) = 1 - \frac{\Delta\text{surplus}(a, \underline{a})}{\Delta\text{surplus}(\bar{a}, \underline{a})} \quad (4)$$

This measure of misallocation is in percentage terms and is zero when the observed placement coincides with the surplus-maximizing placement. For each district, I can evaluate the time profile of misallocation by computing this measure for each three-year period  $t$ . Note that although I compute consumer surplus gains by district, the optimal placement is defined as that which is optimal for society – namely, the placement that fully internalizes the spillovers to non-constituents

**Figure 3:** Misallocation over time



Misallocation is defined as one minus the proportion of the maximum achievable consumer surplus gain that is achieved by the observed placement. It is zero when the actual placement coincides with the surplus-maximizing placement. The plot on the right shows the consumer surplus generated by the surplus-maximizing placements (top line) and the actual placements (bottom line) over time, controlling for district fixed effects. For each period, the benchmark placement  $\underline{a}$  is the facility placement in 1990. The vertical dashed lines mark the period of reform, and the error bars are 95% confidence intervals.

living outside of a district’s borders.

Figure 3 shows how, on average, these values evolve over time. The left panel shows that misallocation levels are lower after the reform period. The data cover a relatively short period of time before the reform, so future work will focus on collecting additional data to extend the length of the panel. The right panel shows the estimated consumer surplus gains that I use to estimate misallocation.<sup>4</sup> The dashed line shows the maximum consumer surplus gains achievable under the budget constraint, while the solid line shows those achieved by the observed placements. For all years I take facility placements in 1990 as the benchmark placement, and as such these figures present the time path of the stock of misallocation. In the appendix, figures A3 and A4 show the contribution of each facility type to these trends. In these calculations, I optimize over each facility type in turn while holding the other facility types fixed. Clinics contribute the most to overall misallocation, while public hospitals contribute the least.

## 5.4 Challenges in interpreting misallocation

Interpreting this measure of misallocation requires some caution. In general, any quantification of misallocation is dependent on a economic model, which either delivers an optimal allocation directly or shows that marginal products are unequalized across agents (and therefore that the current allocation is suboptimal). This dependence on a model means that model misspecification will spuriously suggest misallocation. For example, building on the seminal quantification of firm-

<sup>4</sup> To see why the misallocation figures are not exactly determined by dividing the solid line by the dashed line, note that misallocation first divides the actual by achievable gains, then takes the average, while dividing the solid and dashed lines would take the average of the actual gains and the average of the achievable gains, then divide them.

level factor misallocation in [Hsieh and Klenow \(2009\)](#), [Bartelsman et al. \(2013\)](#) and [Asker et al. \(2014\)](#) discuss how a model with adjustment frictions can explain a significant portion of the productivity dispersion observed among firms.

Applied to misallocation in infrastructure investment, these insights suggest that the level of misallocation I find in figure 3 can change depending on how I specify the benchmark objective function. That is, the question of what agents are maximizing is central to measuring misallocation. I find that facilities are about 40% misallocated relative to an allocation that maximizes consumer surplus. But another objective function may place weight on equity over space, and yet another objective function may value gains to the wealthy, who have higher willingness to pay. Measured against these benchmarks, misallocation may be higher or lower than 40%. As such, while papers like [Fajgelbaum and Schaal \(2017\)](#) and [Balboni \(2019\)](#) make significant progress in quantifying misallocation under an assumed government objective function, these exercises are not geared toward determining whether the misallocation they measure is a result of irrational behavior or simply a government objective function that differs from the ones used in their models. Instead, it is the changes over time that deserve emphasis, as these changes are all measured relative to the same benchmark.

A potential confounder of the changes over time seen in figure 3 is misspecification of the agent’s dynamic horizon – another example of model misspecification. The concern is that a forward-looking government will make facility placement decisions accounting for future placements, such that placements that look suboptimal today are in fact optimal given placements in subsequent periods. Any misspecification of how far the government looks into the future will therefore be spuriously attributed to misallocation. Appendix figure A7 further illustrates this concern with a simple example. The revealed-preference approach I take in the following section remains subject to this concern, so planned work will check robustness across different assumptions on how forward-looking the government is.

Another potential confounder is mismeasurement. [Bils et al. \(2018\)](#) and [Rotemberg and White \(2017\)](#) document how mismeasurement of data can be attributed to misallocation across firms. In this setting, improvements in survey technology over time will generate a decrease in measured misallocation over time simply because optimal allocations in earlier periods are incorrectly recorded as suboptimal allocations. It is difficult to credibly rule out this possibility for the purposes of figure 3, but the revealed-preference approach that I take in the following section is robust to this concern.

## 6 Determinants of Misallocation

The previous section quantifies misallocation, but it does not explain why misallocation exists. This section models the facility placement problem as a dynamic discrete choice problem and estimates the government’s objective function by revealed preference.

## 6.1 The facility allocation problem

I consider the dynamic facility location problem at the level of the district. At time  $t = 0$ , the district mayor chooses a policy  $a$  that specifies where to build facilities in every future period. The mayor chooses construction locations, but is subject to a budget constraint that specifies the number of facilities to be constructed in each period. Recall that  $a^t$  denotes the placement resulting from policy  $a$  as of time  $t$ , and  $n_v(a^t)$  the number of facilities in village  $v$  given placement  $a^t$ . The mayor's objective function is a discounted sum of payoffs over time.

$$\pi(a) = \sum_{t=1}^{\infty} \beta^{t-1} \left( S(a^t) + X(a^t) + \xi(a^t) \right) \quad (5)$$

It nests maximization of consumer surplus  $S$ , which I used to define measure misallocation in section 5, but it can also include other preferences – both observed  $X$  and unobserved  $\xi$ .

The mayor considers the consumer surplus generated by a given placement. I set social surplus as the numeraire, and I distinguish between surplus for in-district and out-of-district villages.

$$S(a^t) = \sum_{v \in \mathcal{V}_{\text{in}}} \text{surplus}_{vt}(a^t; \omega) + \underbrace{\sum_{v \in \mathcal{V}_{\text{out}}} \tau^S \text{surplus}_{vt}(a^t; \omega)}_{\text{internalized spillovers}} \quad (6)$$

Consumer surplus is a function of the demand system estimated in section 4, and as such depends on demand parameters  $\omega$ . Patients can travel, so villages can benefit from new facilities even if they do not receive the facilities themselves. The parameter  $\tau^S \in [0, 1]$  captures the extent to which a mayor internalizes spillover benefits for out-of-district villages. At one extreme, a mayor that acts as the social planner does internalizes spillovers fully, such that  $\tau^S = 1$ . At the other extreme, a mayor focusing on in-district villages will fully discount spillover benefits to out-of-district villages, which provide neither votes nor tax revenue, such that  $\tau^S = 0$ . A richer model could allow for more nuanced internalization of spillovers, such as among neighboring districts headed by mayors of the same political party.

I consider the extent to which observables can rationalize deviations from surplus maximization. In particular, the mayor's decision may also be subject to favoritism.

$$X(a^t) = \sum_{v \in \mathcal{V}_{\text{in}}} \sum_f n_{fv}(a^t) \cdot D_v \left( \underbrace{\tau_f^P \text{patronage}_v + \tau_f^G \text{golkar}_v + \tau_f^E \text{ethnicity}_v}_{\text{favoritism}} \right) \quad (7)$$

Favoritism includes the military patronage network and a village's underlying support for the *Golkar* party, both of which potentially give rise to distorted allocations under the Suharto regime. I also include ethnic composition, which is another source of favoritism in this context. I sum over in-district villages because the mayor considers only in-district placements, and I allow for

differential impacts by facility type  $f \in \{\text{hosp}, \text{clin}, \text{sub}\}$  that enter linearly in the number of facilities constructed in a given village. To capture the spatial nature of the problem, building in a village involves favoritism associated with both that village and the surrounding villages, down-weighting by distance.

$$D_v(X_v, X_{-v}) = \sum_{v' \in \mathcal{V}_{\text{in}}} W_{v'} X_{v'} / \sum_{v' \in \mathcal{V}_{\text{in}}} W_{v'}, \quad W_{v'} = (1 + \eta \cdot \text{distance}(v, v'))^{-2},$$

where I suppress  $X_{-v}$  in the objective function. I assume a single weighting function parameterized by  $\eta$ , but this function can in principle differ for each variable.

I further allow for unobservable choice factors. I consider a decomposition of village, time, and village-time factors, and I assume that these factors enter linearly as the observables do.

$$\xi(a^t) = \sum_{v \in \mathcal{V}_{\text{in}}} \sum_f n_{fv}(a^t) \cdot D_v(\lambda_{fv} + \psi_{ft} + \varepsilon_{fvt}), \quad (8)$$

where  $\mathbb{E}(\varepsilon_{fvt}) = 0$  by construction. In general, the unobservables reflect underlying heterogeneity in preferences over villages. They absorb unobserved sources of favoritism, but they also arise from mismeasured consumer surplus. For example, demand shocks arising from disease outbreaks and misspecification of the demand system will both enter here. Unobservables also absorb variation in costs unaccounted for by a budget constraint based on the number of facilities. For example, high land and labor costs in one region make building a hospital there more expensive than building in another region. Accommodating unobservables is therefore critical in this setting where ignoring unobservables induces selection bias in estimation, and furthermore where the econometrician is unlikely to observe all choice factors underlying the observed placement decisions.

Finally, I set consumer surplus as the numeraire, such that preferences over other factors are denominated in dollars of consumer surplus. I omit the fixed costs of each facility type without loss of generality because these costs are fixed across placement choices.

## 6.2 Additional assumptions for estimation

I assume a fixed relationship across parameters by facility type. Doing so avoids the need to estimate the full set of parameters separately for each facility type.

$$\theta_{\text{hosp}} = \theta, \quad \theta_{\text{clin}} = \gamma_{\text{clin}} \cdot \theta_{\text{hosp}}, \quad \theta_{\text{sub}} = \gamma_{\text{sub}} \cdot \theta_{\text{hosp}},$$

for  $\theta \in \{\tau^P, \tau^G, \tau^E, \{\lambda_v\}\}$ . I also set the annual discount factor to  $\beta = 0.95$ , as it is generically unidentified ([Magnac and Thesmar 2002](#)).

I further assume that the government has perfect foresight over population growth, which evolves exogenously. Given future populations, the government can evaluate the consumer surplus

generated over time by any given placement decision. I do not need to make assumptions on the government’s expectations over time unobservables  $\psi_{ft}$  or village-time unobservables  $\varepsilon_{fvt}$ . Furthermore, new facilities in Indonesia are allocated using population-based rules, and so exogenous population growth implies that facility budgets also evolve exogenously over time. Otherwise, if budget allocations tomorrow depended on placement choices today, then the payoff associated with a given placement would need to account for these future budget effects above and beyond the objective function described above.

I also make several assumptions for computational feasibility. I assume that the government is unable to reallocate its budget across periods. It therefore spends the entirety of its budget in each period, simplifying the choice space by restricting it to the number of facilities observed in the data in each period.

$$\mathcal{A} = \left\{ b \mid \sum_{v \in \mathcal{V}_{in}} n_{fv}(b^t) = \sum_{v \in \mathcal{V}_{in}} n_{fv}(a^t), \forall f, t \right\}. \quad (9)$$

Next, I assume homogeneity across newer and older facilities, which focuses the analysis on the number of facilities of each type in each village, and not on the age composition of facilities. Finally, I assume away strategic responses placement decisions. In reality, the private sector may respond to the placement of public facilities, although private hospitals serve a different market than public hospitals do (i.e., the very wealthy). Under decentralization, the responses of other districts may also be important.

### 6.3 Estimation

I estimate government preference parameters  $\tau$  with the revealed-preference approach described in [Pakes \(2010\)](#) and [Pakes et al. \(2015\)](#). Let  $a$  denote the actual placement policy observed in the data. For true preference parameters  $\tau^0$ ,

$$\pi(a; \tau, \omega^0) \geq \pi(b; \tau, \omega^0) \quad \forall b \in \mathcal{A} \quad (10)$$

for alternative policies  $b$  given demand parameters  $\omega^0$ . That is, at the true parameter values, the chosen option at least weakly dominates all other options. I operationalize this insight by choosing a set of alternatives  $b$ , constructing the associated revealed-preference inequality for each alternative, and ruling out candidate values of parameters  $\tau$  that violate one or more of these inequalities. This approach achieves dimension reduction by evaluating only a subset of the possible alternatives. By contrast, a nested fixed-point approach picks a candidate value for parameter  $\tau$  and evaluates every policy to find the optimal policy given  $\tau$ ; it then chooses the value of  $\tau$  that produces the predicted optimal policy most similar to the observed policy. The latter approach is infeasible in this context given the computational complexity of computing the optimal policy.

Expanding inequality 10 and rearranging to show the difference in payoffs between actual policy  $a$  and alternative policy  $b$ , the revealed-preference inequality becomes

$$\sum_{t=1}^{\infty} \beta^{t-1} \left( \Delta SX(a^t, b^t; \tau) + \Delta \xi(a^t, b^t; \lambda, \psi, \varepsilon) \right) \geq 0, \quad (11)$$

where I suppress the demand parameters  $\omega^0$ , and I define

$$\Delta SX(a^t, b^t; \tau) \equiv S(a^t; \tau) - S(b^t; \tau) + X(a^t; \tau) - X(b^t; \tau).$$

The challenge is in constructing a sample analogue to inequality 11 given (1) dynamics and (2) unobservables. To address these issues, I use techniques described in applications [Holmes \(2011\)](#) and [Ho and Pakes \(2014\)](#). For simplicity, in discussing identification I assume only one facility type and suppress all  $f$  subscripts. Estimation uses information from all facilities types – hospitals, clinics, and subclinics – jointly, and the extension is straightforward.

The first challenge arises in capturing the dynamic effects of a given policy. A new facility in period  $t$  impacts payoffs both in period  $t$  and in all future periods, but calculating future payoffs requires knowing future placements. For example, placing a facility in village  $v$  may have large benefits today, but these marginal benefits will decrease if the neighboring village is slated to receive a facility tomorrow. One approach is to solve within a limited lookahead window as in [Zheng \(2016\)](#), but this approach sacrifices the last periods of the panel data and only captures dynamics within the lookahead window.

Instead, I sidestep these issues with “pairwise resequeing” as in [Holmes \(2011\)](#). I select alternatives that swap the construction order of two facilities in the actual policy. For example, for actual policy  $a = (v_1, \{v_2, v_3\}, v_4, \dots)$ , a pairwise resequenced alternative is  $b = (v_2, \{v_1, v_3\}, v_4, \dots)$ . After the second period, these policies result in the same number of facilities in every village. Formally, the set  $\mathcal{S} \subset \mathcal{A}$  of “swapped” alternatives to actual policy  $a$  are such that

$$n_v(a^t) - n_v(b^t) = \begin{cases} 0 & \text{for } t \notin s(b), v \in \mathcal{V} \\ -1 & \text{for } t \in s(b), v = w_1(b) \\ 1 & \text{for } t \in s(b), v = w_2(b) \\ 0 & \text{for } t \in s(b), v \in \mathcal{V} \setminus \{w_1(b), w_2(b)\} \end{cases} \quad (12)$$

where  $s(b) = \{s_1(b), s_1(b) + 1, \dots, s_2(b) - 1\}$ ,  $s_1(b)$  is the earlier period involved in the swap,  $s_2(b)$  is the later period,  $w_1(b)$  is the village that receives its facility earlier in the swap, and  $w_2(b)$  is the village that receives it later.<sup>5</sup> For swapped alternatives  $b \in \mathcal{S}$ , applying the first line of condition

<sup>5</sup> If facilities are constructed at the beginning of the period, then in period  $s_2$  both of the facilities involved in the swap have been built. As such, the equality holds in period  $s_2$ . If facilities are instead constructed at the end of the period, then the equality holds at all  $t \notin \{s_1(b) - 1, s_1(b), \dots, s_2(b)\}$ .

