Export agriculture and rural poverty: evidence from Indonesian palm oil

Ryan B. Edwards*
Dartmouth College

March 25, 2019

Abstract

This paper measures the impacts of the world's largest modern agricultural export expansion—that of Indonesian palm oil since 2000—on poverty and consumption in producing communities. Identification exploits external demand growth and geographic differences in cultivation suitability in a difference-in-difference instrumental variables design. The main finding is that growth in palm oil sector lifted up to 2.6 million rural Indonesians from poverty this century. The median expansion led to 2.7 percentage points faster poverty reduction and 4 percent faster consumption growth. Divergent regional development trajectories can be explained by rising returns to labor and land, and indirect effects through increased household investment, government revenue, and rural economic and social infrastructure in producing communities.

*Contact address: Department of Economics, 6106 Rockefeller Hall, Hanover, NH 03755 USA. Email: ryan.b.edwards@dartmouth.edu. I thank Eric Edmonds, Paul Burke, Lisa Cameron, Wally Falcon, Doug Gollin, Grace Hadiwidjaja, Will Johnson, Naureen Karachiwalla, Erzo Luttmer, Roz Naylor, Paul Novosad, Budy Resosudarmo, Na'ama Shenhav, Chris Snyder, Robert Sparrow, Doug Staiger, Sudarno Sumarto, Daniel Suryadarma, and seminar participants at the Vatican, Dartmouth, Stanford, Duke, ANU, QUB, Monash, CIFOR, Bank Indonesia, World Bank, NEUDC Cornell, and the 2017 AEA Meetings for helpful comments and discussions. I also thank the ANU Indonesia Project, Stanford SEED, Wahli, Wilmar, Asian Agri, and WWF-Indonesia for field assistance, Victoria Meyer for phenomenal research assistance under the Dartmouth Presidential Scholars program, and Mayang Rizky and Ridho Al Izzati at the SMERU Research Institute for generously compiling and sharing recent poverty data. Any errors are of course my own.

1 Introduction

While growth in trade has been shown to increase incomes and reduce poverty in poor countries in a wide variety of contexts, agricultural export growth is more controversial. Several studies argue that globalization of agriculture discourages structural transformation, leaving areas induced to specialize in agriculture worse off (Mokyr, 1976; Krugman, 1987; Matsuyama, 1992). Others highlight a lack of price pass-through to the farm gate, because of market power in distribution networks and surplus labor on the farm (de Janvry, Fafchamps, and Sadoulet 1991; Fafchamps, Hill, Kauda, and Nsibirwa, 2001). The view that export agriculture—especially when involving large commerical farms—is unhelpful for the poor remains widely held (Byerlee, de Janvry, and Sadoulet, 2009; Easterly, 2007; Engerman and Solokoff, 2002; Bhagwati, 1958; Carter, Barham, and Mesbah 1996; Barham et al. 1992). Yet, there is limited evidence on how modern agricultural export growth affects poverty and the distribution of income within countries.

This study examines the impact of Indonesia's palm oil expansion on poverty and household consumption in rural communities that produce palm oil. Palm oil is the world's leading vegetable oil, found in around half of the products in supermarkets and almost exclusively grown in developing countries. Indonesia's four-fold increase in production since 2000 is the world's largest modern agricultural expansion. Oil palms cover around 7 percent of Indonesia's 1.9 million square kilometers of land area.[1] The view that palm oil is not only harmful for the environment, but also the economy and society, is common. Coalitions of activists are mobilized around the world arguing that palm oil production is environmentally and socially damaging and should be limited through government policy or consumer boycotts. In response to these concerns, the World Bank placed a moratorium on palm oil related investments in 2009. The European Parliament voted to ban palm oil imports for biofuels in 2017.

[1]The area under cultivation for palm continues to increase even though the price has declined since 2011. In this sense, this paper is not about a [price] boom per se but a sustained sectoral expansion.

Examining the impact of agricultural growth on poverty is complicated because agricultural output depends on a production process that will depend on correlates of poverty, and because farm gate prices are apt to be correlated with local demand. An ideal natural experiment might leverage an external shock and some plausibly exogenous geographic characteristics affecting the distribution of agricultural activity across space. Indonesia's expansion provides a useful approximation of this experimental ideal. Since 85% of Indonesia's palm oil is exported, the relevant demand is outside growing communities. To address endogeneity in production, I take advantage of recent growth in global palm oil demand—principally from China, India, and other emerging economies—coupled with the fact that regions differ in their productive potential. District area expansion, in a difference in difference framework, is instrumented with its average agro-climatically attainable palm oil yield interacted with the demand shock. Hence, I examine changes in poverty and household consumption over time across districts that vary in cultivation intensity due to their potential rather than actual production.

The main finding is that increased palm cultivation delivered strong poverty reduction and broad consumption gains for producing regions. A ten percentage point increase in the share of district area under cultivation for palm oil corresponds to an additional 5.36 percentage point poverty reduction and eight percent faster consumption growth relative to districts that increased cultivation less or not at all. The median expansion was five percent of district area. Relative gains were strongest for the bottom 20–60% and I find no evidence of urban households becoming worse off. Magnitudes are economically significant. With national poverty declining from 18.2% to 11.2% from 2002–2015, a non-trivial portion of Indonesia's regional development performance can be explained by palm cultivation. A simple policy simulation suggests that this unique episode of export growth accounts for up to 2.6 million of the ten million Indonesians lifted from poverty this century.

I trace the declines in poverty to direct and indirect mechanisms. Since most of the increase in production has come through cultivation area rather than yield increases, a first-order question is whether the impact is simply due to expanding the agricultural frontier. I find that the poverty impacts of increasing the share of farmland under cultivation for oil palm are similar to those using total area as the denominator. Poverty reductions are not only due to expansion onto marginal lands, but parallel changes in agricultural production increasing returns to labor and land: changing crops or practices. Higher labor productivity in agriculture and manufacturing, higher agricultural wages, and larger elasticities for more labor-intensive smallholder cultivation support this interpretation.

Three indirect channels reinforce the direct labor income gains. First, rising household expenditures are concentrated on health, education, and durable good expenditures, which correspond to more household assets and floorspace. I interpret these changes as household-level capital accumulation, a classic theoretical channel linking agricultural productivity to economic development only recently finding empirical support (Johnson and Mellor, 1961; Bustos, Caprettini, and Ponticelli, 2018; Marden, 2018). A second indirect channel relates to local governments, with revenue and expenditure growing faster in expanding districts. Since demand for public services is likely lower with rising consumption and falling poverty, fiscal windfalls may be directed to more productive public investments and amplify regional inequalities, as Caselli and Michaels (2013) find in Brazil and Feler and Senses (2017) in the United States. Palm expansion led to improved electrification, increased use of modern cooking fuels, and more marketplaces, schools, health clinics, and places of worship. Complementary economic and social infrastructure could allow economic returns to ratchet up over time through further market integration, as Donaldson (2018) finds for colonial India and Dell and Olken (2018) for Dutch sugar processing on Java. Although my study does not speak to impacts beyond my 15 year horizon, the health, education, and infrastructure investments here and fertility effects documented in Kubitza and Gehrke (2018) together suggest at least some positive long-run impacts.