12 to inequality 11 gives

$$\sum_{t=s_1(b)}^{s_2(b)-1} \beta^{t-1} \left( \Delta SX(a^t, b^t; \tau) + \Delta \xi(a^t, b^t; \lambda, \psi, \varepsilon) \right) \geq 0, \quad (13)$$

since  $n_v(a^t) = n_v(b^t)$  for all villages  $v$  outside of the swap-relevant periods. That is, the per-period payoffs of actual policy  $a$  and swapped alternative  $b$  are identical in the periods before and after the swap, thereby eliminating dynamics beyond period  $s_2(b)$ . Figure A5, appended, presents the intuition visually.

The second challenge lies in the unobservable terms of inequality 13. Ignoring these terms leads to selection bias: given that policy  $a$  was actually chosen, the unobserved payoff of  $a$  relative to unchosen policy  $b$  is unlikely to be zero in expectation. To proceed, I apply the assumed functional form on the unobservables as described in equation 8.

$$\Delta \xi(a^t, b^t; \lambda, \psi, \varepsilon) = \sum_{v \in \mathcal{V}_{\text{in}}} (n_v(a^t) - n_v(b^t)) \cdot D_v \left( \lambda_v + \psi_t + \varepsilon_{vt} \right). \quad (14)$$

Substituting this expression, applying the rest of condition 12, and applying the budget-spending assumption of equation 9, inequality 13 simplifies to

$$\sum_{t=s_1(b)}^{s_2(b)-1} \beta^{t-1} \left( \Delta SX(a^t, b^t; \tau) + \Delta D(\lambda) + \Delta D(\varepsilon_t) \right) \geq 0, \quad (15)$$

where I define

$$\Delta D(X) \equiv D_{w_2(b)}(X_{w_2(b)}, X_{-w_2(b)}) - D_{w_1(b)}(X_{w_1(b)}, X_{-w_1(b)}).$$

I address the  $\Delta D(\lambda)$  terms by estimating the village fixed effects  $\lambda_v$  directly, and the  $\Delta D(\varepsilon_t)$  terms with an aggregation step. Time effects  $\psi_t$  cancel because the number of facilities is held constant.

I estimate the model separately on the pre- and post-reform data. As such, the parameter values – including the village fixed effects – are allowed to differ across the two periods. To summarize, the parameters to estimate are  $\tau \equiv \{\tau^S, \tau^P, \tau^G, \tau^E, \{\lambda_v\}, \eta, \gamma_{\text{clin}}, \gamma_{\text{sub}}\}$ . First, I identify  $K$  valid swaps and form the left-hand side of inequality 15 for each. The computationally intensive part of the estimation procedure is in calculating these inequality values for all  $K$  swaps, although the process is readily parallelizable.

$$I_k(\tau) = \sum_{t=s_1(b)}^{s_2(b)-1} \beta^{t-1} \left( \Delta SX(a^t, b^t; \tau) + \Delta D(\lambda) \right) \quad (16)$$

This expression contains only predicted values, observed values, and parameters to be estimated.

Second, I form aggregated moments by averaging these values in groups  $\mathcal{G}_{g+}$  and  $\mathcal{G}_{g-}$  for  $g \in \{\tau^S, \tau^P, \tau^G, \tau^E, \{\lambda_v\}\}$ . That is, in order to identify upper and lower bounds for each parameter of interest, I group swap values based on whether the swap increases or decreases the variables associated with the parameter.<sup>6</sup>

$$M_g(\tau) = \frac{\sum_{k=1}^K \left( \mathbb{1}(k \in \mathcal{G}_g) \cdot I_k(\tau) \right)}{\sum_{k=1}^K \left( \mathbb{1}(k \in \mathcal{G}_g) \right)} \quad (17)$$

Since village-time shocks are mean-zero by construction, these moments coincide with the aggregated revealed-preference inequalities implied by the model.<sup>7</sup> Third, by revealed preference  $M_g(\tau) \geq 0$  for all groups  $g$  at the true preference parameters  $\tau^0$ . I therefore choose parameters  $\tau$  that minimize violations of the moment inequalities.

$$\hat{\tau} = \arg \min_{\tau} \left\{ \frac{1}{G} \sum_{g=1}^G \left( \min\{M_g(\tau), 0\} \right)^2 \right\} \quad (18)$$

Because the moments can be formed offline, at this point the optimization problem is computationally light.

## 6.4 Estimates

Table 3 shows how many valid swaps can be identified based on the actual policy  $a$  observed in the data. Swaps are constrained to be within districts both before and after the reforms.<sup>8</sup> Where redistricting occurs, I omit swaps that span redistricted borders. Table 4 shows estimates of government preference parameters  $\tau$ . Estimation yields point estimates, which occur when not all inequalities can be simultaneously satisfied. Indeed, point identification is common in cases with a large number of moments.

District governments are only somewhat less likely to internalize spillovers increases in the post-

<sup>6</sup> Grouping inequalities with different types of identifying variation may eliminate it. For example, combining an inequality with an alternative that increases costs and one with an alternative that decreases them will eliminate the identifying variation.

<sup>7</sup> Since  $\mathbb{E}(\varepsilon_{vt}) = 0$ , the unobserved  $\varepsilon_{vt}$  terms can be averaged out. The selection issue is that these shocks, which are unconditionally mean-zero, may not be mean-zero after conditioning on the chosen policy  $a$ . In this case, however, the assumed linearity of the unobservable component terms delivers an inequality that is additive in the  $\lambda_v$  and  $\varepsilon_{vt}$  terms no matter the chosen policy  $a$ . That is, regardless of where policy  $a$  places facilities, the alternative policy involving swap villages  $w_1(b)$  and  $w_2(b)$  yields an inequality containing the same  $\varepsilon_{w_2(b),t}$  and  $\varepsilon_{w_1(b),t}$  terms. Thus, the inequality need not condition on chosen policy  $a$ , and the unconditional average is sufficient for addressing the  $\varepsilon_{vt}$  terms.

<sup>8</sup> Swaps compare the actual policy  $a$  to some alternative policy  $b$  within the decision-maker's choice set. If alternative  $b$  is not within the decision-maker's consideration set, then the resulting revealed-preference inequality will not necessarily hold. After decentralization, district governments choose facility placement within their districts. Before decentralization, in principle the central government chooses over the full set of villages. The full set of swaps is therefore very large. In practice, however, population-based rules govern the allocation of facilities to districts, so placement choices may still be constrained to be within districts. As such, I consider only the subset of swaps that occur within districts.

**Table 3:** Number of potential swaps for estimation

	Data	Swaps		
	PODES	Hospitals	Clinics	Subclinics
Pre-reform	1990, 93, 96, 2000	1,089	39,522	420,755
Post-reform	2003, 06, 08, 11, 14	5,013	45,267	299,056

Swaps are pairs of facility allocations – one observed in the data and one hypothetical – in which the placement order for two facilities has been swapped. In other words, a swapped allocation is a pairwise resequencing of an observed allocation. Swaps are all within districts and are restricted to be either within the pre-reform period or within the post-reform period. I form swaps separately for each facility type. The table shows the total number of swaps available, but where the total is large I sample a subset of swaps for use in estimation.

**Table 4:** Government preference parameter estimates

	Pre-reform	Post-reform
Internalized spillovers ( $\tau^S$ )	0.79	0.66
Patronage ( $\tau^P$ )	3.41	0.85
<i>Golkar</i> ( $\tau^G$ )	2.85	1.03
Ethnicity ( $\tau^E$ )	0.04	0.07
“R-squared”	0.64	0.70

Each row in the upper panel corresponds to an estimated parameter of the government objective function. Internalized spillovers indicates the extent to which a district government values spillover benefits to non-constituents. A value of one represents full internalization and is what the social planner would do. Patronage refers to villages within the Suharto patronage network, as proxied by military presence. *Golkar* is whether a village is an historical supporter of Suharto’s political party, as proxied by vote shares in Suharto-era elections. The R-squared shown in the bottom panel indicates the proportion of revealed-preference inequalities that can be satisfied under the estimated parameters.

reform period. Patronage and *Golkar* support play a larger role in the pre-reform period, while ethnicity does not play a major role in either period. Figure A6, appended, maps the estimated village preferences  $\lambda_v$  for Jakarta in the post-decentralization period. The majority of villages have a large, negative values of  $\lambda_v$ , which reflects that they do not receive many facilities despite potentially large welfare benefits.

I examine the goodness of fit of the model by evaluating the set of inequality values given by equation 16 at the estimated parameters, and I report the percentage that are positive as an “R-squared” in table 4. The percentage that are positive reflects the degree to which the model and the estimated parameters explain the observed placement, at least relative to its one-step deviations.

## 6.5 Discussion

I assume village populations are fixed in calculating counterfactual consumer surplus under alternative placements. Estimates will be biased if village populations respond endogenously to changes in infrastructure. In particular, surplus gains to alternative placements will be understated: with migration, some of the people in villages that lose facilities will move to villages that gain

facilities, tempering the surplus loss for these people (and therefore increasing the surplus gains). A countervailing force is increased congestion in the villages that gain facilities, although the demand estimates suggest that this effect will be smaller in magnitude. Accommodating this force requires a separate model of individuals' location choices, as I adopt in related work on schooling infrastructure in Indonesia (Hsiao 2020). Combining the models is difficult because it requires modeling the strategic game between individuals and governments.

Another implicit assumption is that the government makes decisions over healthcare infrastructure independently of other infrastructure. Estimates will be biased if decisions are made jointly across all types of infrastructure, although in practice these decisions are made within individual government departments. Village fixed effects absorb part of this non-health infrastructure, and I can use the PODES data to control directly for a range of observed infrastructure, including schools and roads.

I also require that unobserved village preferences  $\lambda_v$  be fixed over time within the pre- and post-reform periods. The tension is that district mayors may change in either the pre- or post-periods, and as such their preferences may differ. One way forward is to place some parametric structure on the  $\lambda_v$  terms, for example based on political party, hometown, or some other observed mayor characteristics. Another approach is to test robustness by estimating the model separately on districts in which mayors do and do not change.

Finally, I note that village preference  $\lambda_v$  subsume a variety of structural objects that make interpreting these terms difficult. While they may contain true unobserved motivations for deviating from the surplus-maximizing allocation, they may also contain misspecification error and unobserved costs. I estimate these terms as nuisance parameters in order to obtain unbiased estimates for other parameters, but their catch-all nature means that they may not be indicative of misallocation themselves. Instead, I turn to observables like patronage, ethnicity, and constituent status to understand the sources of misallocation in this setting.

## 7 Supporting Evidence

Using quasi-experimental variation in electoral accountability and spillover effects, I provide reduced-form evidence in support of the structural findings. I control for the number of facilities constructed, the initial stock of facilities, populations, and island group dummies.

For electoral accountability, I compare districts with differential exposure to the Suharto-appointment mayors, who were arguably less constrained by electoral concerns. [Martinez-Bravo et al. \(2017\)](#) establishes that this variation is quasi-random both because the end of the Suharto regime was unexpected and because the variation in district term dates is a vestige of Dutch colonial rule. I compare districts with more and less exposure to a Suharto mayor, and I study how this

**Figure 4:** Effects of Suharto mayors and redistricting on misallocation

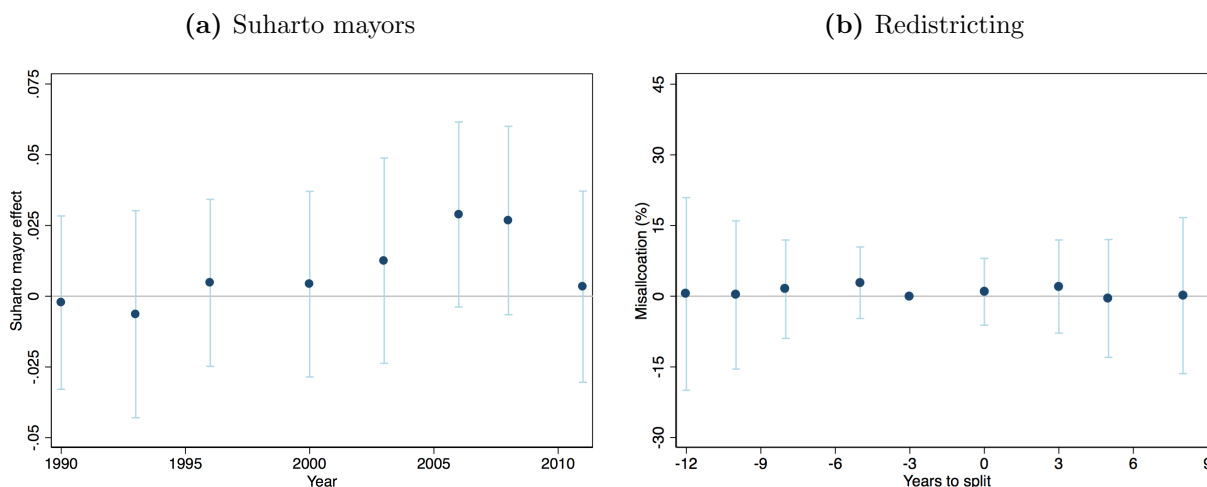


Figure 4a shows the difference in facility misallocation between districts in which Suharto mayors remained in power after the fall of Suharto and districts in which they did not. Before 2000, Suharto was in power and therefore all districts were headed by Suharto mayors. Figure 4b shows the impact of redistricting on misallocation as an event study. Redistricting splits a parent district into child districts and therefore increases the potential for uninternalized spillovers as a single constituency becomes multiple constituencies.

effect varies over time with the specification

$$\text{misallocation}_{dt} = \beta^e \left( \sum_t \text{end\_date}_{dt} \cdot \delta_t \right) + x_{dt}\beta + \delta_t + \varepsilon_{dt}. \quad (19)$$

The treatment variable  $\text{end\_date}_{dt}$  captures exposure: districts with later term expiration dates for Suharto-appointed mayors are districts with greater exposure.

For spillover effects, I study the effect of redistricting on misallocation in an event-study framework. Burgess et al. (2012) and Bazzi and Gudgeon (2017) argue that the timing of redistricting is plausibly exogenous around two national moratoria placed on redistricting from 2004 to 2006 and from 2009 to 2012. Approximately 32% of the districts in my sample undergo redistricting. Restricting attention to these districts, the following specification compares misallocation before and after redistricting.

$$\text{misallocation}_{dt} = \beta^r \text{redistricted}_{dt} + x_{dt}\beta + \delta_d + \delta_t + \varepsilon_{dt} \quad (20)$$

The treatment variable  $\text{redistricted}_{vt}$  takes a value of zero before the first instance of redistricting for a given district and a value of one afterwards.

Figure 4 presents the impact of these channels. Exposure to Suharto mayors has no effect in the pre-reform period, and this balance in the pre-treatment period is consistent with the treatment being as if randomly assigned. In the post-reform period, the effect is positive and significant. That is, greater exposure to these mayors that were unconstrained by electoral accountability

corresponds to greater misallocation in healthcare facility placements. For redistricting, I find no evidence that uninternalized spillovers generated significant misallocation in terms of social welfare. Taken together, the benefits of electoral accountability seem to outweigh the costs of uninternalized spillovers in the post-reform period. These district-level findings are consistent with the structural findings, which provide further texture by taking advantage of variation at the village level.

## 8 Conclusion

Infrastructure investment is central to economic development, but it is also a major target of corruption. This paper asks whether electoral accountability helped limit corruption in Indonesia's national expansion of healthcare infrastructure – one of the largest such efforts in recent history. I draw on spatial panel data on healthcare facility access and usage to study the construction of new hospitals, clinics, and subclinics in Indonesia, both before and after democratization in 1999. I quantify the consumer surplus generated by new facilities with a spatial model of demand for healthcare. Particularly prior to democratization, I find that the actual allocation of new facilities falls far short of the optimal allocation. To understand why, I model the facility placement decision as a dynamic discrete choice problem, and I estimate the government's objective function by revealed preference.

My main finding is that democratization decreases misallocation overall. My structural estimates suggest that, after democratization, there is less bias toward Suharto-era villages, such as those within the patronage network. A countervailing force is that spillover effects were less internalized as local electoral accountability pushed local governments to prioritize constituents over non-constituents. However, the magnitude of this second effect is relatively small. Using district-level variation in electoral accountability and district boundaries, I also find reduced-form evidence in support of this narrative.

I leave several directions open for future work. First, public and private healthcare facilities may interact in ways that I do not currently accommodate. Private facilities may compete in prices or compete spatially with public facilities, and as such may respond endogenously to changes in the placement of public infrastructure. Second, I focus on healthcare infrastructure, but healthcare may interact with other investments, such as in education or roads. Future work could study investment across several types of infrastructure jointly. Third, path dependence arises in spatial settings when infrastructure is durable because the marginal effect of new investment depends on the placement of prior investment. Thus, whether they be from corruption or otherwise, distortions today limit the gains from investment in later periods, and future work might focus on these cumulative effects.


## References

- Acemoglu, Daron, Suresh Naidu, Pascual Restrepo, and James Robinson. Democracy Does Cause Growth. *Journal of Political Economy*, 2019.
- Arcidiacono, Peter and Robert Miller. Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity. *Econometrica*, 79(6):1823–1867, 2011.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker. Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy*, 122(5):1013–1063, 2014.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker. (Mis)Allocation, Market Power, and Global Oil Extraction. *American Economic Review*, 2019.
- Balboni, Clare. In Harm’s Way? Infrastructure Investments and the Persistence of Coastal Cities. 2019.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1):305–334, 2013.
- Bazzi, Samuel and Matthew Gudgeon. The Political Boundaries of Ethnic Divisions. 2017.
- Bils, Mark, Peter Klenow, and Cian Ruane. Misallocation or Mismeasurement? 2018.
- Bryan, Gharad and Melanie Morten. The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy*, 2019.
- Burgess, Robin, Matthew Hansen, Benjamin A. 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, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel. The Value of Democracy: Evidence from Road Building in Kenya. *American Economic Review*, 105 (6):1817–1851, 2015.
- Casey, Katherine. Crossing Party Lines: The Effects of Information on Redistributive Politics. *American Economic Review*, 105(8):2410–2448, 2015.
- Donaldson, Dave and Richard Hornbeck. Railroads and American Economic Growth: A “Market Access” Approach. *Quarterly Journal of Economics*, 131(2):799–858, 2016.
- Fajgelbaum, Pablo and Edouard Schaal. Optimal Transport Networks in Spatial Equilibrium. 2017.
- Ferraz, Claudio and Frederico Finan. Exposing Corrupt Politicians: The Effects of Brazil’s Publicly Released Audits on Electoral Outcomes. *Quarterly Journal of Economics*, 123(2):703–745, 2008.
- Gadenne, Lucie and Monica Singhal. Decentralization in Developing Economies. *Annual Review of Economics*, 6:581–604, 2014.
- Ho, Kate and Ariel Pakes. Hospital Choices, Hospital Prices, and Financial Incentives to Physicians? *American Economic Review*, 104(12):3814–3840, 2014. doi: 10.1257/aer.104.12.3841.

- Holmes, Thomas. The Diffusion of Wal-Mart and Economies of Density. *Econometrica*, 79(1): 253–302, 2011.
- Hsiao, Allan. Allocating Infrastructure Investment: Evidence from School Construction in Indonesia. 2020.
- Hsieh, Chang-Tai and Peter Klenow. Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4):1403–1448, 2009.
- Kahn, Matthew, Pei Li, and Daxuan Zhao. Water Pollution Progress at Borders: The Role of Changes in China’s Political Promotion Incentives. *American Economic Journal: Economic Policy*, 7(4):223–242, 2015.
- Lipscomb, Molly and Ahmed Mushfiq Mobarak. Decentralization and Pollution Spillovers: Evidence from the Re-drawing of County Borders in Brazil. *Review of Economic Studies*, 84(1):464–502, 2017.
- Magnac, Thierry and David Thesmar. Identifying Dynamic Discrete Decision Processes. *Econometrica*, 70(2):801–816, 2002.
- Martinez-Bravo, Monica, Priya Mukherjee, and Andreas Stegmann. The Non-Democratic Roots of Elite Capture: Evidence From Soeharto Mayors in Indonesia. *Econometrica*, 85(6):1991–2010, 2017.
- Mauro, Paolo. Corruption and Growth. *Quarterly Journal of Economics*, 110(3):681–712, 1995.
- McFadden, Daniel. Econometric Models of Probabilistic Choice. In Manski, Charles and Daniel McFadden, editors, *Structural analysis of discrete data with econometric applications*, pages 198–272. MIT Press, Cambridge, MA, 1981.
- Mookherjee, Dilip. Political Decentralization. *2Annual Review of Economics*, 7:231–249, 2015.
- Oates, Wallace. *Fiscal Federalism*. Harcourt Brace Jovanovich, New York, 1972.
- Olken, Benjamin. Monitoring Corruption: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy*, 115(2):200–249, 2007.
- Oxford Economics. Global Infrastructure Outlook. Technical report, 2017.
- Pakes, Ariel. Alternative Models for Moment Inequalities. *Econometrica*, 78(6):1783–1822, 2010.
- Pakes, Ariel, Jack Porter, Kate Ho, and Joy Ishii. Moment Inequalities and Their Application. *Econometrica*, 83(1):315–334, 2015.
- Rotemberg, Martin and T. Kirk White. Measuring Cross-Country Differences in Misallocation. 2017.
- Seabright, Paul. Accountability and Decentralization in government: An Incomplete Contracts Model. *European Economic Review*, 40:61–89, 1996.
- Sigman, Hilary. International Spillovers and Water Quality in Rivers: Do Countries Free Ride? *American Economic Review*, 92(4):1152–1159, 2002.
- Zheng, Fanyin. Spatial Competition and Preemptive Entry in the Discount Retail Industry. 2016.

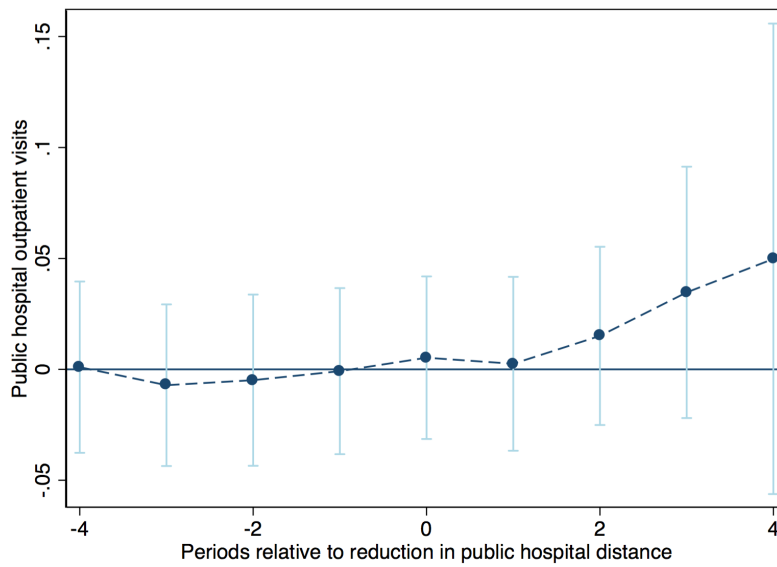
# Appendix: Figures and Tables

**Figure A1:** Correlation between usage and distance by facility type

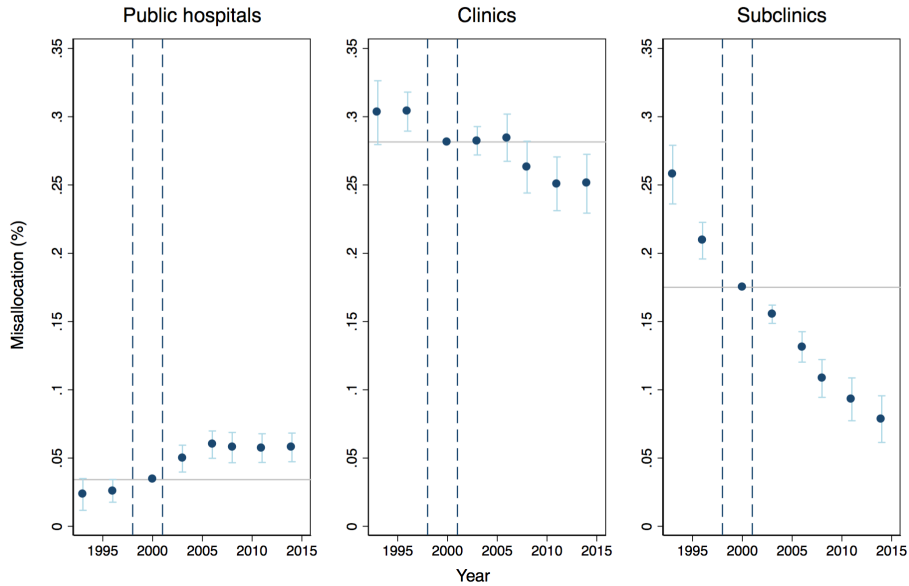


Binned scatter plots controlling for year fixed effects.

**Figure A2:** Pretrends for usage before public hospital construction

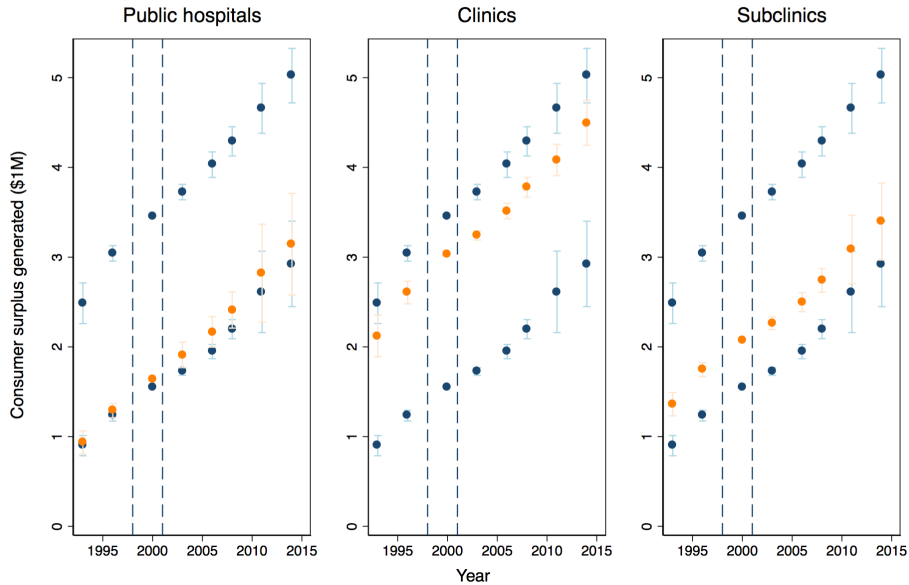


**Figure A3:** Misallocation over time by facility type



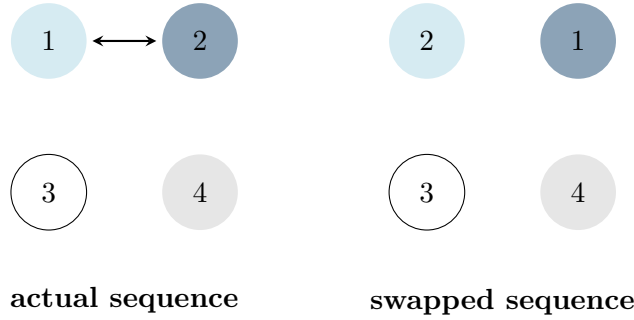
Misallocation is defined as one minus the proportion of the best achievable consumer surplus gain that is achieved by the observed placement. It is zero when the actual placement coincides with the surplus-maximizing placement. These figures show the contribution of each facility type of overall misallocation. For each, I compute the optimal placement of the facility type of interest holding fixed all other facility types.

**Figure A4:** Best achievable vs. actual policy over time by facility type



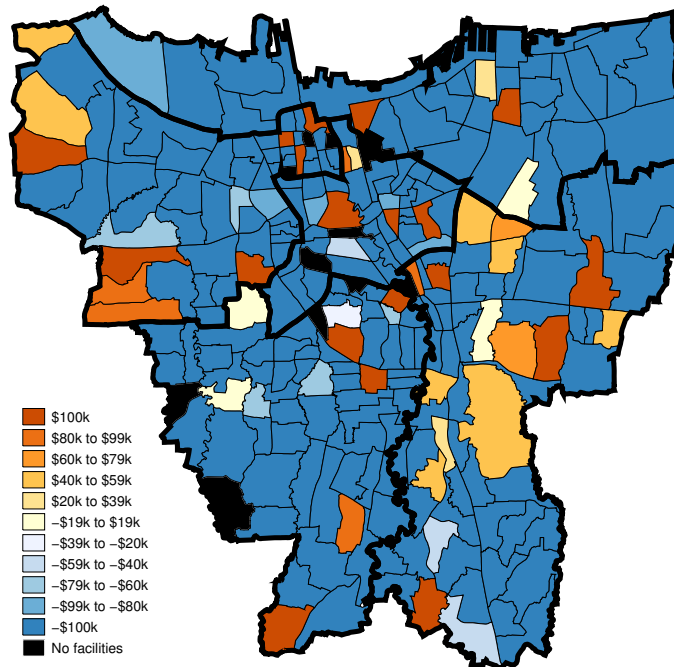
The top blue line is the maximum consumer surplus gain achievable with the facility budget in a given time period. The bottom blue line is the consumer surplus gain achieved by the actual placement. The orange line is the maximum achievable by optimizing over the facility type of interest while holding all other facility types fixed. For each period, the benchmark placement  $a$  is the facility placement in 1990.