This study contributes to four main streams of economics. First, I add to the trade literature new evidence on the distributional impacts of agricultural export growth (Castilho et al., 2012; Hasan et al, 2012; Topalova, 2007; Edmonds et al, 2010; Edmonds and Pavcnik, 2006; Topalova, 2010; Kis-Katos and Sparrow, 2011, 2015; Dix-Carneiro and Kovak, 2016; see Goldberg and Pavcnik (2007) for a review). The traditional view is that agricultural exports are famously short-lived, driven by external capital, and environmentally catastrophic, generating cash for politically-connected industrialists while depriving the poor of their land (Engerman and Solokoff, 2002; Easterly, 2007; World Bank, 2008). Far from adverse effects, the positive impacts of export market access that I find here are more consistent with the work of McCaig (2011) on provincial poverty in Vietnam, Balat et al. (2009) on Ugandan agricultural exports, and Costa, Garret, and Pessoa (2016) exploiting the same demand shock to study recent Brazilian export growth.

My study also contributes novel causal evidence to a classic question in development economics: the extent to which poverty alleviation, non-agricultural growth, and economic development can be driven by changes in agricultural productivity (see Gollin (2010) and Dercon and Gollin (2014) for recent reviews). In many ways, Indonesia epitomizes the sweeping changes in the global food system over the past few decades, with globalized supply chains (de Zegher, Iancu, and Plambeck, 2018), highly integrated smallholder-plantation systems (Hayami, 2010; Bellemare and Bloem, 2018), and unprecedented land expansion (Byerlee, Falcon, and Naylor, 2016). The most closely related study to mine in this regard is Bustos, Caprettini, and Ponticelli (2016), studying the recent expansion of soy in Brazil. Soy is also grown by both small and large farms, processed, and exported. Comparing sectoral employment, wages, and productivity across regions, Bustos, Caprettini, and Ponticelli (2016) show that the soy expansion led to non-agricultural productivity growth and structural change. I complement this work by measuring impacts on poverty and consumption in local communities where these controversial oilseeds are grown.

Third, I add new evidence to the literature on the local impacts of natural resources and demand shocks. Much of this literature focuses on extractive industries (Bound and Holzer, 2000; Feyrer, Mansur, and Sacerdote, 2017; Allcott and Keniston, 2017; Hornbeck and Keskin, 2015; see Cust and Poelhekke (2015) for a recent review). For example, Aragon and Rud (2013) highlight the importance of backward (i.e., input, upstream) linkages in shaping the local labor market impacts of a large gold mine in Peru. Plantation-based cash crops are similar in their processing, infrastructure, and backward linkage requirements but different in upstream labor intensity and geographic concentration.[2] In this sense, this paper closely relates to a stream of work emphasizing the importance of the factor intensities in mediating the impacts of natural resource sectors on local economic outcomes (Dal Bo and Dal Bo, 2013; Dube and Vargas, 2013; Edwards, 2016).

The fourth major literature this paper relates to is that on poverty-environment trade-offs (see Dasgupta, Laplante, Wang and Wheeler (2002) and Greenstone and Jack (2015) for reviews). Academic and public debates on palm oil tend to focus on the sector's often catastrophic environmental impacts. Here, I ask whether local communities benefit, thus helping us better understand the potential trade-offs. Although a full cost-benefit analysis would need to account for much more than just poverty and deforestation, I calculate that each percentage point of palm-driven poverty reduction corresponds to between 1.5 and 3 percent of district area lost in tree cover since 2000. However, this is not to say that such environmental impacts were unavoidable. Oil palm plantations account for around 20% of deforestation since 2001, deforestation accounts for around 20% of new oil palm plantations, and my results suggest that the poverty reduction is mostly driven by intensive margin changes (Austin et al. 2019).[3] Future growth could preserve the gains and minimize environmental damages by focusing more on these margins.

[2] The perennial nature of the crop and sustained increase in demand is also dissimilar to more volatile commodity prices and the different phases in mining life cycles.

[3] Gaveau et al. (2016) also estimate that 55% of the new industrial oil palm plantations on Kalimantan from 2005–2015 were developed on land that lacked forest cover for at least the five preceding years. This share tends to be larger for smallholders, who lack the capital to clear cut.

2 Indonesia's palm oil expansion

Palm oil is derived from the pulp of the fruit of the oil palm, a labor-intensive cash crop which requires little skill or capital to grow and harvest. Harvesting involves pulling fresh fruit bunches from trees with a long sickle. Oil palms bear a relatively consistent amount of fruit around every ten days with limited seasonality, offering a more frequent and predictable income stream than most alternative crops (Corley and Tinker, 2015). The largest costs are land acquisition and capital-intensive processing factories, which must receive fruit within 24 hours after harvest to be marketable to global markets. Yielding more oil per hectare than any other crop (i.e., 4–10 times that of other oilseeds), oil palm is one of the most economically attractive uses for land in the tropics. Sustained growth in emerging economies—particularly the "China shock"—increased global demand from less than 5 million metric tonnes per year in 1970 to over 70 million in 2015 (Autor, Dorn, and Hanson, 2013; Naylor, Higgins, Edwards, and Falcon, 2018). Demand is expected to double again over the next decade (USDA, 2016).

Indonesia accounted for more than 55 per cent of the 65 million metric tons of palm oil produced globally in 2017 (Directorate General of Estate Crops, 2017). Production increased from five to over forty million metric tons from 1997–2017. Palm oil has been Indonesia's largest agricultural export for the last two decades. While an established agroindustrial sector and market proximity positioned Indonesia well to take advantage of rising demand, the devalued rupiah from the Asian financial crisis and subsequent decentralization reforms precipitated the take-off (Rada, Buccola, and Fuglie, 2011; Burgess, Hansen, Olken, Potapov, and Sieber, 2012; Edwards, Falcon, Higgins, and Naylor, 2019).

Indonesia's dramatic increase in palm oil production has come almost exclusively through land area expansion, comprising both (a) farmers shifting crops on existing farmland (i.e., intensive margin changes within agriculture), and (b) new farmland from scrub, degraded land, or forest—that is, expanding the agricultural frontier (extensive margin).[4] The total area under cultivation for oil palm increased from 2.9 million hectares in 1997 to over 12.5 million today, around 7 percent of Indonesia's land area. Over two million hectares of the new oil palm plantations are estimated to have come through deforestation (Austin et al. 2019). Figure 1 illustrates the break down of this growth across sectors: private sector plantation area doubled, state-owned plantation area remained static, and the area managed by small, family farmers tripled. The expansion thus involves parallel growth in industrial and smallholder farms.[5] Smallholder farms are around two hectares each—sometimes managed in partnership with large estates but more commonly by independent farmers—and account for over 40% of the area planted today.[6] The scale of area expansion and rapid smallholder growth was in no small part enabled by the devolution of power, resources, and responsibilities to local governments. Decentralization liberalized land use, allowing local leaders to issue permits for new industrial estates and smallholders to expand their farms with little more than a letter or nod from the village head (Naylor et al. 2018).

Lags governing expansion and impacts motivate a long difference approach. The process from planting to exporting is characterized by long lead times. Smallholders need time to switch livelihood, prepare land, plant trees, and wait two and a half years for the first harvest. Production on industrial estates is characterized by similar lags. Trees take five to seven years to reach a productive state. Replanting occurs after 25 years, when yields

[4]Gaskell (2015) estimates that 92% of the increase in Indonesian palm oil production from 1985–2010 is due to land expansion and the remaining 8% due to yield improvements. Other crops expanded relatively little and several contracted, according to the 2003 and 2013 Agricultural Censuses.