**Figure A5:** Pairwise resequencing eliminates dynamics outside of swap



Suppose the observed order of construction is red, blue, white, and gray. I consider an alternative that swaps the order of red and blue construction. These two sequences result in the same allocation from the second period, so attention can be restricted only to where there are differences – namely, periods within the swap (in this case, the first period). In this way, choosing alternatives by pairwise resequencing eliminates dynamic considerations outside of the swap periods.

**Figure A6:** Estimated village preferences  $\lambda_v$ , Jakarta post-decentralization



**Figure A7:** Spurious misallocation under misspecification of the dynamic horizon



Each circle is a village that is a candidate to receive a facility. On the left, without accounting for future construction, there is only one facility to be placed in the current period. Placing it in the middle village puts it in close proximity to all villages. On the right, the decision maker accounts for having an additional facility to place in the following period. It is therefore optimal to place the first facility in the left village given that the second facility will be placed in the right village. This simple example illustrates the importance of how the dynamic horizon – the look-ahead window – is specified. A forward-looking decision-maker will place the first facility in the left village, but this action looks like misallocation under a model that assumes a myopic decision-maker.

**Table A1:** Usage by facility distance and congestion, population-density interaction

	Estimate	Standard Error
Distance, public hospital	-2.058***	(0.0568)
Distance, private hospital	-0.950***	(0.0298)
Distance, clinic	-2.817***	(0.241)
Distance, subclinic	-4.467***	(0.447)
Distance, public hospital × population density	-0.0424	(0.0353)
Distance, private hospital × population density	-0.0303	(0.0210)
Distance, clinic × population density	-0.155	(0.137)
Distance, subclinic × population density	-0.803	(0.646)
Congestion, public hospital	-0.0206***	(0.00298)
Congestion, private hospital	-0.0133***	(0.00123)
Congestion, clinic	-0.281***	(0.0346)
Congestion, subclinic	-0.468***	(0.0380)
Congestion, public hospital × population density	0.00268	(0.00171)
Congestion, private hospital × population density	0.00196**	(0.000900)
Congestion, clinic × population density	0.0172**	(0.00738)
Congestion, subclinic × population density	0.0426***	(0.0136)
Price	-0.764***	(0.152)
Village FE		x
Facility type-year FE		x
Observations		202,668

Each column is a single conditional multinomial logit regression with village and facility type-year fixed effects. The unit of observation is a village-year-facility type, where the set of facility types represents a village’s choice set in a given year. The outcome is usage by facility type, as recorded in the SUSENAS data. Distance is to the closest facility of each type and is measured in units of 100 km. Congestion of the closest facility is the number of people for whom this facility is the closest of its type. This variable is measured in units of 100,000 people. Price is measured in units of \$100 (in year 2000 USD). Population density is measured in units of 10,000 people per square kilometer. Additional controls include population and ruralness. Standard errors are clustered by village. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

## **Chapter 3**

Allocation Rules for Infrastructure Investment:  
Evidence from School Construction in Indonesia

# 1 Introduction

Optimal infrastructure investment is a difficult problem. In the presence of spatial effects, choosing where to place investments is a classic facility-location problem, known to be NP-hard and often solved with approximation algorithms (Korte and Vygen 2008). Economists have made recent progress by combining sophisticated optimization techniques with the structure of spatial equilibrium models (Fajgelbaum and Schaal 2020). However, these methods differ markedly from what policymakers use in practice – most commonly, simple rules like population cutoffs, ranked lists, and need-based formulas that do not account for spatial interdependence (e.g., Faber 2014; Gertler et al. 2019; Asher and Novosad 2020). How effective are these simple rules compared to more sophisticated approaches?

I study this question in the context of schooling infrastructure. Schools have spatial effects that extend beyond local labor markets because graduates migrate and seek employment nationally. As my empirical setting, I focus on Indonesia’s Sekolah Dasar INPRES program, a massive school construction effort that established 61,807 new primary schools between 1973 and 1978. Representative of the typical approach, policymakers adopted a simple rule that allocated schools across districts in proportion to initial enrollment rates. I use a spatial equilibrium model to quantify the program’s aggregate effects and compute the optimal allocation. I then test a range of alternative allocation rules and benchmark their effectiveness against the optimum.

I begin by evaluating the effects of school construction with the difference-in-difference strategy of Duflo (2001). SUSENAS household survey data from 2011 to 2014 allow me to measure long-term effects. Comparing treated (young) and untreated (old) age cohorts in districts with high and low levels of school construction, I find persistent positive effects on education and wages, particularly in districts with few existing schools and high market access. These findings frame the key forces that determine the optimal spatial allocation of school construction. On one hand, diminishing marginal returns favor spreading schools evenly across districts. On the other hand, the benefits of market access favor concentrating schools in non-isolated districts. Navigating these forces requires targeting districts with both high initial returns and high market access. However, migration complicates doing so by making school construction decisions interdependent across districts.

To this end, I present a spatial equilibrium model that explicitly captures diminishing marginal returns, market access, and migration, building on Bryan and Morten (2019) and Hsieh et al. (2019). First, individuals pursue education, subject to education costs that can be decreased with school construction. Next, individuals seek employment within the national labor market, subject to migration costs. A concave human capital function captures diminishing marginal returns to school construction in a given district. Market access in the form of low migration costs increases the incentives to invest in schooling by expanding the pool of job opportunities. While rural residents may have few high-skill jobs at home, they will still invest in education if low migration costs allow

them access to high-skill jobs elsewhere. Finally, migration implies that school construction in one district affects labor markets in other districts, particularly if agglomeration increases productivity as high-skill migrants move in. These labor market interactions generate the spatial interdependence that makes computing the optimal allocation of school construction difficult.

I estimate the model using the same difference-in-difference variation described above, and I use it to quantify the effects of school construction on aggregate output. Comparing treated and untreated age cohorts in districts with high and low levels of school construction, I obtain credible estimates of how education costs respond to school construction, as well as the returns to schooling. When agglomeration is weak, I show that these parameters alone are sufficient for computing counterfactual aggregate output. There is no need to estimate the rest of the model, and I discuss conditions under which the remaining optimization problem can be solved quickly using standard convex optimization techniques. When agglomeration is strong, wages, amenities, and migration respond to school construction and must be recomputed in equilibrium. I present an iterative algorithm for doing so in a tractable way.

I find that the school construction program increased aggregate output by seven percent in the long run, but that the optimal allocation would have increased aggregate output by another four percentage points. The actual allocation used a simple rule to target regions with high unenrollment. This targeting was warranted, as the actual allocation outperformed the untargeted, uniform allocation, but the difference is not large. The performance is improved by incorporating higher-order nonlinearities, but the largest gains come from further conditioning on distance to the nearest city. In this case, an allocation rule does well in approximating the prescriptions of the full model despite being much simpler to implement. Intuitively, the model features two main forces: diminishing marginal returns and market access. Conditioning on unenrollment and city distance captures both forces, while conditioning only on unenrollment misses the latter.

The cost is that more sophisticated rules are less robust to unintended deviations. First, implementation error may cause the final allocation to differ from that prescribed by the rule. Second, the parameters of the rule – such as the slope and intercept of a linear rule – may be improperly specified or estimated. Third, there may be mismeasurement in the observables used to compute allocations. Simplicity helps to avoid overfitting, and thus simpler rules are more robust. Overfitting is a concern because all allocation rules are approximations: they are straightforward to implement precisely because they abstract from the full spatial interdependence of the problem. Furthermore, simpler rules are more transparent and perhaps more politically feasible.

This work is closely related to a growing literature that uses quantitative spatial equilibrium models to analyze infrastructure investments. [Donaldson and Hornbeck \(2016\)](#) studies American railroads, [Donaldson \(2018\)](#) Indian railroads, [Allen and Arkolakis \(2019\)](#) American highways, [Balboni \(2019\)](#) Vietnamese roads, and [Tsivanidis \(2019\)](#) Columbian buses. These papers focus on transportation infrastructure, where the endogenous location of economic activity makes the use of

spatial equilibrium models quite natural. My contribution is in arguing that schooling infrastructure generates similar spatial interdependence via labor market interactions and migration, as new schools build durable human capital that migrants take with them. The papers in this literature also largely avoid solving for the optimal allocation of investments, which is itself a difficult problem that has only recently been solved (Fajgelbaum and Schaal 2020). I tackle this optimization problem directly and discuss scenarios in which it is relatively easy to solve. I also provide a new focus on the design of simple allocation rules given the stark contrast between what policymakers do in practice and what economists suggest in theory.

To do so, I build on the labor allocation models of Hsieh et al. (2019) and Bryan and Morten (2019). Hsieh et al. (2019) study occupational choice and frictions from discrimination in the US, while Bryan and Morten (2019) study migration choice and frictions from distance in Indonesia. I endogenize the education decision as in Hsieh et al. (2019), and I study migration in Indonesia as in Bryan and Morten (2019). My contribution is to nest these models within a larger optimization problem and to highlight how labor market interactions generate spatial interdependence and complicate the optimization problem. I also show how quasi-experimental variation in school construction identifies the key parameters of interest, including one education-related parameter that these papers must calibrate.

Finally, I contribute to the literature on the INPRES school construction program itself. Following seminal work by Duflo (2001), Duflo (2004) discusses cross-cohort labor market spillovers, Martinez-Bravo (2017) finds positive effects on local public goods provision, and Bazzi et al. (2021) studies the resulting interaction between public and religious schools. Like Bazzi et al. (2021), I document long-run effects of the program. Relative to these papers, I estimate an empirical model and quantify aggregate effects in spatial equilibrium. This approach allows me to simulate counterfactual allocations of resources for the largest school construction program in history.

## 2 Data and Stylized Facts

This section describes the data and how INPRES schools were allocated across districts. I use a difference-in-difference approach to show that INPRES school construction increased educational attainment and wages, particularly in areas with few pre-program schools and high market access.

### 2.1 Data

District-level data on INPRES school construction come from Duflo (2001), which draws on data from the Ministry of National Development Planning (*Bappenas*) and the 1971 population census. INPRES refers to the “presidential instructions” that established the school construction program. The data record the number of INPRES primary schools constructed, the number of pre-program primary schools, 1971 child populations and enrollment rates, and INPRES water

and sanitation spending per capita. For each district, I compute market access by combining the district’s own population in 1971 with a weighted sum of other districts’ populations. Weights are inversely proportional to distance: districts in close proximity to large populations have high market access, while isolated districts have low market access.<sup>1</sup>

The main individual-level data come from the 2011, 2012, 2013, and 2014 National Socio-economic Surveys (SUSENAS). I observe districts of both residence and birth, with the latter providing the link to INPRES program exposure. The data record a range of educational and employment outcomes, including educational attainment and monthly wages. Non-wage income from self-employment is not observed, although self-employment activity is. I restrict attention to male heads of household ages 2 to 24 in 1974 – when the first INPRES schools were completed – and I adjust districts to 1971 boundaries for consistency over time.<sup>2</sup> Finally, the 1976 and 1995 Intercensal Population Surveys (SUPAS) provide similar individual-level data and thus allow for an additional set of placebo experiments in the analysis that follows.

## 2.2 School construction targeted high-unenrollment districts

The INPRES program had the stated goal of constructing 62,000 schools across Indonesia: 6,000 in the fiscal year beginning in 1973, 6,000 in 1974, 10,000 in 1975, 10,000 in 1976, 15,000 in 1977, and 15,000 in 1978 (*Inpres* No. 10/1973, 6/1974, 6/1975, 3/1976, 3/1977, 6/1978). The allocation rule was as follows. In 1973 and 1974, schools were distributed across districts in proportion to pre-program unenrollment rates for children of primary school age. From 1975 to 1978, unenrollment was instead defined relative to a 15% threshold, and the allocation rule prescribed no new schools to districts with unenrollment rates below 15%. Figure 1 shows that the data indeed reflect this linear proportionality between school construction and unenrollment rates.

## 2.3 School construction increased education and wages

I evaluate the long-run effects of INPRES school construction with the difference-in-difference strategy of [Duflo \(2001\)](#). Since new primary schools only benefit those of primary school age, individuals ages 2 to 6 in 1974 form the treatment group while those ages 12 to 17 in 1974 form the control group. I compare these groups in regions with high versus low levels of school construction.


$$Y_{ijk} = \delta_j + \delta_k + \beta S_j T_k + \mathbf{C}_j T_k \phi + \varepsilon_{ijk}, \quad (1)$$

for individuals  $i$  born in district  $j$  and age cohort  $k$ . It includes outcome variable  $Y_{ijk}$ , district-of-birth fixed effect  $\delta_j$ , year-of-birth fixed effect  $\delta_k$ , school construction intensity  $S_j$ , treatment dummy

<sup>1</sup> Weights are  $(1 + \text{dist}_{dd'})^{-2}$ , with distance between districts  $d$  measured in hundreds of kilometers.

<sup>2</sup> Rapid decentralization transformed Indonesia from 26 provinces, 234 districts, and 64 municipalities in 1999 to 33 provinces, 398 districts, and 98 municipalities today. I refer to sub-provincial administrative units – both non-urban districts (*kabupaten*) and urban municipalities (*kota*) – as “districts” in this paper.

**Figure 1:** INPRES school construction vs. unenrollment rates by district



Each figure is a binned scatter plot, and each observation is one district. The  $y$ -axis is the proportion of total school construction allocated to each district. In line with the actual allocation rule used, the  $x$ -axis in 1973/1974 is the pre-program unenrollment rate among children of primary school age, and in other years is the extent to which the pre-program unenrollment rate exceeds a target of 15%. I omit outliers by dropping the 5% of districts with the largest unenrollment rates.

$T_k$ , district-of-birth controls  $C_j$ , and error term  $\varepsilon_{ijk}$ .<sup>3</sup> Program intensity is given by the number of INPRES schools constructed per 1,000 children. I also include survey-year fixed effects because I pool SUSENAS data from multiple waves. The coefficient of interest is  $\beta$ , which captures the causal effect of school construction under the common-trends assumption that high- and low-construction regions would have evolved similarly absent the program.

Placebo experiments corroborate the identifying assumption. The main placebo experiment uses the SUSENAS data to compare individuals ages 12 to 17 and those ages 18 to 24 in 1974. This comparison should yield null effects because both groups were too old to be exposed to INPRES primary schools. The SUPAS data allow for another set of similar experiments. In the 1995 SUPAS data, I compare individuals ages 12 to 17 and those ages 18 to 24 in 1974 – the same cohorts in the primary placebo experiment. In the 1976 SUPAS data, I compare individuals ages 12 to 17 and those ages 18 to 24 in 1955. The cohorts in the primary placebo experiment are not yet of working age in these data, and thus I focus on cohorts analogous to those in the 1995 experiment.

The first panel of table 1 shows that INPRES school construction increased both educational attainment and wages in the long run. These results are consistent with the medium-run findings of [Duflo \(2001\)](#). The first three columns show the results of the main experiment: school construction increases years of schooling, both in the full sample and in the selected sample of wage earners, and it also increases log monthly wages. The last three columns show the results of the main

<sup>3</sup> Following [Duflo \(2001\)](#), controls  $C_j$  include 1971 child populations, 1971 enrollment rates, and INPRES spending on water and sanitation projects. The first two controls account for the fact that INPRES schools were allocated based on these measures, such that high- and low-construction regions differed along these dimensions to begin with. The third control addresses the concern that high-construction regions also received other forms of treatment, the effects of which should not be attributed to school construction.