[5]Large and small farms are usually geographically close. Smallholders need a mill close by to process and market their fruit. Virtually all palm oil processing plants depend on smallholder supply.

[6]In the Suharto era, industrial "nucleus" estates allocated a portion of new developments to company-supported smallholders, known as "plasma" or "scheme" smallholders (Pramudya, Hospes, and Termeer, 2016). Many plasma farmers were relocated from Java as part of the national transmigration program, examined in Bazzi, Gaduh, Rothenberg, and Wong (2016).

decline and fruit becomes difficult to reach.[7] Expansion is thus mostly determined by future demand and alternative rural livelihood opportunities, rather than short-term changes in socioeconomic conditions or commodity prices.[8]

Geographic differences in growing conditions led to large differences in cultivation intensity across regions, shown in Figure 2. Not all land is equally suitable for oil palm cultivation. Humid low-lying tropical areas with ample rainfall provide the ideal growing conditions, and navigable terrain allows for easier planting, harvesting, and transporting (Corley and Tinker, 2015). Districts with above-median suitability (described further below) increased the share of district under cultivation by 8.4 percentage points (92,000 hectares) more than those below the median and virtually all districts on main producing islands of Sumatra and Kalimantan cultivated some oil palm by 2015.[9] The median expansion from 2000–15 was around five percent of district area, or 42,000 hectares.[10]

3 Empirical approach

Using newly digitized data on local palm oil acreage, I compare development trajectories in districts with large increases in oil palm cultivation against those with smaller increases or none at all. I use two years of data, year 2000 because it predates the expansion and 2015 as the present, and estimate:

$$y_{d,t} = \beta P_{d,t} + \delta_d + \delta_t + \gamma X_{d,2000} * post + \varepsilon_d \tag{1}$$

[7] The price paid for a fresh fruit bunch increases with tree maturity. Prices are set weekly and published in local newspapers, reflecting limited pass-through of the world palm oil price to local markets (Boyabatli, Nguyen, and Wang, 2017). District fixed effects capture systematic differences across markets.

[8] Since Indonesia is the world's largest supplier, the palm oil price is unlikely to be an appropriate source of identifying variation (Dube and Vargas, 2013).

[9] Estimating marginal effects with a continuous treatment thus seems more appropriate than sacrificing treatment variation and geographic comparability to arbitrarily bin districts into a treated group.

[10] This figure is for expansion, i.e., districts that increased their area under cultivation. Including districts that did not increase their area under cultivation, the median change is 1 percent of district area (6,500 ha). 60/179 rural districts, 2000 district boundaries excluding Java, did not expand palm acreage.

$P_{d,t}$ is the share of district area used for oil palm farming in 2000 and 2015.[11] The temporal bandwidth of 15 years reflects the lags from planting, to harvesting, to exporting described in Section 2. $y_{d,t}$ is an outcome of interest in district d at the closest feasible periods to 2000 and the present. My primary outcomes are the district poverty rate and average monthly per capita household expenditures—two key policy targets capturing welfare and distribution well. Variable construction and data sources are detailed in the Data Appendix.

District fixed effects (FEs) δ_d absorb district-specific heterogeneities affecting the local extent of adoption (e.g., geography and climate; historical, cultural, and political institutions; and government policies). δ_t is a 2015 dummy capturing common changes. To account for potential convergence dynamics and adjust all estimates for initial observable differences in the most palm-suitable districts, $X_{d,2000}$ includes initial poverty rates, literacy rates, rural population shares, agricultural and manufacturing employment shares, and the share of villages in each district with paved roads, all interacted with a post period indicator. All results thus depend upon comparisons between districts with the same initial levels of development, industrialization, urbanization, transport, and employment.[12] Robust standard errors are clustered by district.

I modify the two-period panel in three important ways to make control districts more suitable and improve counterfactual comparisons. First, I apply 2000 district definitions to work with a balanced panel of constant-area geographic units, given my focus on land.[13] Second, I remove cities, where little palm is grown but palm oil companies are often

[11] Total district oil palm acreage is hand-digitized from the Tree Crop Statistics of Indonesia for Oil Palm yearbooks, produced annually by the Directorate General of Estate Crops at the Department of Agriculture. District palm acreage is divided by total district area to scale cultivation intensity by district size. I focus on palm acreage because (a) most of the increase in production was through land expansion, (b) land use is the central policy issue, and (c) this approach compares palm farming against all alternative rural land uses. Alternative parameterizations yield similar results, specifically palm acreage or production (in tons) per person, both either in level terms or taken as the inverse hyperbolic sine to transform while retaining zeros (Tables A2 and A3).

[12] Results are similar conditioning only on initial poverty, with initial conditions controls for remaining observable level differences, and with a polynomial in latitude and longitude (Tables A11 and A12).

[13] Decentralization saw the number of districts proliferate from 282 in 1998 to 514 in 2015. Fitriani, Hofman, and Kaiser (2005) and Bazzi and Gudgeon (2018) describe the balkanization.

headquartered. Including cities would violate the stable unit treatment value assumption, i.e., no interference. Third, I eliminate districts on the most populous island of Java, which grows little palm and is historically richer.[14] Finally, I stress that Equation 1 does not identify aggregate effects for Indonesia as a whole, but rather the general equilibrium effects at the district level, assuming no spillovers across districts.

3.1 Identification

To I address endogenous adoption by exploiting exogenous differences in cultivation suitability across districts and the external demand shock. Average district agro-climatically attainable palm oil yield is calculated from the Food and Agriculture Organization's (FAO) Global Agro-Ecological Zones (GAEZ) dataset. GAEZ uses agronomic models and high resolution geographic and climatic data to predict attainable yields for different crops on each piece of land regardless of whether the land is cultivated. It does not rely on actual cultivation in its estimates, nor does it involve estimating any sort of statistical relationship between observed inputs, outputs, and agro-climatic conditions.[15] I map gridded data on crop-specific potential yields to district boundaries, take district means, and interact with a post-period indicator to induce temporal variation reflecting the last two decades' rapid increase in palm oil demand.

The first stage intuition is that higher potential yields increase the likelihood of developing palm processing infrastructure and planting trees.[16] Panel A of Figure 3 shows this graphically, with a binned scatter plot of potential palm yields against the share of

[14]Results are nonetheless similar including Java, cities, or even island-by-year fixed effects to compare only across districts within island groups (Tables A5, A6, A11, and A12).

[15]Fischer, van Nelthuizen, Shah, and Nachtergaele (2002) detail GAEZ construction. Costinot, Donaldson, and Smith (2016) and Nunn and Qian (2011) discuss additional benefits of GAEZ for identification.

[16]Qualitative evidence gathered from interviews and focus group discussions with firms suggests that suitability is the first-order concern when making palm oil processing and plantation investments. Empirically, this pattern is also observed at a finer spatial level within districts with respect to the optimal palm processor placement (see Figure 9). Farmers are also highly attuned to the relative profitability of adopting, usually from observing neighbors.

district area under cultivation. The weak but positive relationship between potential yields and cultivation area in 2000 came to life by 2015, particularly in the most suitable districts.

The crucial identification assumption is that the interaction of potential palm yields and a post period indicator does not affect poverty trajectories through any channel other than palm adoption. Clearly, the primary channel for potential palm yields to affect economic outcomes must be through growing palms. However, one might still be concerned that suitable districts differ in other ways potentially correlated with development trajectories.

To clarify the plausibly exogenous nature of palm suitability, I estimate the following well-saturated panel specification:

$$poverty_{d,t} = \sum_{t=2002}^{2017} \sum_{c \in C} \alpha_{c,t} suit_{c,d,t} + \delta_d + \tau_{i,t} + \gamma_d * T + \epsilon_{d,t} \tag{2}$$

where poverty rates are observed each district-year from 2002–17, δ_d and $\tau_{i,t}$ are district and time fixed effects, and $\gamma_d * T$ are district-specific trends. α is the year-specific effect of suitability for palm oil and Indonesia's three other major cash crops: coffee, cocoa, and teas. Figure 4 plots the suitability*year coefficients. The absence of any statistically significant effects for other crops highlights the centrality of palm suitability. Statistically insignificant effects in the early years suggests similar pre-period trajectories. Standard pre-trends tests support this interpretation. Tables A7 and A8 ask whether pre-period trends in poverty and consumption are related to subsequent oil palm expansion and my instrument, finding no evidence of any statistically significant "placebo" effects. My main estimates are thus unlikely to be picking up preexisting trends.

I conduct four additional identification checks. First, I show that IV and reduced-form estimates are similar using alternative *relative* suitability instruments directly accounting for any underlying differences in suitability for other cash crops or agriculture overall. Specifically, I take the normalized difference between palm suitability and (a) Indonesia's other key cash crops and (b) all other crops. This exercise, like Figure 4, suggests that other agricultural potential, potentially picked up by shared inputs in the palm oil GAEZ productivity model, is not a major concern for my identification strategy (Tables A9 and A10). Second, I show that the main IV estimates are robust to the addition of a host of additional trends (Tables A11 and A12). Third, I first conduct a "zero-first-stage" falsification test with an auxiliary regression, examining the relationship between suitability and poverty in places that do not grow much palm oil—that is, the sub-sample where the first stage is effectively zero by construction (Table A13). The fact that the reduced form relationship between palm suitability and poverty reduction exists only in palm growing regions suggests that (a) my results are unlikely to be confounded by an unobserved correlate of palm suitability also correlated with poverty, and (b) impacts are coming through adoption in suitable regions (Bound and Jaeger, 2000; Altonji et al. 2005; Angrist et al. 2010; Nunn and Wantchekon, 2011). Finally, after the main results I show that the statistical relationships underpinning my identification strategy—between suitability, adoption, and poverty—are also observed across villages *within* districts, a much finer level of spatial aggregation (Figure 9).

3.2 First stage regression results

Table 1 presents first stage results. Column 1 includes the instrument, district and year FEs, and the baseline initial conditions trends. A potential yield of an additional metric ton per year corresponds to 2.1% more of the district being planted. Panel B of Figure 3 shows this graphically. To explicitly factor in pre-trends, Column 2 adds the change in the district poverty rate from 1993 (the first year SUSENAS became representative at the district level) to 2002 (after the Asian Financial Crisis). Column 3 includes trends related to the initial observable differences in the most palm-suitable districts not already included as baseline controls: ethnolinguistic fractionalization, the share of villages in each district with palm farmers, district production in tons, population density, and the percentage of households with access to electricity (Table A1). Column 4 includes a completed polynomial in latitude and longitude interacted with a post indicator. Across these demanding specifications, the point estimate is stable, standard errors small, and first stage robust.

Exploiting the variation in expansion arising from crop-specific agro-climatic suitability isolates the effects of cultivating oil palm on land where natural agro-climatic characteristics are best, not other sources of profitability like market access, trade costs, or input costs. This LATE may be different to those relating to these other sources of profitability, adopting in places less suitable, or the average treatment effect (ATE). A reduced form approach, by contrast, understates adoption effects by including suitable districts that do not cultivate much palm. With cultivation data I go a step further and estimate the impact of palm adoption induced by external demand and exogenous crop-specific geographic characteristic—an ideal policy parameter.

4 Poverty reduction and consumption growth

The main finding is that Indonesian districts converting more of their land for oil palm cultivation since 2000 achieved more rapid poverty reduction. Figure 5 shows a simplified version of the main result in the raw data over the 2000s, comparing the average poverty rate of rural districts with the most oil palm expansion against those without. Rural districts had similar poverty levels and trends in the early 2000s, but districts more intensively increasing palm oil production diverged as the decade progressed.

Table 2 presents the main regression estimates of the impacts on poverty (Columns 1–3) and average per capita household consumption (Columns 4–6). Each column reports a different version of Equation 1. Columns 1 and 4 give the ordinary least squares (OLS) relationships with cultivated area. Columns 2 and 5 report my preferred IV estimates. Columns 3 and 6 report the reduced form using average district agro-climatically attainable palm yield interacted with a post-period indicator. The unidentified OLS point estimate on oil palm land in Column 1 is -0.081, but increases in magnitude to -0.536 when instrumented with post*suitability. A ten percentage point increase in the share of district land under cultivation for oil palm, due to that district being more suitable, corresponds to an additional 5.36 percentage point reduction in district poverty. Although OLS and IV estimates are strictly incomparable, the increase from Column 1 to Column 2 is likely a combination of OLS bias (e.g., due to planting in areas with lower unobservable land costs, or weaker institutions) and the LATE exceeding the ATE (e.g., due to planting palms where most profitable). My IV estimates are thus best interpreted as upper bounds.[17] The reduced form relationship between suitability*post and poverty (Column 3) shows that districts with an average potential palm yield of an additional metric ton per hectare per year higher reduced poverty by 1.2 percentage points more. A simple policy simulation based on Column 2 suggests that up to 2.6 million of the 10 million Indonesians lifted from poverty this century were lifted exclusively due to growth in the palm oil sector.[18] Echoing Suryahadi, Suryadarma, and Sumarto (2009), my results highlight the continued importance of agricultural sector growth for rural poverty reduction in Indonesia.

[17]Results are also similar taking the natural log of the district poverty rate, which de-weights high poverty districts making higher level reductions in the estimation (Table A19), and using alternative parameterizations of district poverty (e.g., acreage or production per capita).

[18]This back-of-the-envelope calculation was done by multiplying the change in area under cultivation by the estimated coefficient on palm land in Column 2 of Table 3 to get the predicted percentage poverty reduction for each district. I then multiplied that by district population and summed over rural, non-Java districts to get the total number of poor lifted from poverty.

Columns 4–6 of Table 2 present estimates on average per capita household expenditure.[19] The OLS coefficient of 0.001 again illustrates the biases that OLS might introduce relative to the identified IV and reduced form specifications. The IV coefficient is 0.008, meaning the median area expansion of 5 percent of district area corresponds to a 4 percent faster increase in average per capita household expenditure. The reduced form estimate finds that a potential yield of an additional metric ton corresponds to 1.8% faster consumption growth.

4.1 Effect heterogeneity across households

I delve into the household surveys to understand which groups drive the poverty reduction and consumption growth. I classify SUSENAS households based on whether they derive most of their income from agriculture and whether they live in rural or urban areas.[20] Since cities are dropped, urban households refer to those those living in urban villages—that is, small towns in rural districts.