**Table 1:** Effects of INPRES school construction on education and wages

	Treatment			Placebo		
	Years of schooling	Years of schooling	Log wages (month)	Years of schooling	Years of schooling	Log wages (month)
INPRES × young	0.103** (0.0424)	0.121** (0.0495)	0.0195** (0.00916)	-0.0176 (0.0318)	0.0120 (0.0566)	-0.00765 (0.00890)
INPRES × young	0.164*** (0.0554)	0.221*** (0.0679)	0.0378*** (0.0124)	-0.0338 (0.0411)	0.00121 (0.0807)	-0.00858 (0.0133)
— × existing schools	-0.0737* (0.0388)	-0.116** (0.0505)	-0.0212** (0.00906)	0.0196 (0.0302)	0.0126 (0.0573)	0.00109 (0.00972)
INPRES × young	0.0445 (0.0411)	0.0246 (0.0591)	0.00694 (0.0113)	0.00266 (0.0335)	0.0458 (0.0645)	-0.0130 (0.00939)
— × market access	0.122*** (0.0349)	0.197*** (0.0447)	0.0257*** (0.00904)	-0.0411 (0.0285)	-0.0711 (0.0492)	0.0112 (0.00786)
Observations	233,517	89,404	89,404	196,308	55,091	55,091
Mean	7.75	9.57	14.35	6.50	8.55	14.28

Each column of each panel is one regression. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. The first three columns display treatment estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1974. The last three columns display placebo estimates that compare individuals ages 12 to 17 and those ages 18 to 24 in 1974. Column headings list outcome variables, with the second and fifth columns restricted to the sample of wage earners. The first panel presents baseline treatment effects, while the second and third panels present interacted treatment effects. The interaction terms involve dummy variables for having an above-median number of existing primary schools or an above-median degree of market access. The numbers of observations and means are common to all three panels. Regressions control for birth district fixed effects, birth year fixed effects, survey year fixed effects, 1971 child population, 1971 enrollment rates, and INPRES spending on water and sanitation projects. Standard errors are clustered by birth district, as defined by 1971 boundaries. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

placebo experiment, which indeed yields null effects. Furthermore, while I observe wage data only for the sample of wage earners, the first panel of table 2 shows that selection into this sample is not meaningfully affected by school construction. There is also no impact on weekly hours, and thus the wage effect must imply higher wage rates. Finally, appendix table A1 shows the results of the SUPAS placebo experiments and finds null effects for all outcomes.

## 2.4 Trade-off between diminishing marginal returns and market access

The second and third panels of table 1 show how the INPRES treatment effect interacts with the number of existing primary schools and the degree of market access. The specification is

$$Y_{ijk} = \delta_j + \delta_k + \beta S_j T_k + \beta' S_j T_k X_j + C_j T_k \phi + \varepsilon_{ijk}, \quad (2)$$

which adds an interaction term with district-of-birth covariate  $X_j$ . First, there are diminishing marginal returns to primary school construction. The education and wage effects of INPRES schools are half as large in districts with many existing primary schools as they are in districts with

**Table 2:** Effects of INPRES school construction on employment, school completion, and migration

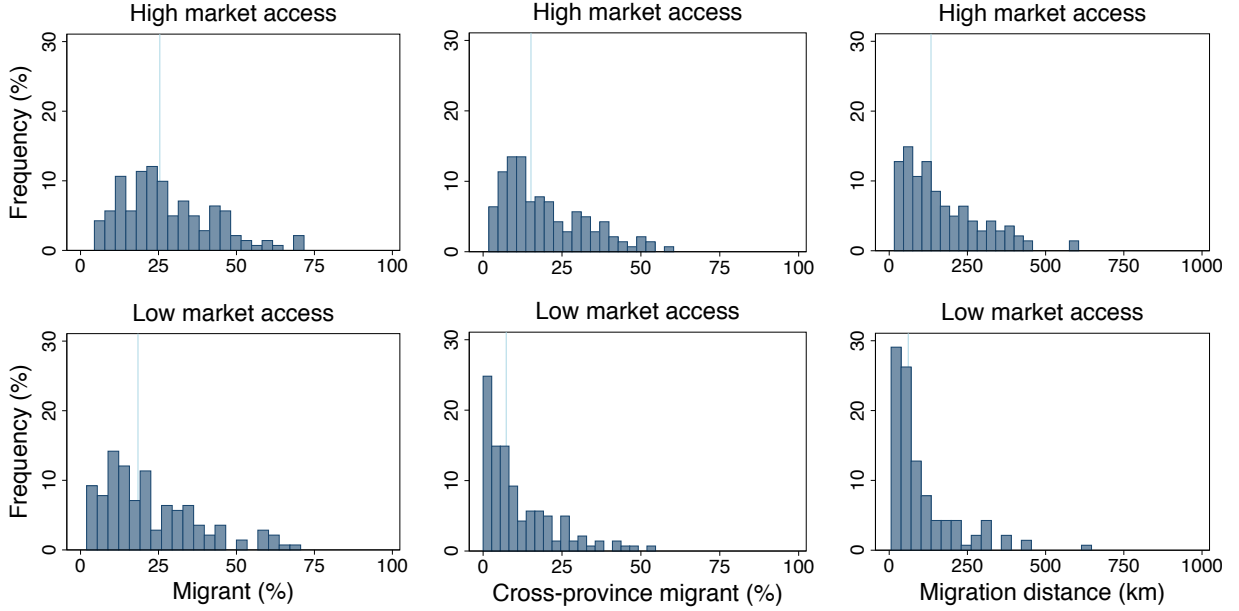
	Treatment			Placebo		
	Estimate	SE	Obs	Estimate	SE	Obs
Employment	0.0304	(0.0278)	241,173	0.0309	(0.0216)	203,995
Wage employment	0.000376	(0.0131)	241,173	-0.0204	(0.0189)	203,995
Self-employment	-0.00219	(0.0119)	241,173	0.0140	(0.0142)	203,995
Weekly hours	-0.136	(0.102)	229,662	-0.00968	(0.109)	183,840
Primary school completion	0.0585**	(0.0291)	233,517	-0.0134	(0.0167)	196,308
Middle school completion	0.0480**	(0.0207)	233,517	0.00573	(0.0156)	196,308
High school completion	0.0292	(0.0180)	233,517	-0.00167	(0.0140)	196,308
University completion	-0.0236	(0.0196)	233,517	-0.00791	(0.0214)	196,308
Migrant	0.0244	(0.0194)	244,793	-0.0249*	(0.0129)	210,543
Migrant to Jakarta	0.156*	(0.0867)	188,622	-0.0353	(0.0687)	158,719

Each row is one treatment regression and one placebo regression. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. The first three columns display treatment estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1974. The last three columns display placebo estimates that compare individuals ages 12 to 17 and those ages 18 to 24 in 1974. Row headings list outcome variables. I run logit regressions for dummy outcomes – all outcomes except for weekly hours. Regressions control for birth district fixed effects, birth year fixed effects, survey year fixed effects, 1971 child population, 1971 enrollment rates, and INPRES spending on water and sanitation projects. Standard errors are clustered by birth district, as defined by 1971 boundaries. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

few existing primary schools (with “many” and “few” defined as above- and below-median). The first panel of table 2 suggests why: new primary schools only increase primary and middle school completion, and so the effects of primary school construction are bounded. Second, the effects of INPRES schools are driven by construction in districts with a high degree of market access (with “high” again defined as above-median). Figure 2 illustrates why: birth districts with high market access experience more out-migration, particularly across provinces and over long distances, as individuals seek opportunities nationally. By contrast, those in isolated districts with low market access have high barriers to migration, limiting both the pool of job opportunities and the incentives to pursue education. Lastly, the placebo experiments yield insignificant estimates in each case.

These interactions suggest a trade-off in allocating INPRES schools. Diminishing marginal returns favor spreading schools evenly across districts, while market access favors concentrating schools in non-isolated districts. Navigating these forces requires targeting districts with both few existing schools and high market access, and appendix figure A1 shows that such districts do exist. Districts with fewer existing schools in fact tend to have higher market access, although the correlation is too weak to focus only on one or the other. The model that follows therefore incorporates both forces explicitly. At the same time, migration complicates the school allocation problem by making school construction interdependent across districts. In particular, schools have spatial effects as students graduate and seek work nationally, such that school construction in one district depends on and affects labor markets in all other districts. Indeed, figure 2 shows

**Figure 2:** Out-migration by birth district



Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household ages 2 to 24 in 1974. I compute mean values by birth district, such that each observation in the above figures corresponds to one birth district as defined by 1971 district boundaries. Migrants are individuals with districts of residence that differ from their districts of birth. Cross-province migrants are defined analogously based on provinces of residence and birth. Migration distances are Euclidean distances between district centroids. “High” and “low” market access refer to districts with above- and below-median market access. Market access is the sum of a district’s own population in 1971 and an inverse-distance-weighted sum of other districts’ populations. The light blue vertical lines mark the median values in each figure: 28%, 23%, 19%, and 11% from left to right.

substantial migration in the data, and the third panel of table 2 shows that school construction may even increase migration to Jakarta (although perhaps not overall). The model captures these migration decisions and allows me to evaluate school construction in spatial equilibrium.

### 3 Model

Building on Bryan and Morten (2019) and Hsieh et al. (2019), this section presents a spatial equilibrium model in which individuals invest in education, then choose to migrate for work.

#### 3.1 Utility and migration

Consider an individual  $i$  born in origin district  $j$  and age cohort  $k$ . The individual realizes a skill draw for each possible destination  $\ell \in \mathcal{L}$ , and given these skill draws chooses schooling and a single destination. Utility for a given destination is

$$U(e, \epsilon) = \alpha_\ell \left[ \underbrace{(1 - \tau_{j\ell}^m) w_\ell h_{jk} e^\eta \epsilon}_{\text{net labor income}} - \underbrace{(1 + \tau_{jk}^e) c e}_{\text{cost of education}} \right], \quad (3)$$

where I suppress subscripts  $ijk\ell$ . The individual considers amenities  $\alpha_\ell$  and consumption, where consumption is net labor income less the total cost of education. Total labor income is the product of base wage  $w_\ell$  and human capital, which in turn combines base human capital  $h_{jk}$ , schooling  $e$ , returns to schooling  $\eta$ , and skill draw  $\epsilon$ . Individuals across cohorts are perfect substitutes conditional on human capital and thus face common base wages. Net labor income is net of migration costs  $\tau_{j\ell}^m$ , which capture the consumption-denominated costs – financial, psychological, and otherwise – of being away from home. The total cost of education is the product of base cost  $c$  and schooling  $e$ , amplified by education costs  $\tau_{jk}^e$ . I interpret INPRES school construction as lowering education costs for treated age cohorts through increasing accessibility. Human capital is concave in schooling for  $\eta < 1$ , reflecting diminishing marginal returns.

For each destination, the individual chooses schooling conditional on the skill draw associated with that destination. Schooling increases human capital and thus labor income, but also increases the total cost of education. The optimal schooling choice is

$$e^* = \arg \max_e \{U(e, \epsilon)\} = \left[ \frac{(1 - \tau_{j\ell}^m)w_\ell h_{jk} \eta \epsilon}{(1 + \tau_{jk}^e)c} \right]^{\frac{1}{1-\eta}}.$$

Schooling is decreasing in education costs, which make schooling costly, and in migration costs, which lower net labor income and reduce the returns to schooling. Schooling choices are subject to optimization error  $\varepsilon_{j\ell}^e$ , which I can interpret as misspecification in amenities, wages, and the cost of education.<sup>4</sup> Schooling choice and utility are therefore

$$\begin{aligned} \hat{e} &= e^*(\varepsilon_{j\ell}^e)^{\frac{1}{1-\eta}}, \\ U(\epsilon) = U(\hat{e}, \epsilon) &= (1 - \eta)\eta^{\frac{\eta}{1-\eta}}\alpha_\ell \left[ \frac{(1 - \tau_{j\ell}^m)w_\ell h_{jk} \epsilon}{[(1 + \tau_{jk}^e)c]^\eta} \right]^{\frac{1}{1-\eta}} (1 - \varepsilon_{j\ell}^e)(\varepsilon_{j\ell}^e)^{\frac{\eta}{1-\eta}}. \end{aligned} \quad (4)$$

Conditional on skill draws and schooling choices, the individual compares utilities across destinations and chooses the utility-maximizing destination. Skill draws are Fréchet distributed as in [Hsieh et al. \(2019\)](#) and [Bryan and Morten \(2019\)](#), which in turn follow [McFadden \(1974\)](#) and [Eaton and Kortum \(2002\)](#).

$$F(\epsilon_1, \dots, \epsilon_L) = \exp \left\{ - \sum_\ell \epsilon_\ell^{-\theta} \right\},$$

for shorthand  $\sum_\ell \equiv \sum_{\ell=1}^L$ . Skill dispersion is captured by  $\theta$ , with a high value indicating low dispersion. This distributional assumption yields closed-form choice probabilities across destinations.

$$\pi_{jk\ell} = \frac{\tilde{w}_{jk\ell}^\theta}{\sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta} \quad \text{for} \quad \tilde{w}_{jk\ell} \equiv \alpha_\ell^{1-\eta} (1 - \tau_{j\ell}^m) w_\ell \underbrace{(1 - \varepsilon_{j\ell}^e)^{1-\eta} (\varepsilon_{j\ell}^e)^\eta}_{\tilde{\varepsilon}_{j\ell}^e} \quad (5)$$

<sup>4</sup> Consider utility  $U(e, \epsilon) = \alpha_\ell \varepsilon_{j\ell}^\alpha [(1 - \tau_{j\ell}^m)w_\ell \varepsilon_{j\ell}^w h_{jk} e^\eta \epsilon - (1 + \tau_{jk}^e)c \varepsilon_{j\ell}^c e]$ .

Destinations with low migration costs attract migrants, as do destinations with high amenities and base wages. Education costs do not enter directly because they affect destinations equally. But education costs can enter indirectly in the presence of agglomeration that affects base wages, as I define below. Congestion is an opposing force, however, such that migration may not respond significantly to school construction – as in table 2 – even under strong agglomeration.