Figure 6 reports IV estimates for total, food, and non-food expenditure for all households and each of the four groups. The first point from the top reports the average effect on total per capital household consumption (i.e., from Column 5 of Table 3) for reference. Average effects are driven by rural households and more elastic non-food expenditures, which increase by over three percent for a single percentage point increase in palm area. Since most rural poor rely on agriculture for a livelihood, low-income households capturing rents from labor and land intensive growth is one explanation for the poverty reduction. Similar impacts on non-agricultural households suggest spillovers beyond agriculture, for example through demand for local goods and services (Foster and Rosenzweig, 2004; Emerick, 2018).

[19] Using households as the unit of analysis allows me to control for household size, and urban and sector FEs. The identifying variation is at the district level and all else is the same, with SUSENAS pooled over two waves.

[20] The share of households in agriculture and in urban and rural areas across districts over time is endogenous, and as such estimates represent effects for the average household in each group each year rather than comparisons of the same households, which SUSENAS does not allow over this time horizon.

Despite positive impacts for the average household, my main poverty findings could be explained by people near the poverty line being lifted just above, with little effect on the extreme poor who are apt to be marginalized in land and labor markets. Figure 7 presents the distribution of per capita household expenditures in 2015 for households in non-producing, mild producing, and major producing districts with over 20% of their area planted. The distribution shifts progressively to the right with cultivation intensity. The consumption "floor" is also higher in producing districts. To explore distributional impacts more formally, Figure 8 presents IV estimates of the effects on household expenditure for each decile. Households in each district-year are divided into deciles based on their total per capita expenditures. Each is used in the same manner as in Figure 6 to reveal the change in consumption for the average household in each decile.[21] Panel A of Figure 8 finds that the poorest 10% consume 2.5% more in the median expansion district relative to the poorest 10% in a counterfactual district with no expansion. This is not particularly surprising since the landless often work on large industrial estates and assist smallholders, whose largest production-related expenditure is hired labor (BPS, 2013). The bottom 20–60% experience the largest relative gains, with effects tapering off for the upper-middle class and ratcheting up again for top 10%. However, none of these decile impacts are statistically different from average effect in Column 5 of Table 3. In Panel B, I present the same estimates with expenditure in Indonesian rupiah (i.e., not logged) to highlight how the relative gains in Panel A translate into absolute dollar terms. The median household, experiencing the median expansion, has roughly an additional $3.5 USD per person per month—four days more consumption above the poverty line.

[21]This approach is analagous to extracting out percentiles for each district, which, although common in the literature (see, e.g., Topalova, 2010), assumes rank equivalence and a stable distribution over time, and should be interpreted as such.

4.2 Robustness—within-district estimates

The relationship between palm cultivation and local poverty is not unique to my cross-district trend comparisons. An annual many-way fixed effects model—exploiting only variation within districts over time, rather than changes across districts—yields similar results (Table A14). Within-panel estimates find effects slowly emerging over time, consistent with the crop life cycle and the lags motivating my longer-term specifications.

Indonesia's rich village data—compiled for a companion paper (Edwards, 2019)—allows me to also look for the empirical patterns underpinning my identification strategy at a much finer level of spatial variation: across villages in the same district. I compare villages near palm oil factories with those slightly farther away and unable to market palm fruits. Figure 9 presents "distance band" coefficients every 5 kilometers from a factory, adjusting estimates for locality fixed effects (district or nearest factory) and a rich vector of geographic characteristics capturing relative suitability within localities. Factories clearly place in the most suitable villages, and oil palm adoption is concentrated near factories. Consistent with the main district level comparisons, poverty decreases with proximity to processors. I focus on aggregate district effects and the long-difference approach for the rest of this article.

5 Potential explanations

This section attempts to explain the poverty reduction through (a) rising returns to land and labor, (b) indirect effects reinforcing the gains over time, and (c) migration. Since agriculture is relatively labor-intensive, any poverty benefits from expansion could be purely a direct labor income story for smallholders, workers on industrial estates, or people employed elsewhere in the supply chain—like a classic agricultural productivity shock (Evenson and Gollin, 2003; Emerick, 2018). However, in a setting of relatively abundant labor and reliance on land as a factor of production, increasing farmland alone (cf., raising productivity) could increase agricultural output and reduce poverty. I first confirm whether expanding the agricultural frontier explains most of the effect (cf., crop-switching and rising returns to land) and whether returns to labor are rising in expansion districts. I then explore three ways local agricultural surpluses might reinforce the direct income gains: (a) households could invest in productive assets and human capital (Foster and Rosenzweig, 1996); (b) revenue-flush local governments could do the same; and (c) export orientation and immediate processing mean that local infrastructure development may be a necessary condition to expand production (Donaldson, 2018; Dell and Olken, 2018), and that I may be capturing returns to that effect. I conclude this section by attempting to rule out migration as an alternative explanation.

5.1 Frontier expansion and factor intensities

The main results could be driven by expanding the agricultural frontier rather than a more efficient use of agricultural resources and rising returns to labor and land. To explore this possibility, I denominate palm acreage with the total area under cultivation for all types of agriculture (cf.., total district area) to adjust estimates for changes in the agricultural frontier and focus on the intensive margin. Total district farmland is calculated as the sum of village farmland reported in the 2003 and 2008 villages censuses (PODES).[22] Columns 1 and 2 of Table 3 report the main OLS and IV poverty results estimated over this shorter time window for comparison (i.e., 2003–08 instead of 2000–15). Marginal effects are larger than in Table 2, perhaps reflecting the higher palm oil price during the 2008 food price crisis. Columns 3 and 4 denominate palm oil acreage with total district farmland. The OLS estimate in Column 3 is not statistically different from zero. The identified IV estimate in Column 4 is indiscernible from that using total district area (Column 2), implying that the main results are not only due to new farmland but its particular use—that is, changes in

[22] Concordance between more recent data is poor, due to missing variables (PODES 2014) or different coverage and variable definitions (Agricultural Census 2013). I adjust the other variables to periods reflecting this shorter time horizon.

crop mix and overall agricultural productivity.[23] Columns 5–8 use naive OLS estimates to probe this conjecture from slightly different angles. Column 5 uses farmland as a share of total area as the explanatory variable to look at whether increasing farmland, regardless of its use, corresponds to faster poverty reduction. The point estimate is one third of the OLS estimate in Column 1. The final two columns run a "horse race" between an additional hectare of oil palm versus any farmland. Palm wins by a factor of eight.

Table 4 returns to the causally-identified IV estimates to examine labor productivity and wages. The goal here is confirm that labor is in fact capturing rents. Columns 1 and 2 use average district output per worker in agriculture and manufacturing as dependent variables. Columns 3–6 use average wages. A one percentage point increase in area under cultivation for palm oil corresponds to 160 million rupiah (12,000 USD) more output per worker per year in agriculture, 685 million (45,000 USD) more in manufacturing, and four percent faster wage growth across all sectors. Wage growth is almost entirely driven by agriculture, suggesting at least some pass-through of the productivity gains to the workers.[24] An alternative way to gauge the importance of labor intensity is to focus on small, family farms, which account for twice the jobs per hectare. Specifically, the median family farm is 2 hectares, while industrial farms employ around two farm laborers for every five hectares. However, most oil palm smallholders are also "part-time" farmers, since oil palm tends to be labor-saving relative to alternatives (Kubitza and Gerhke, 2018). Indeed, smallholder acreage yields considerably larger cultivation-poverty elasticities (Table A15).[25]

[23] Most palm-producing districts also expanded cropland, making extensive vs. intensive margin effects difficult to disentangle much further.