### 3.2 Education and wages

Average education and wages for individuals of origin  $j$  and cohort  $k$  in destination  $\ell$  are

$$\begin{aligned}\overline{educ}_{jkl} &\equiv \mathbb{E}[e^* \mid \text{individuals choose } \ell] \\ &= \gamma \left( \frac{1}{\alpha_\ell} \right) \left[ \frac{h_{jk}\eta}{(1 + \tau_{jk}^e)c} \right]^{\frac{1}{1-\eta}} \left( \sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta \right)^{\frac{1}{\theta(1-\eta)}} \left( \frac{\varepsilon_{jkl}^e}{1 - \varepsilon_{jkl}^e} \right),\end{aligned}\quad (6)$$

$$\begin{aligned}\overline{wage}_{jkl} &\equiv \mathbb{E}[w_\ell h_{jk} e^{\eta\epsilon} \mid \text{individuals choose } \ell, e = e^*] \\ &= \gamma \left( \frac{1}{\alpha_\ell} \right) \left( \frac{1}{1 - \tau_{jk}^m} \right) \left[ \frac{h_{jk}^{1/\eta} \eta}{(1 + \tau_{jk}^e)c} \right]^{\frac{\eta}{1-\eta}} \left( \sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta \right)^{\frac{1}{\theta(1-\eta)}} \left( \frac{1}{1 - \varepsilon_{jkl}^e} \right),\end{aligned}\quad (7)$$

for  $\gamma = \Gamma(1 - \frac{1}{\theta(1-\eta)})$ , noting that  $\mathbb{E}[\epsilon^{\frac{1}{1-\eta}} \mid \text{individuals choose } \ell] = \pi_{jkl}^{-\frac{1}{\theta(1-\eta)}} \gamma$ . Conditional on destination choice, education costs decrease both education and wages by reducing schooling and thus human capital. Migration costs directly increase wages because those that overcome higher barriers are positively selected. Migration costs indirectly decrease both education and wages through the summation terms, as high migration costs mean fewer effective Fréchet draws and thus lower returns to investment in schooling. Base wages do not directly enter because higher base wages attract individuals with increasingly poor skill draws, decreasing average education and wages. With Fréchet draws, this countervailing force exactly offsets higher base wages. Amenities decrease both education and wages because they attract individuals independent of schooling choices.

Summing across destinations gives the origin-cohort quantities considered in section 2.

$$\overline{educ}_{jk} = \sum_{\ell} \overline{educ}_{jkl} \pi_{jkl}, \quad \overline{wage}_{jk} = \sum_{\ell} \overline{wage}_{jkl} \pi_{jkl}$$

Education costs decrease both quantities by reducing schooling and thus  $\overline{educ}_{jkl}$  and  $\overline{wage}_{jkl}$  across destinations. Furthermore, the effects of education costs interact with migration costs and therefore depend on market access. The partial derivatives with respect to education costs are

$$\frac{\partial \overline{educ}_{jk}}{\partial \tau_{jk}^e} = \frac{\partial}{\partial \tau_{jk}^e} \left\{ \gamma \left[ \frac{h_{jk}\eta}{(1 + \tau_{jk}^e)c} \right]^{\frac{1}{1-\eta}} \underbrace{\left( \sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta \right)^{\frac{1}{\theta(1-\eta)}} \left( \sum_{\ell} \frac{\pi_{jkl}}{\alpha_\ell} \right)}_{f(\tau_j^m)} \left( \frac{\varepsilon_{jkl}^e}{1 - \varepsilon_{jkl}^e} \right) \right\},$$

$$\frac{\partial \overline{wage}_{jk}}{\partial \tau_{jk}^e} = \frac{\partial}{\partial \tau_{jk}^e} \left\{ \gamma \left[ \frac{h_{jk}^{1/\eta} \eta}{(1 + \tau_{jk}^e)c} \right]^{\frac{\eta}{1-\eta}} \underbrace{\left( \sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta \right)^{\frac{1}{\theta(1-\eta)}} \left( \sum_{\ell} \frac{\pi_{jk\ell}}{\alpha_\ell(1 - \tau_{j\ell}^m)} \right)}_{g(\tau_j^m)} \left( \frac{1}{1 - \varepsilon_{jk\ell}^e} \right) \right\},$$

grouping terms dependent on migration costs  $\tau_j^m = \{\tau_{j\ell}^m\}$  for  $\ell \in \mathcal{L}$ . I study the impact of migration costs by comparing the extreme cases of full isolation and mobility.

$$\tau_{j\ell}^{m, \text{isolated}} = \mathbb{1}[\ell \neq j], \quad \tau_{j\ell}^{m, \text{mobile}} = 0,$$

for all destinations  $\ell$ . For these migration costs,  $f(\tau_j^m)$  and  $g(\tau_j^m)$  coincide.

$$f(\tau_j^{m, \text{isolated}}) = w_j^{\frac{1}{1-\eta}}, \quad f(\tau_j^{m, \text{mobile}}) = \left( \sum_{\hat{\ell}} \alpha_{\hat{\ell}}^{\theta(1-\eta)} w_{\hat{\ell}}^\theta \right)^{\frac{1}{\theta(1-\eta)}} \left( \sum_{\ell} \frac{\pi_{jk\ell}}{\alpha_\ell} \right)$$

The limiting case allows direct comparison. For  $\bar{w}_\ell = \alpha_\ell^{1-\eta} w_\ell$ , bounded amenities, and  $j' \neq j$ ,

$$\lim_{\bar{w}_{j'} \rightarrow \infty} f(\tau_j^{m, \text{mobile}}) = w_{j'}^{\frac{1}{1-\eta}} > f(\tau_j^{m, \text{isolated}}).$$

Full mobility allows individuals to pursue high base wages in destination  $j'$ , increasing the incentives to invest in schooling and thus the effects of reducing education costs. This force adds value to building schools in districts with high market access. If amenities are not bounded, however, individuals may be drawn to  $j'$  despite low base wages, in which case free mobility can mute the effects of schooling construction. But this case is inconsistent with the findings of section 2, and I will also show that high-amenity places tend to be high-wage in the data. Indeed, urban places often provide high wages and amenities in other settings as well ([Diamond 2016](#)).

### 3.3 Equilibrium and aggregate output

In equilibrium, base wages  $w_\ell$  clear human capital markets for each cohort in each destination.

$$\sum_{j,k} H_{jkl}^{\text{supply}} = H_\ell^{\text{demand}}$$

Schooling and migration choices by individuals determine the supply of human capital.

$$H_{jkl}^{\text{supply}} = N_{jk} \pi_{jk\ell} \underbrace{\mathbb{E}[h_{jk} e^\eta \epsilon \mid \text{individuals choose } \ell, e = e^*]}_{\bar{h}_{jk\ell}},$$

where  $N_{jk}$  is labor force by origin-cohort,  $\pi_{jk\ell}$  captures migration decisions, and  $\bar{h}_{jk\ell}$  is average worker quality. Representative firms use human capital to produce output, and profit maximization

determines the demand for human capital. For productivity  $A_\ell$ , which firms take as given,

$$H_\ell^{\text{demand}} = \arg \max_{H_\ell} \underbrace{(A_\ell H_\ell - w_\ell H_\ell)}_{\Pi_\ell}.$$

Perfect competition among firms implies zero profits, such that base wages reflect productivity.

$$w_\ell = A_\ell$$

Destinations produce differentiated goods with (constant) elasticity of substitution  $\sigma > 1$ , and summing over destinations gives CES aggregate output.

$$Y = \left( \sum_{\ell} (A_\ell H_\ell)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (8)$$

Productivity allows for agglomeration  $\kappa$ , and amenities incorporate congestion  $\lambda$ .

$$A_\ell = \bar{A}_\ell H_\ell^\kappa, \quad \alpha_\ell = \bar{\alpha}_\ell \left( \sum_{j,k} N_{jk} \pi_{jkl} \right)^{-\lambda} \quad (9)$$

### 3.4 Spatial interdependence

Migration causes the effects of school construction to be interdependent across regions. That is, the effects of school construction in a given district depend on the extent of school construction in other districts. The reason is that school construction interacts across districts via labor markets: individuals migrate to Jakarta from districts across Indonesia, and so the labor market in Jakarta depends on school construction in each of these origin districts. Without such labor market interactions, school construction decisions would not be interdependent across districts because I assume individuals are educated in their birth districts.

These labor market interactions take two opposing forms in the model. On one hand, aggregate output exhibits diminishing marginal returns to human capital in each destination under imperfect substitution. There are limited benefits to increasing Jakarta's output because others prefer their local goods. This mechanism depends on the elasticity of substitution and is eliminated when  $\sigma \rightarrow \infty$ . On the other hand, there are increasing returns to human capital in each destination under agglomeration. Jakarta schools increase human capital among natives, raising productivity and base wages. Jakarta migrants then have greater incentives to invest in schooling, thus raising the value of school construction in their origin districts. This mechanism depends on the strength of agglomeration and is eliminated when  $\kappa = 0$ . Together, these mechanisms determine how to maximize the effects of school construction in spatial equilibrium. They also reveal why the allocation problem is combinatorial and therefore not straightforward to solve.

## 4 Estimation

This section describes identification of key parameters and the estimation procedure. It also presents the estimates themselves.

### 4.1 Moments and identification

Equations 5, 6, and 7 describe  $\pi_{jkl}$ ,  $\overline{educ}_{jkl}$ , and  $\overline{wage}_{jkl}$ , which I measure with data on migration, education, and wages. These equations provide the basis of estimation, and I add the following additional structure on education and migration costs.

$$1 + \tau_{jk}^e = (1 + S_j T_k)^{-\beta} \delta_j \delta_k (1 + \mathbf{C}_j T_k)^\phi, \quad (10a)$$

$$1 - \tau_{j\ell}^m = (1 + d_{j\ell}^P)^{-\omega_1} (1 + d_{j\ell}^D)^{-\omega_2} \quad (10b)$$

School construction  $S_j$  decreases education costs for treated cohorts ( $T_k = 1$ ), subject to origin- and cohort-specific factors  $\delta_j$  and  $\delta_k$  and controlling for confounders  $\mathbf{C}_j$ . This relationship maps counterfactual school construction onto education costs and thus onto outcomes. The underlying assumption is that school construction changes education costs but not other parameters, including amenities. Physical and demographic distances ( $d_{j\ell}^P, d_{j\ell}^D$ ) increase migration costs, which are zero for non-migrants and bilaterally symmetric for migrants by construction. Physical distance is Euclidean, and demographic distance captures (pre-INPRES) dissimilarity in religion and language.<sup>5</sup>

Equations 5, 6, and 7 contain summation terms  $\sum_{\hat{\ell}} \tilde{w}_{jk\hat{\ell}}^\theta$  that complicate estimation because they are mechanically correlated with the error terms. Differencing eliminates these terms.

$$\log \overline{educ}_{jkl} - \log \overline{wage}_{jkl} = \log \frac{\eta}{c} - \log(1 + \tau_{jk}^e) + \log(1 - \tau_{j\ell}^m) + \log \varepsilon_{jkl}^e, \quad (11a)$$

$$\Delta_\ell \log \overline{educ}_{jkl} = -\Delta_\ell \log \alpha_\ell - \Delta_\ell \log \left( \frac{\varepsilon_{jkl}^e}{1 - \varepsilon_{jkl}^e} \right), \quad (11b)$$

$$\Delta_\ell \log \overline{wage}_{jkl} = -\Delta_\ell \log \alpha_\ell - \Delta_\ell \log(1 - \tau_{j\ell}^m) - \Delta_\ell \log \left( \frac{1}{1 - \varepsilon_{jkl}^e} \right), \quad (11c)$$

$$\Delta_\ell \log \pi_{jkl} = \theta \Delta_\ell \log(1 - \tau_{j\ell}^m) + \theta \Delta_\ell \log(\alpha_\ell^{1-\eta} w_\ell) + \theta \Delta_\ell \log \tilde{\varepsilon}_{jkl}^e \quad (11d)$$

The first equation differences equations 6 and 7, and the last three equations difference with respect to a reference destination for  $\Delta_\ell \log X_{jkl} \equiv \log X_{jkl} - \log X_{jk0}$ .

Given an initial estimate of returns to schooling  $\eta$ , these moments identify education and migration costs, amenities, the Fréchet dispersion parameter, and base wages. Equation 11a identifies

<sup>5</sup> Physical distance captures differences in latitude and longitude. Demographic distance captures differences in Muslim share and Indonesian share in 1971.

education and migration costs. By equations 10a and 10b, it becomes

$$\begin{aligned} \log \overline{educ}_{jkl} - \log \overline{wage}_{jkl} &= \beta \log(1 + S_j T_k) - \log \delta_j - \log \delta_k - \phi \log(1 + C_j T_k) \\ &\quad - \omega_1 \log(1 + d_{j\ell}^P) - \omega_2 \log(1 + d_{j\ell}^D) + \log \frac{\eta}{c} + \log \varepsilon_{jkl}^e. \end{aligned}$$

School construction  $S_j$  is endogenous: it explicitly targeted regions with low enrollment, and thus is correlated with error  $\varepsilon_{jkl}^e$  in educational attainment. I address this endogeneity with the same difference-in-difference variation described in section 2. Rather than directly comparing districts with high and low levels of school construction, I instead compare how treated and untreated cohorts differ across such districts. Given migration costs, equations 11b and 11c identify relative amenities  $\frac{\alpha_\ell}{\alpha_0}$ , and equation 11d identifies Fréchet parameter  $\theta$ . Given relative amenities and returns to schooling, destination  $\ell$  fixed effects in equation 11d further identify relative base wages  $\frac{w_\ell}{w_0}$ . Base human capital  $h_{jk}$  and base cost of education  $c$  do not enter the counterfactuals of interest and thus do not need to be estimated. Note that although returns to schooling  $\eta$  are estimated separately, as I discuss below, they only affect estimates of base wages.

## 4.2 Estimation procedure

First, I use the difference-in-difference variation in school construction to obtain an estimate of returns to schooling  $\eta$ . As in Duflo (2001), the INPRES program generates exogenous variation in educational attainment that identifies the causal effect of education on wages. I consider the effect of log education on log wages given how returns to schooling are parameterized in the model, where wages are proportional to  $e^\eta$ . For individuals  $i$  of origin  $j$  and cohort  $k$ ,

$$\log \text{wage}_{ijk} = \delta_j + \delta_k + \eta \log \text{educ}_{ijk} + C_j T_k \phi + \varepsilon_{ijk}, \quad (12a)$$

$$\log \text{educ}_{ijk} = \delta_j + \delta_k + \beta S_j T_k + C_j T_k \phi + \varepsilon_{ijk}. \quad (12b)$$

I instrument for log education with the interaction of school construction and treatment cohort.

Second, I estimate the rest of the model by Poisson pseudo-maximum likelihood (PPML), as is common in spatial models (Santos Silva and Tenreyro 2006). Estimation in logs cannot accommodate zeros in observed migration probabilities, and taking logs is a non-linear transformation that introduces bias from heteroskedasticity. The PPML approach addresses both concerns. For a model  $y_i = \exp(x_i \beta) + \varepsilon_i$ , PPML uses the set of first-order conditions

$$\sum_{i=1}^n [y_i - \exp(x_i \hat{\beta})] x_i = 0$$

as the basis of estimation. I form these conditions for each of equations 11a, 11b, 11c, and 11d, and I use them as moments in a generalized method of moments (GMM) estimation framework.

**Table 3:** Returns to education  $\eta$ 

	Treatment			Placebo		
	OLS	IV	First stage	OLS	IV	First stage
Log years of schooling	0.393*** (0.00721)	0.688** (0.311)		0.394*** (0.00678)	-1.357 (3.523)	
INPRES $\times$ young			0.0284*** (0.00899)			0.00564 (0.0110)
Observations	89,404	89,404	89,404	55,091	55,091	55,091
F-statistic			9.97			0.26

Each column is one regression. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. The first three columns display treatment estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1974. The last three columns display placebo estimates that compare individuals ages 12 to 17 and those ages 18 to 24 in 1974. The outcome variable is log monthly wages, and the instrument for log years of schooling is the interaction of INPRES program intensity and treatment cohort. Regressions control for birth district fixed effects, birth year fixed effects, survey year fixed effects, 1971 child population, 1971 enrollment rates, and INPRES spending on water and sanitation projects. Standard errors are clustered by birth district, as defined by 1971 boundaries. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Finally, general equilibrium parameters  $(\sigma, \kappa, \lambda)$  for output, agglomeration, and congestion are set exogenously. I follow [Bryan and Morten \(2019\)](#) in taking  $(\sigma, \kappa, \lambda) = (8, 0.05, 0.075)$  as baseline values and in considering robustness to these choices. These parameters affect counterfactuals, but they affect neither estimation nor identification of other parameters.

### 4.3 Estimates

For returns to education  $\eta$ , [table 3](#) presents an IV estimate of 0.7 compared to an OLS estimate of 0.4. The IV estimate is larger than the OLS estimate, as is the case in [Duflo \(2001\)](#). There is a relatively strong first stage that indeed disappears in the placebo experiment. For the US, [Hsieh et al. \(2019\)](#) choose a value of 0.1 that corresponds to the fraction of output spent on human capital accumulation. They obtain this value by dividing education spending by the labor share of GDP. I take my IV estimate of  $\eta = 0.7$  as a baseline value, but also I consider robustness to the OLS value of 0.4 and the [Hsieh et al. \(2019\)](#) value of 0.1.

[Table 4](#) shows parameter estimates that correspond to the main objects of interest: education and migration costs. The  $\beta$  parameter captures the relationship between school construction and education costs. This parameter affects the role of diminishing marginal returns, as higher values imply higher initial gains from school construction. The estimate is positive and significant in the treatment sample, which is exposed to school construction, and insignificant in the placebo sample, which is not exposed. The  $\omega$  parameters capture the relationship between distance and migration costs. These parameters affect the role of market access, as higher values imply higher

**Table 4:** Education and migration costs

	Treatment		Placebo	
	Estimate	SE	Estimate	SE
$\beta$	0.110**	(0.0467)	0.0514	(0.0457)
$\omega_1$	0.0415***	(0.00353)	0.0388***	(0.00423)
$\omega_2$	0.0184	(0.0500)	-0.0299	(0.0658)

These parameters determine education and migration costs as given by equations 10a and 10b. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. The first two columns display treatment estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1974. The last two columns display placebo estimates that compare individuals ages 12 to 17 and those ages 18 to 24 in 1974. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

relative gains in districts with market access. The estimates suggest that migration costs are driven by physical distance, with demographic distance having little effect. The treatment and placebo groups produce similar estimates because distance measures do not differ for these groups.

## 5 Counterfactuals

This section estimates the general equilibrium effects of school construction on aggregate output. I evaluate the actual rule for allocating school construction across districts, as well as alternative rules with varying degrees of flexibility.

### 5.1 Aggregate output and budget constraint

Counterfactuals consider the effects of school construction on aggregate output, as defined in equation 8, subject to a budget constraint. For an allocation  $a$  within a space of allocations  $\mathcal{A}$ ,

$$\max_{a \in \mathcal{A}} \{Y(a)\} = \max_{a \in \mathcal{A}} \left\{ \left[ \sum_{\ell} \underbrace{\left( \sum_{j,k} N_{jk} \pi_{jk\ell}(a) \overline{wage}_{jk\ell}(a) \right)}_{Y_{\ell}(a)} \right]^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\sigma}} \}. \quad (13)$$

The budget constraint is the total number of schools constructed, as observed in the data.

$$\mathcal{A} = \left\{ a \mid \sum_{\ell} a_{\ell} = \sum_{\ell} a_{\ell}^{\text{obs}} \right\}$$

School construction costs vary spatially, and I account for this variation by subtracting cost increases (or adding cost savings) to aggregate output. The underlying assumption is that the government is not credit constrained, and therefore able to allocate more schools to expensive districts when the benefits to doing so are high. An alternative is to use cost data to compute total expenditures, then

to take total expenditures as the budget constraint. Indeed, the INPRES program specified these costs by district, ranging in 1973 from 2.5M IDR for non-urban districts in Sumatra, Java, Bali, and Kalimantan to 7M IDR for districts in Greater Jakarta. But because the law was specifically written to target a total number of schools and not total expenditures, I define the budget constraint in terms of the number of schools.

The problem simplifies in the absence of agglomeration. Productivity, base wages, amenities, and migration –  $A_\ell$ ,  $w_\ell$ ,  $\alpha_\ell$ , and  $\pi_{jk\ell}$  – are invariant in school construction when  $\kappa = 0$ . Thus, counterfactual aggregate output  $Y'_\ell$  in a given destination is a simple function of counterfactual school construction  $S'_j$ , parameter  $\beta$ , and observed quantities.

$$Y'_\ell = \sum_{j,k} N_{jk} \pi_{jk\ell} \overline{wage}_{jk\ell} \left( \frac{1 + S'_j T_k}{1 + S_j T_k} \right)^{\frac{\beta\eta}{1-\eta}} \quad (14)$$

Observed quantities proxy for the fundamentals of the model, as in the exact-hat algebra of [Dekle et al. \(2008\)](#). Estimating parameter  $\beta$  is sufficient for counterfactuals, and there is no need to estimate the rest of the model. There is also no need to specify congestion  $\lambda$  (having specified  $\eta$ ,  $\sigma$ , and  $\kappa$ ). The optimal allocation depends on the trade-off between diminishing marginal returns, via returns to schooling  $\eta$ , and market access, via migration probabilities  $\pi_{jk\ell}$ . Migration costs do not enter directly because migration probabilities do not change in response to school construction, but note that market access still determines the migration probabilities themselves.

The problem is further simplified when output is perfectly substitutable:  $\lim_{\sigma \rightarrow \infty} Y' = \sum_\ell Y'_\ell$ . Without either agglomeration or imperfect substitution, school construction is not interdependent across districts, and the optimization problem can be solved by sequentially allocating schools to the district with the highest marginal return. But when output is not perfectly substitutable ( $\sigma < \infty$ ), it generates spatial interdependence that demands joint optimization over the allocation of schools to each district. However,  $\frac{\beta\eta}{1-\eta} \leq 1$  is a sufficient condition for the concavity of the objective function, as I show in [appendix A.2](#). In this case, standard convex optimization techniques apply and allow the problem to be solved quickly nonetheless.

Agglomeration complicates the problem by causing productivity, base wages, amenities, and migration to respond to school construction. Thus, evaluating counterfactual school construction requires solving for how these quantities change in equilibrium. I do so with the following algorithm.

1. Given  $S'_j$ , compute  $H'_\ell = Y'_\ell / A'_\ell$  ignoring changes in  $w_\ell$ ,  $\alpha_\ell$ , and  $\pi_{jk\ell}$  (as if  $\kappa = 0$ ).
  - (a)  $A'_\ell = w_\ell$ , and  $Y'_\ell$  is given by [equation 14](#).
2. Given  $H'_\ell$ , compute  $w'_\ell = \bar{A}_\ell (H'_\ell)^\kappa = w_\ell (H'_\ell / H_\ell)^\kappa$  for  $H_\ell = Y_\ell / w_\ell$ .
3. Given  $w'_\ell$ , compute  $\pi'_{jk\ell}$  and  $\alpha'_\ell$ 
  - (a) Given  $w'_\ell$ , compute  $\pi'_{jk\ell} = (\tilde{w}'_{jk\ell})^\theta / \sum_{\hat{\ell}} (\tilde{w}'_{jk\hat{\ell}})^\theta$  for  $\tilde{w}'_{jk\ell} = (\alpha'_\ell)^{1-\eta} (1 - \tau_{j\ell}^m) w'_\ell \tilde{\varepsilon}_{jk\ell}^c$ .
  - (b) Given  $\pi'_{jk\ell}$ , compute  $\alpha'_\ell = \bar{\alpha}_\ell (\sum_{j,k} N_{jk} \pi'_{jk\ell})^{-\lambda}$ .

- (c) Repeat until convergence, updating  $\pi'_{jkl}$  based on  $\alpha'_\ell$ .
4. Given  $w'_\ell$ ,  $\alpha'_\ell$ ,  $\pi'_{jkl}$ , and  $S'_j$ , compute  $H'_\ell = Y'_\ell/A'_\ell$ .
- (a)  $A'_\ell = w'_\ell$  and

$$Y'_\ell = \sum_{j,k} N_{jk} \pi'_{j\ell} \overline{wage}_{jkl} \left( \frac{1 + S'_j T_k}{1 + S_j T_k} \right)^{\frac{\beta\eta}{1-\eta}} \left( \frac{w'_\ell}{w_\ell} \right)^{\frac{1}{1-\eta}} \left( \frac{\pi'_{jkl}}{\pi_{jkl}} \right)^{\frac{1}{\theta(1-\eta)}}. \quad (15)$$

5. Repeat (2) to (4) until convergence to obtain  $Y'_\ell$ .

I compute aggregate output with equation 8. Agglomeration generates spatial interdependence that demands joint optimization over the allocation of schools to each district, regardless of the substitutability of output across districts. Note that relative base wages and relative amenities, as estimated previously, are sufficient for computing counterfactuals. Base wages enter aggregate output through  $Y_\ell$  (equation 15), and base wages and amenities enter migration probabilities through  $\tilde{w}_{jkl}$  (equation 5). The normalizations cancel in both cases, so there is no loss from not estimating base wages and amenities in levels.

More broadly, my focus here is on maximizing aggregate output, but there exists an objective function for which the actual allocation is optimal. Political concerns are one potential factor, and I pursue this line of inquiry in related work on healthcare infrastructure in Indonesia (Hsiao 2021). Political concerns are less relevant in this setting because the actual allocation was determined by a specific formula, but equity concerns may explain why districts with high unenrollment were not targeted more aggressively. Indeed, equity concerns likely contributed to the actual allocation rule in later years that assigned zero new schools to districts with enrollment above the 85% cutoff. However, note that equity concerns are already captured in part by diminishing marginal returns, which also encourage investment in underdeveloped areas. Furthermore, equity concerns would result in more schools in low-enrollment districts relative to the output-maximizing allocation, but my results suggest the opposite.

## 5.2 Allocation rules

I quantify how well alternative allocation rules do in approximating the output-maximizing allocation. Simple rules are common in practice because they are more transparent and much simpler to implement than optimizing over a full spatial equilibrium model. In this context, simple rules can still do well in targeting districts with few existing schools and high market access, but they may not fully capture spatial interdependence across districts.

I study two kinds of simple allocation rules. First, I consider rules based solely on unenrollment as in the actual rule. I consider a linearly proportional rule, but relative to the actual rule I allow for a parameter governing the slope of the rule. For example, while the actual rule imposes that districts with twice as much unenrollment receive twice as many schools, I allow this ratio to vary and choose it optimally. I also consider more flexible functional forms: quadratic rules and rules

**Table 5:** Counterfactual effects on aggregate output

Allocation	Agglomeration ( $\kappa$ )			
	0	0.025	0.05	0.075
Actual	1.00	1.00	1.00	1.00
None	0.96	0.95	0.93	0.90
Uniform	0.99	0.99	0.98	0.96
Optimal	1.02	1.03	1.04	1.07

Each cell is one counterfactual scenario for aggregate output. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. Columns vary the strength of agglomeration parameter  $\kappa$ . Rows show different allocations. The first row shows the actual allocation observed in the data, and I normalize these values to one. The second row shows aggregate output in the absence of INPRES school construction. The third row allocates all districts the same number of new schools per child. The fourth row maximizes aggregate output net of cost differences. All allocations are subject to a budget constraint defined by the total number of schools constructed in the data.

with cutoffs. The actual rule from 1975 to 1978 features a cutoff in which districts with less than 15% unenrollment are allocated zero new schools.

$$\begin{aligned} \mathcal{A}^{\text{lin}} &= \mathcal{A} \cap \{a \mid a_\ell = \rho X_\ell\}, & \mathcal{A}^{\text{lin, cut}} &= \mathcal{A} \cap \{a \mid a_\ell = \rho X_\ell \cdot \mathbb{1}[X_\ell > c]\}, \\ \mathcal{A}^{\text{quad}} &= \mathcal{A} \cap \{a \mid a_\ell = \rho X_\ell + \rho_2 X_\ell^2\}, & \mathcal{A}^{\text{quad, cut}} &= \mathcal{A} \cap \{a \mid a_\ell = (\rho X_\ell + \rho_2 X_\ell^2) \cdot \mathbb{1}[X_\ell > c]\}, \end{aligned}$$

Second, I expand these rules to account for additional covariates – namely, ruralness and distance to the nearest urban district. Note that each of these rules simplifies the optimization problem by further constraining the choice space: searching over  $\mathcal{A}$  requires optimizing over school construction  $a_\ell$  in each location, while searching over  $\mathcal{A}^{\text{lin}}$  only requires optimizing over a single parameter  $\rho$ .

While more sophisticated rules can be more effective, they may also be less robust. I consider how three potential sources of error undermine each allocation rule. First, implementation error captures deviation from the rule. For each district, I draw an allocation from a uniform distribution ranging from 25% below to 25% above the allocation implied by the rule. Second, I suppose that the parameters of each counterfactual rule are misspecified (or misestimated) relative to their optimal values. I simulate errors drawn from a uniform distribution ranging from 25% below to 25% above the optimal values. Third, I study mismeasurement in the data used to determine allocations, such as unenrollment rates, drawing measurement error again from a uniform distribution ranging from 25% below to 25% above each actual value.

### 5.3 Results

Table 5 shows the impact of school construction on aggregate output under agglomeration of varying strength. In the baseline, I assume agglomeration  $\kappa = 0.05$ . Aggregate output falls by seven percent when school construction is eliminated entirely, indicating meaningful long-term

**Table 6:** Counterfactual effects on aggregate output, alternative rules


Allocation rule	Proportional rule				Cutoff rule			
	$\kappa = 0$	0.025	0.05	0.075	$\kappa = 0$	0.025	0.05	0.075
Unenrollment								
Linear	1.01	1.01	1.01	1.02	1.01	1.01	1.01	1.02
Quadratic	1.01	1.01	1.02	1.03	1.01	1.01	1.02	1.04
Unenrollment + ruralness								
Linear	1.01	1.01	1.02	1.02	1.01	1.01	1.02	1.03
Quadratic	1.01	1.02	1.02	1.03	1.01	1.02	1.03	1.04
Unenrollment + urban distance								
Linear	1.02	1.02	1.03	1.04	1.02	1.02	1.03	1.05
Quadratic	1.02	1.03	1.04	1.06	1.02	1.03	1.04	1.07

Each cell is one counterfactual scenario for aggregate output. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. Columns vary the strength of agglomeration parameter  $\kappa$ . Each column is relative to aggregate output under the actual allocation, which corresponds to a value of one, given agglomeration in that column. The first four columns involve proportional rules without cutoffs, while the last four columns allow for optimally chosen cutoffs. Rows show allocation rules of varying forms, with the parameters of each allocation rule chosen optimally. The first panel shows linear and quadratic rules based only on unenrollment, and the second and third panels allow rules to condition on ruralness and urban distance. Urban distance is Euclidean distance to the nearest urban area. All allocations are subject to a budget constraint defined by the total number of schools constructed in the data.

effects of the program in general equilibrium. The actual allocation was targeted, allocating more schools to districts with higher levels of unenrollment. This targeting was warranted: subject to the budget constraint, spreading resources uniformly across districts would have diminished the effect of the program by two percentage points. However, the targeting could have gone farther, and indeed the actual allocation rule did not greatly outperform the uniform allocation. Under the optimal allocation, which I compute unstrained by any allocation rule, I find that aggregate output would have increased by an additional four percentage points. These effects are muted when agglomeration is weak because in this case there are lower returns to targeting specific labor markets. For the same reason, there are smaller losses in moving from the actual to the uniform allocation. As a result, the counterfactuals under weak agglomeration are well approximated by the zero-agglomeration case ( $\kappa = 0$ ), which takes a much simpler form and can be computed without estimating the full model.