[24] Rural services, by comparison, are often unskilled, unproductive, and informal. Manufacturing labor is typically skilled, mobile, and a much smaller share, with wages more likely to equalize across districts.

[25] Similar to trying to separately identify intensive margin effects from land expansion, smallholder and industrial cultivation is closely intertwined in most all districts, making this an highly imperfect exercise.

5.2 Rural savings, investment, and infrastructure

This section explores indirect mechanisms potentially reinforcing the poverty reduction. I begin by asking what households are doing with their rising incomes. The first three panels of Figure 10 disaggregate impacts on non-food expenditure by expenditure and household type. All types of non-food expenditure increase, particularly health and education. The remaining panels examine whether higher durables expenditure corresponds to physical asset accumulation. I find that households in the median expansion district are twenty percent more likely to own a major asset and have on average three percent more floorspace.[26]

Local demand shocks can offer windfall revenues for local governments. In a system of highly decentralized government and fiscal affairs, such fiscal windfalls could amplify regional inequalities with more productive public investments and public services in growing regions. Indonesia's 2001 decentralization reforms devolved significant fiscal and policy autonomy to local governments, making this a real possibility. Districts are responsible for budgeting and service delivery and held accountable through local elections every five years. Panel A of Table 5 reports effects on district government revenue and expenditure. Columns 1 and 2 show that total district government revenue and spending are almost twenty percent higher in the median expansion district.[27] Columns 3 and 4 turn to villages, which are also able to raise revenue and provide basic services and infrastructure.[28] The median district agricultural expansion allowed the average village in that district to generate 35% more own source revenue and increase expenditure by around 25%.

[26]Home extensions—in addition to motorcycles, counted in assets—are often the first thing a rural household will buy following an income windfall and thus a good proxy for rural financial health. I cannot distinguish between productive and non-productive assets across SUSENAS 2002 and 2015, and more detailed analyses of human capital accumulation are slightly beyond the scope of this paper.

[27]Although Indonesia's system of intergovernmental transfers makes it difficult to rule out greater revenue and expenditure coming at the expense of other regions, statistically insignificant effects are returned using district transfer payments as "placebo tests". These estimates and separate impacts by revenue and expenditure type are in Tables A16 and A17.

[28]Own source revenue is the smallest revenue stream for villages. Most comes in a grant from the central government (i.e., Dana Desa). Districts provide additional transfers, often in-kind in the form of health clinics, schools, and other infrastructure. Village fiscal data are observed in the 2014 village census, so unaffected by the 2014 Village Law, which increased village funding and autonomy significantly.

Against a background of rising household incomes, increased fiscal capacity, and the need for new supply chain infrastructure, the second panel of Table 5 examines the main economic infrastructure variables reported in the village censuses: access to energy, road quality, and physical markets. Column 5 of Table 5 finds an economically large improvement in village access to clean cooking fuel—that is, using gas or kerosene provided through utilities and markets, instead of self-collected firewood or dung. Columns 6 and 7 consider village road quality: whether roads have been upgraded from dirt to hardened gravel or asphalt, and whether roads are fitted with street lights. Column 6 finds no evidence of improved *village* roads (i.e., not major roads), while Column 7 reports that the average village road in the median expansion district is 6.5 percent more likely to be fitted with a street light (consistent with the lower costs of fitting the light versus upgrading and maintaining the local road). Using household data in SUSENAS, I also find the average household in the median expansion district five percent more likely to be connected to the electricity grid and serviced by Perusahaan Listrik Negara (PLN), the main electricity company (Table A18). Column 8 uses an indicator for whether a village has built a permanent, physical market as a dependent variable. Markets are centers of commercial exchange, helpful for organizing agricultural activities and aggregating harvests. A ten percentage point increase in district palm cultivation leads to the average village in that district being four percent more likely to have built a market since 2000. With only sixteen percent of rural villages having markets in 2014, up from 12 percent in 2003, the effect size is economically significant.

The final panel of Table 5 turns to village social infrastructure: education, health, and religious facilities (e.g., churches and mosques). Unlike the economic infrastructure examined in Panel B, these public goods are less plausibly related to agricultural value chains. I find broad improvements in social infrastructure in expansion districts. A ten percentage point increase in palm acreage corresponds to an additional school, half a health clinic, and an extra mosque in the average village within expansion districts. PODES allows me to disaggregate education provision by public and private sectors, but not the other outcomes. The impact on education facilities is mostly explained by non-government schools. I cautiously conclude that at least some of the new public goods in palm producing regions may be privately provided, for example through in-kind transfers or new infrastructure to process, transport, and export palm oil.[29]

5.3 Migration

Three types of population changes could contaminate my results: (a) differential population growth altering compositions; (b) inward migration of non-poor people from non-producing districts; and (c) outward migration of poor people. Table 6 examines different population outcomes. The explanatory variable is the share of district area under cultivation for oil palm in 2000 and 2010, reflecting the years of the population censuses used to calculate the dependent variables. Column 1 cannot reject the null hypothesis of no effect on population. Columns 2 and 3 find slightly less recent inward migration in expansion districts. Although a local demand shock might be expected to increase in-migration (e.g., through a Harris and Todaro (1970) "labor pull" effect), the labor-saving nature of palm adoption (Kubitza and Gehrke, 2018) and relaxed liquidity constraints (Bryan et al. 2014; Bazzi, 2017) appear to dominate.

That migration to expansion districts is less common than elsewhere is reassuring, but does not tell us whether low-income people are leaving. Population censuses do not have data on income but I can examine migration status by education level. The probability of migrating increases with education, with poor households less likely to move (Figure A1). Moreover, districts are large geographic units and most migration is local. District-level analysis captures such sorting. A displaced individual is unlikely to move beyond the district capital (in no small part due to financial constraints), and cross-district migration is twice as

[29]Estimates in Panels B and C of Table 5 are similar when adjusted for log total district government revenue and expenditure and log village expenditure. These estimates available on author request but omitted here since fiscal variables would be "bad controls" and likely induce bias.

common as cross-province migration at all education levels. Migration patterns are similar across high and low suitability districts (Figure A2) and province-level estimates, which remove the influence of any cross-district migration within provinces, are also qualitatively similar to the main district-level results (Table A19). Although I cannot rule out poor people systematically leaving palm-growing districts and being replaced by non-poor inward migrants, it seems unlikely to fully explain my findings.

6 Impacts on tree cover loss and fire

I conclude my analysis by circling back to the public debate on palm oil and estimating the local environmental trade-off arising from an oil palm-driven change in poverty. I relate state-of-the-art satellite-based measures of district tree cover loss and fire to changes in cultivation area with the following long difference specification:

$$y_d = \beta(P_{d,2015} - P_{d,2000}) + \gamma X_{d,2000} + \varepsilon_d \tag{3}$$

where y_d is either gross tree cover loss (excluding regeneration) as a share of district area or thermal hotspot detections since 2000. $P_{d,2015} - P_{d,2000}$ is the change in the share of district area under cultivation for oil palm (instrumented with palm suitability) and $X_{d,2000}$ includes the same initial conditions.[30] I stress two points regarding the interpretation of these estimates. First, tree cover loss is an imperfect measure of deforestation, including changes in forestry, tree crops, and wildfires in addition to any primary forest loss. Second, since many oil palms were planted after clearing forest and fire is used to clear or prepare land for agriculture, the following results are somewhat mechanical.