Table 6 shows how well simple allocation rules perform in approximating the optimal allocation. The actual allocation was based on a linear cutoff rule that conditioned only on enrollment. Optimizing over the parameters of this rule – in particular, the slope governing how quickly the allocation increases in unenrollment – increases the effect on aggregate output by another one to two percentage points. Adopting a more flexible quadratic rule boosts the effect further. But the largest gains come from conditioning on additional observables. Ruralness is a rough proxy for

**Figure 3:** Pre-INPRES unenrollment vs. existing schools and market access



Each figure is a binned scatter plot, and each observation is one district as defined by 1971 boundaries. The left figure compares 1971 unenrollment rates and the number of existing, pre-INPRES primary schools per 1,000 children (relative to the mean). It controls for 1971 market access and population. The right figure compares 1971 unenrollment rates and 1971 market access (relative to the mean). It controls for the number of existing schools and 1971 population. Market access is the sum of a district's own population in 1971 and an inverse-distance-weighted sum of other districts' populations.

market access that is already commonly considered in regional planning. Distance to the nearest urban area is a better proxy, and indeed allocation rules that incorporate urban distance perform substantially better than those that consider unenrollment alone. This difference is magnified when agglomeration is strong, as agglomeration increases both the returns to targeting and the strength of the spatial interdependence ignored by simple rules.

Intuitively, unenrollment and urban distance together capture the two forces that drive the trade-off in the full model: diminishing marginal returns and market access. The former favors spreading resources across districts to capitalize on high initial returns to school construction. The latter favors concentrating resources in districts with high market access, where the barriers to migration are low and thus the incentives to invest in schooling are high. Navigating this trade-off requires targeting districts with both high initial returns and high market access. In this sense, allocation rules based on unenrollment alone are insufficient. Figure 3 shows that unenrollment only reflects the number of existing schools, and thus the marginal return to new schools. Unenrollment is not indicative of market access.

While more sophisticated allocation rules can improve targeting, table 7 shows that they also come at a cost. The benefit of simple rules is that they are more robust to implementation, misspecification, and mismeasurement errors. That is, issues of overfitting can arise for rules with higher-order nonlinearities that conditional on many covariates. Overfitting is a concern because all allocation rules are approximations: even those that account for diminishing marginal returns and market access still abstract from the full spatial interdependence of the problem. This abstraction is what makes such rules straightforward to implement in the first place. Sophisticated rules incur

**Table 7:** Counterfactual effects on aggregate output, robustness

Allocation rule	Agglomeration ( $\kappa$ )			
	0	0.025	0.05	0.075
Implementation error				
Unenrollment, linear	1.00	1.00	0.99	0.98
Unenrollment + urban distance, quadratic	0.99	0.98	0.97	0.96
Misspecification error				
Unenrollment, linear	1.00	0.99	0.97	0.95
Unenrollment + urban distance, quadratic	0.98	0.97	0.96	0.93
Mismeasurement error				
Unenrollment, linear	0.99	0.98	0.96	0.93
Unenrollment + urban distance, quadratic	0.98	0.97	0.95	0.92

Each cell is one counterfactual scenario for aggregate output. Data come from SUSENAS 2011, 2012, 2013, and 2014 and focus on male heads of household. Columns vary the strength of agglomeration parameter  $\kappa$ . Each column is relative to aggregate output under the actual allocation, which corresponds to a value of one, given agglomeration in that column. Panels show allocation rules of two forms subject to different kinds of error. The first panel adds implementation error in which each district's allocation of new schools differs by up to 25% of that implied by the rule. The second panel adds misspecification error in which the parameters of the rule differ by the 25% of their optimal values. The third panel adds mismeasurement error in which the data used to compute allocations differs by up to 25% from the true values. Urban distance is Euclidean distance to the nearest urban area. All allocations are subject to a budget constraint defined by the total number of schools constructed in the data.

the greatest losses under misspecification and mismeasurement, as they have more parameters to be misspecified and more inputs to be mismeasured. Losses are also increasing in agglomeration, which amplifies the gains from precise targeting. More broadly, overly sophisticated rules can backfire. The outcome in many cases is worse than that of the uniform allocation, and in some cases only somewhat better than constructing no schools at all.

## 6 Conclusion

Spatial effects make optimal infrastructure investment a difficult problem to solve. In practice, however, policymakers often allocate infrastructure investments according to simple rules based on a small set of observable characteristics. This paper quantifies the effectiveness of these simple rules and provides guidance on how to make them better. I do so in the context of one of the largest school construction efforts in history. Indonesia built 62,000 primary schools in the mid-1970s, allocating schools across districts with a simple rule that prioritized regions with low rates of school enrollment. I use a spatial equilibrium model to quantify the effects of the program, compute the optimal allocation, and evaluate a range of allocation rules.

My baseline estimates suggest that the school construction program increased aggregate output by seven percent in the long run, but that the optimal allocation would have increased aggregate

output by another four percentage points. The optimal allocation navigates the two key forces in the model – diminishing marginal returns and market access – while also accounting for migration in equilibrium that generates spatial interdependence in the allocation decision. Allocation rules can capture diminishing marginal returns and market access by conditioning on enrollment rates and distance to the nearest city with flexible functional forms. But while more sophisticated rules better approximate the optimal allocation, simpler rules are more robust to unintended deviations arising from implementation error, misspecification, and mismeasurement. All allocation rules abstract from spatial interdependence, and these errors compound the loss from doing so.

Several lines of inquiry are left for future work. First, future work might quantify the complementary effects of joint investment in schools and roads. Such an approach may be valuable given the interaction between education and migration costs. Second, I assume that school construction lowers education costs by increasing physical access, but the effects of new schools might also depend on factors like school quality and interactions with existing schools. Third, I find the observed allocation to be suboptimal from the perspective of an output-maximizing social planner, but there exists an objective function that rationalizes the observed allocation. One could also treat observed choices as optimal and consider what they reveal about government preferences.

## References

- Allen, Treb and Costas Arkolakis. The Welfare Effects of Transportation Infrastructure Improvements. 2019.
- Asher, Sam and Paul Novosad. Rural Roads and Local Economic Development. American Economic Review3, 110(3):797–823, 2020.
- Balboni, Clare. In Harm’s Way? Infrastructure Investments and the Persistence of Coastal Cities. 2019.
- Bazzi, Samuel, Masyhur Hilmy, and Benjamin Marx. Islam and the State: Religious Education in the Age of Mass Schooling. 2021.
- Bryan, Gharad and Melanie Morten. The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. Journal of Political Economy, 127(5):2229–2268, 2019.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum. Global Rebalancing with Gravity: Measuring the Burden of Adjustment. IMF Staff Papers, 55(3):511–540, 2008.
- Diamond, Rebecca. The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980?2000. American Economic Review, 106(3):479–524, 2016.
- Donaldson, Dave. Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. American Economic Review, 108(4-5):899–934, 2018.
- Donaldson, Dave and Richard Hornbeck. Railroads and American Economic Growth: A ”Market Access” Approach. Quarterly Journal of Economics, 131(2):799–858, 2016.
- Duflo, Esther. Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. American Economic Review, 91(4):795–813, 2001.
- Duflo, Esther. The Medium Run Effects of Educational Expansion: Evidence from a Large School Construction Program in Indonesia. Journal of Development Economics, 74(1):163–197, 2004.
- Eaton, Jonathan and Samuel Kortum. Technology, Geography, and Trade. Econometrica, 70(5): 1741–1779, 2002.
- Faber, Benjamin. Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System. The Review of Economic Studies, 81(3):1046–1070, 2014.
- Fajgelbaum, Pablo and Edouard Schaal. Optimal Transport Networks in Spatial Equilibrium. Econometrica, 88(4):1411–1452, 2020.
- Gertler, Paul, Marco Gonzalez-Navarro, Tadeja Gračner, and Alexander Rothenberg. Road Quality, Local Economic Activity, and Welfare: Evidence from Indonesia’s Highways. 2019.
- Hsiao, Allan. Democratization and Infrastructure Investment: Evidence from Healthcare in Indonesia. 2021.
- Hsieh, Chang-Tai, Erik Hurst, Charles Jones, and Peter Klenow. The Allocation of Talent and U.S. Economic Growth. Econometrica, 87(5):1439–1474, 2019.

Korte, Bernhard and Jens Vygen. Facility Location. In Combinatorial Optimization: Theory and Algorithms, pages 563–598. Springer, 2008.

Martinez-Bravo, Monica. The Local Political Economy Effects of School Construction in Indonesia. American Economic Journal: Applied Economics, 9(2):256–289, 2017.

McFadden, Daniel. Conditional Logit Analysis of Qualitative Choice Behavior. In Zarembka, Paul, editor, Frontiers in Econometrics, volume 1, chapter 4, pages 105–142. Academic Press, 1974.

Santos Silva, J.M.C. and Silvana Tenreyro. The Log of Gravity. The Review of Economics and Statistics, 88(4):641–658, 2006.

Tsivanidis, Nick. Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota’s TransMilenio. 2019.

## A Appendix


### A.1 Data and Stylized Facts

**Table A1:** Placebo effects of INPRES school construction

	SUPAS 1976			SUPAS 1995		
	Estimate	SE	Obs	Estimate	SE	Obs
Years of schooling	-0.0175	(0.0703)	18,173	0.0411	(0.0470)	64,392
Years of schooling	0.169	(0.164)	6,461	-0.0185	(0.0765)	25,159
Log monthly wages	0.00219	(0.0280)	6,461	0.000511	(0.00808)	25,159
Primary school completion	-0.0489	(0.0479)	18,061	0.00535	(0.0252)	64,392
Middle school completion	-0.01000	(0.0535)	18,135	0.0213	(0.0219)	64,392
High school completion	0.103	(0.0738)	17,838	0.0408	(0.0287)	64,392
University completion	0.239	(0.151)	12,598	-0.0357	(0.0520)	63,828
Employment	0.0900	(0.119)	16,135	-0.0125	(0.0487)	68,595
Wage employment	-0.000908	(0.0526)	18,098	0.00264	(0.0171)	69,114
Self-employment	0.00755	(0.0568)	18,067	-0.00622	(0.0180)	69,114
Weekly hours	-0.250	(0.270)	16,903	0.0277	(0.128)	66,423
Migrant	0.0480	(0.0590)	16,860	-0.0224	(0.0232)	69,114
Migrant to Jakarta	0.250	(0.163)	11,993	0.0337	(0.0585)	56,750

Each row is two placebo regressions. Data come from SUSENAS 1976 and 1995 and focus on male heads of household. The first three columns display placebo estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1955. The last three columns display placebo estimates that compare individuals ages 2 to 6 and those ages 12 to 17 in 1974. Row headings list outcome variables. I run logit regressions for dummy outcomes – all outcomes except for years of schooling, log monthly wages, and weekly hours. Regressions control for birth district fixed effects, birth year fixed effects, survey year fixed effects, 1971 child population, 1971 enrollment rates, and INPRES spending on water and sanitation projects. Standard errors are clustered by birth district, as defined by 1971 boundaries. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

**Figure A1:** Existing schools vs. market access



Each observation is one district as defined by 1971 boundaries. The figure compares the number of existing, pre-INPRES primary schools per 1,000 children to 1971 market access. Market access is the sum of a district's own population in 1971 and an inverse-distance-weighted sum of other districts' populations. The figure controls for the 1971 child population, 1971 enrollment rates, and INPRES spending on water and sanitation projects. I omit outliers by dropping the 1% largest and smallest values for the number of existing schools and market access.

## A.2 Counterfactuals

**Claim.** If  $\kappa = 0$  and  $\frac{\beta\eta}{1-\eta} \leq 1$ , then aggregate output  $Y$  is concave in school allocation  $a$ .

**Proof.** Consider aggregate output when  $\kappa = 0$ . Defining  $X_{jkl} \equiv N_{jk}\pi_{jkl}\overline{wage}_{jkl}(1 + S_j^{\text{obs}}T_k)^{-\frac{\beta\eta}{1-\eta}}$  and given  $S_j(a) = a_j$ ,

$$Y(a) = \left[ \sum_{\ell} (Y_{\ell}(a))^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1-\sigma}} \quad \text{for} \quad Y_{\ell}(a) = \sum_{j,k} X_{jkl}(1 + a_j T_k)^{\frac{\beta\eta}{1-\eta}}$$

If  $\frac{\beta\eta}{1-\eta} \leq 1$ , then  $Y_{\ell}$  is concave in  $a$  for all  $\ell$ . Furthermore,  $Y$  is (strictly) concave in  $\{Y_{\ell}\}$ . Using shorthand  $Y(Y_{\ell}) \equiv Y(Y_1, \dots, Y_L)$  and noting that  $\sigma > 1$ ,

$$g(Y_{\ell}) = \sum_{\ell} (Y_{\ell}(a))^{\frac{\sigma-1}{\sigma}}$$

is the sum of concave functions and therefore concave itself. It follows that

$$Y(Y_{\ell}) = h(g(Y_{\ell})) \quad \text{for} \quad h(x) = x^{\frac{\sigma}{1-\sigma}}$$

is quasi-concave because  $h(\cdot)$  is an increasing function. The quasi-concavity and homogeneity of  $Y$  together imply that it is concave. To see this, note that  $Y$  exhibits homogeneity of degree one,

which implies

$$Y(tY_\ell) = tY(Y_\ell) \quad \Rightarrow \quad Y\left(\frac{Y_\ell}{Y(Y_\ell)}\right) = 1.$$

Quasi-concavity implies

$$f\left(\tilde{c}\frac{Y'_\ell}{Y(Y'_\ell)} + (1 - \tilde{c})\frac{Y''_\ell}{Y(Y''_\ell)}\right) > \min\left\{Y\left(\frac{Y'_\ell}{Y(Y'_\ell)}\right), Y\left(\frac{Y''_\ell}{Y(Y''_\ell)}\right)\right\} = 1$$

for all  $\tilde{c}$ . Applying  $\tilde{c} = \frac{cY(Y'_\ell)}{cY(Y'_\ell) + (1-c)Y(Y''_\ell)}$  and homogeneity,

$$Y\left(\frac{cY'_\ell + (1-c)Y''_\ell}{cY(Y'_\ell) + (1-c)Y(Y''_\ell)}\right) > 1 \quad \Rightarrow \quad Y(cY'_\ell + (1-c)Y''_\ell) > cY(Y'_\ell) + (1-c)Y(Y''_\ell),$$

Finally,  $Y(Y_\ell(a))$  is a composition of concave functions and thus concave in  $a$ .