[30] I opt for the cross-sectional long-difference analogue of the main panel specification for two reasons. First, tree cover loss data represents the change in pixels since 2000. Second, since fire hotspots are highly seasonal (mostly due to El Niño), using the total detections over the expansion is closer to the spirit of my main approach and avoids rely on one end line year.

Environmental impacts are presented in Table 7. Columns 1–3 present OLS, IV, and reduced form estimates for district forest loss from palm expansion since 2000. A one percentage point increase in district area under cultivation on average corresponds to between an 0.8–1.7 percentage point loss in forest cover. Columns 4–6 use district hotspot detections since 2000 as the dependent variable and Poisson estimation since data are counts. Hotspot detections increased by roughly eight percent for each percentage point increase in the share of a district planted with palm since 2000, with major health impacts (Jayachandran, 2009; Rosales-Rueda and Triyana, 2018). Together with my main results, these estimates suggest that each percentage point of poverty reduction that has been achieved through extensive palm oil expansion since 2000 has come at the cost of between 1.5 and 3 percent of district area lost in tree cover and around ten percent more fire. These large and precisely estimated effects suggest that agricultural growth, forest loss, and fire have—at least over the last fifteen years—gone hand-in-hand in the Indonesian countryside.

7 Conclusion

This paper measured the impacts of Indonesia's rapid increase in palm oil exports from 2000 to 2015 on welfare in producing communities. Although national poverty continued to decline since the fall of Suharto in 1998, rural areas more intensively increasing palm oil production experienced faster poverty reduction. The magnitude of the effect is not trivial. National poverty declined from 18.2% to 11.2% from 2002–2015, but the median expansion district reduced poverty around five percentage points faster than an otherwise similar rural district. Consumption impacts are also significant, with four percent faster consumption growth in the median expansion district.

My findings line up behind a large body of work emphasizing the benefits of trade, export market access, and agricultural growth for managing and alleviating poverty in developing countries. I find little empirical support for widely-held views that export-oriented agriculture functions as an economic enclave and brings little benefit to local communities, at least in the context of Indonesian palm oil. Evidence on the channels at work clarify why. Direct impacts are coming through a broad rise in farm gate incomes through rising returns to land and labor, offering a contrast to demand shocks where production is less labor intensive and resource extraction more concentrated. An increasingly outward, market-oriented agricultural sector appear to be reinforcing these gains through more local government revenue and complementary economic and social infrastructure.

Globalization is in retreat. Several major economies are turning inwards and invoking trade discriminatory trade policies, particularly against products from developing countries. This study highlights the potential benefits of export growth and the importance of continued integration into global value chains for producing regions in developing countries. However well intentioned, policy actions that shift demand away from palm oil and other commodities produced predominantly in low and middle income countries are likely to be detrimental for poverty and economic development in producing regions, at least in the short to medium run. The significant yield differences across crops also highlight the potential for unintended adverse environmental consequences. For example, meeting estimated future vegetable oil demand in 2050 (around 310 Mt) without palm oil—with most oil instead produced by soy, rapeseed, and sunflower—is estimated to require around 300 million additional hectares of farmland, placing immense additional pressure on forests and limiting our ability to curb climate change (Meijaard et al., 2018).

References

Allcott, H., and Keniston, D. (2017) “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America”, Review of Economic Studies, 85: 2, pp. 596–731.

Altonji, J.G., Elder, T.E., and Taber, C.R. (2005) “ An evaluation of instrumental variable strategies for estimating the effects of Catholic schooling”, Journal of Human Resources, 40, pp. 791–821.

Angrist, J.D., Lavy, V., and Schlosser, A. (2010) “Multiple experiments for the causal like between the quantity and quality of children”, Journal of Labor eEconomics, 28, pp. 773–824.

Aragon, F., and Rud, J.P. (2013) “Natural resource and local communities: evidence from a Peruvian gold mine”, American Economic Journal: Economic Policy, 5:2, pp. 1–25.

Austin, K. G., Schwantes, A., Gu, Y., and Kasibhatla, P. S. (2019) “What causes deforestation in Indonesia”, Environmental Research Letters, 12:19, 024007.

Autor, D., Dorn, D., and Hanson, G. (2013), “The China Syndrome: Local Labor Market Effects of Import Competition in the United States”, American Economic Review, 103:6, pp. 2121–68.

Badan Pusat Statistik (2013), “Sensus Pertanian 2013”, available online: st2013.bps.go.id/

Balat, J., Brambilla, I, and Porto, G. (2009) “Realizing the gains from trade: export crops, marketing costs, and poverty”, Journal of International Economics, 78:1, pp. 21–31.

Barham, B., M. Clark, E. Katz, and R. Schurman (1992), “Nontraditional Agricultural Exports in Latin America,” Latin American Research Review 27:43–82.

Bazzi, S. (2017) “Wealth heterogeneity and the income elasticity of migration”, American Economic Journal: Applied Economics 9:2, pp.224–228.

Bazzi, S., Gaduh, A., Rothenberg, A., and Wong, M. (2016), “Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia,” American Economic Review, September 2016.

Bazzi, S., and Gudgeon, M. (2018), “The political boundaries of ethnic divisions,” working paper.

Bellemare, M. and Bloem, J. (2018), Does contract farming improve welfare? A Review,” World Development, forthcoming.

Bhagwati, J. (1958), “Immiserizing Growth: A Geometrical Note,” Review of Economic Studies 25, (June), pp. 201–205.

Bound, J., and Holzer, H.J. (2000), “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s”, Journal of Labor Economics, 18:1, pp. 20–54.

Bound, J., and Jaeger, D.A. (2000) "Do compulsory school attendance laws alone explain the association between quarter of birth and earnings?" Research in Labor Economics, 93, pp. 83–108.

Boyabatli, O., Nguyen, J., and Wang, T. (2017), "Capacity Management in Agricultural Commodity Processing and Application in the Palm Industry," Management and Service Operations Management, pp. 1–39.

Bryan, G., Chordhury, S., and Mobarak, A.M. (2014), "Underinvestment in a profitable technology: the case of seasonal migration in Bangladesh", Econometrica, 82:5, pp. 1671–1748.

Burgess, R., Hansen, M., Olken, B. A., Potapov, P., and Sieber, S. (2012), "The Political Economy of Deforestation in the Tropics," Quarterly Journal of Economics, 122:1, pp. 73–117.

Bustos, P., Caprettini, B., and Ponticelli, J. (2016), "Agricultural Productivity and Structural Transformation: Evidence from Brazil," American Economic Review, 106:1, pp. 130–1365

Bustos, P., Caprettini, B., and Ponticelli, J. (2018), "Capital accumulation and structural transformation," working paper.

Byerlee, D., de Janvry, A., Sadoulet, E. (2009), "Agriculture for Development: Toward a New Paradigm," Annual Review of Resource Economics," 1: 15–31.

Byerlee, D., Falcon, W., and Naylor, R. (2016), "The Tropical Oil Crops' Revolution: Food, Feed, Fuel, and Forests," Oxford University Press, Oxford.

Carter, M.R., B.L. Barham, and D. Mesbah. (1996) "Agricultural Export Booms and the Rural Poor in Chile, Guatemala, and Paraguay." Latin American Research Review 31:33–65.

Caselli, F., and Michaels, G. (2013), "Do oil windfalls improve living standards? Evidence from Brazil," American Economic Journal: Applied Economics, 5:1, pp. 208–38.

Castilho, M., Menendez, M., and Sztulman, A. (2012), "Trade Liberalization, Inequality, and Poverty in Brazilian States", World Development 40:4, pp. 821–835.

Corley, R., and Tinker, R. (2015), "The Oil Palm," 5th Edition, Blackwell, Oxford.

Costa, F., Garret, J., and Pessoa, J. (2017), "Winners and losers from a commodities for manufactures trade boom," Journal of International Economics, 102, pp. 50–69.

Costinot, A., Donaldson, D., Smith, C. (2016), "Evolving Comparative Advantage and the Impact of Climate Change on Agricultural Markets: Evidence form 1.7 Million Fields around the World," Journal of Political Economy, 24:1, pp. 205–248.

Cust, J. and Poelhekke, S. (2015) "The Local Economic Impacts of Natural Resource Extraction", Annual Review of Resource Economics, Vol. 7 pp. 251–268.

Dal Bo, E., and Dal Bo, P. (2011) "Workers, Warriors, and Criminals: Social Conflict in General Equilibrium", Journal of the European Economic Association, 9:4.

Dasgupta, S., Laplante, B., Wang, H., and Wheeler, D. (1992), "Confronting the Environmental Kuznets Curve,," Journal of Economic Perspectives, 16:1, pp. 147–168.

de Janvry, A., Fafchamps, M., and Sadoulet, E. (1991), "Peasant Household Behavior with Missing Markets: Some Paradoxes Explained," Economic Journal, 101, 1400–17.

de Zegher, J.F., Iancu, D.A., Plambeck, E.L. (2018) "Sustaining Rainforests and Smallholders by Eliminating Payment Delay in a Commodity Supply Chain: It Takes a Village", unpublished working paper.

Dell, M., and Olken, B. (2018), "The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java," working paper.

Dercon, S. and Gollin, D. (2014), "Agriculture in African Development: Theories and Strategies," Annual Review of Resource Economics, 6, pp. 471–492

Directorate General of Estate Crops (2017), "Tree Crop Estate Statistics of Indonesia, Palm Oil, 2015–2017," Department of Agriculture, December 2017, Jakarta.

Dix-Carneiro, R., and Kovak, B. (2017) "Trade Liberalization and Regional Dynamics," American Economic Review, 107:109, pp. 2908–2946.

Donaldson, D., (2018), "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure, American Economic Review, 108:4-5, pp. 899–934.

Dube, A., and Vargas, J.F. (2013) "Commodity Price Shocks and Civil Conflict: Evidence from Columbia", Review of Economic Studies, 80, pp. 1384–1421.

Easterly, W. (2007), "Inequality does cause underdevelopment: Insights from a new instrument," Journal of Development Economics, 85, pp. 755–776.

Edmonds, E.V., and Pavcnik, N. (2006) "Trade liberalization and the allocation of labor between households and markets in a poor country", Journal of International Economics, 69:2, pp. 272–295.

Edmonds, E. V., Pavcnik, N. and Topalova, P. (2010), "Trade Adjustment and Human Capital Investments: Evidence from Indian Tariff Reform", American Economic Journal: Applied Economics, 2, pp. 42-75.

Edwards, R. B. (2016), "Mining away the Preston curve," World Development, 78, February, pp. 2236.

Edwards, R. B. (2019), "Spillovers from agricultural processing," working paper.

Edwards, R. B., Falcon, W.P., Higgins, M.M., and Naylor, R.L. (2019), "Causes of Indonesia's forest fires", working paper.

Emerick, K. (2018) "Agricultural productivity and the sectoral reallocation of labor in rural India", Journal of Development Economics, 135, pp. 488–503.

Engerman, S., and Sokoloff, K. (2002) "Factor endowments, inequality, and paths of development among new world economies," NBER Working Paper Series 9259.

Evenson, R.E., and Gollin, D. (2003), “ Assessing the impact of the Green Revolution, 1960—2000”, Science 300: 5620, pp. 758–762.

Fafchamps, M., Hill, R.V., Kauda, A., and Nsibirwa, R.W., (2001), “The transmission of international commodity prices to domestic producers,” Mimeograph 2003, Oxford University, CSAE Working Paper Series, 2003–14.

Feler, L., and Senses, M., (2017), “Trade Shocks and the Provision of Local Public Goods,” American Economic Journal: Economic Policy, 9:4, pp. 101-54.

Feyrer, J., Manson, E.T., and Sacerdote, B. (2017) “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution”, American Economic Review, 107: 4, pp. 1313–34.

Fischer, G., van Nelthuizen, H., Shah, M., and Nachtergaele, F. (2002), “Global Agro-ecological Assessment for Agriculture in the 21st Century: Methodology and Results,” Food and Agriculture Organisation, Rome.

Fitriani, F., Hofman, B., and Kaiser, K. (2005), “Unity in diversity? The creation of new local governments in decentralising Indonesia,” Bulletin of Indonesian Economic Studies, 74: 1, pp. 70–93.

Foster, A., and Rosenzweig, M.R. (1996) “Technical change and human-capital returns and investments: evidence from the green revolution”, American Economic Review, pp 931–953.

Foster, A., and Rosenzweig, M.R. (2004) “Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970–2000”, Economic Development and Cultural Change, 52:3, pp. 509–542.

Gaddis, I. and Pieters, J. (2016), “The gendered impacts of trade liberalization: evidence from Brazil”, Journal of Human Resources, forthcoming

Gaskell, J. C. (2015), “The Role of Markets, Technology, and Policy in Generating Palm-Oil Demand in Indonesia,” Bulletin of Indonesian Economic Studies, 51:1, pp. 29–45.

Gaveau, D., Sheil, D., Husnayaen, Salim., M.A., Arjasakusuma, S., Ancrenaz, M., Pacheco, P., and Meijaard, E. (2016) “Rapid conversions and avoided deforestation: examining four decades of industrial plantation expansion in Borneo”, Scientific Reports, 6, Article 32017.

Goldberg, P. K., and Pavcnik, N. (2007), “Distributional effects of globalization in developing countries,” Journal of Economic Literature, 45: 1, pp. 39–82.

Gollin, D. (2010), “Agricultural Productivity and Economic Growth,” in Handbook of Agricultural Economics, pp 3825–3866.

Greenstone, M., and Jack, K. (2015)“Envirodevonomics: A Research Agenda for an Emerging Field,” Journal of Economic Literature, 53:1, pp. 5–42.

Hasan, R., Mitra, D., Ranjan, P. and Ahsan, R. (2012) “Trade liberalization and unemployment: Theory and evidence from India”, Journal of Development Economics, 97: 2, pp. 269-280.

Hayami, T. (2010), “Plantations Agriculture,” in Handbook of Agricultural Economics, pp. 3305–3322.

Hornbeck, R. and Keskin, P. (2015) “Does agricultura generate local economic spillovers? Short-run and long-run evidence from the Ogallala aquifer”, American Economic Journal: Economic Policy, 7:2, pp. 192–213.

Jayachandran, S. (2009) “Air quality and early-life mortality: evidence from Indonesia’s wildfires”, Journal of Human Resources, 44, pp. 916–954

Johnston, B.F., and Mellor, J.W. (1961), “The Role of Agriculture in Economic Development,” American Economic Review, 51:4, pp. 566–593.

Kis-Katos, K., and Sparrow, R. (2011), “Child labor and trade liberalization in Indonesia”, Journal of Human Resources, 46, pp. 722–49.

Kis-Katos, K., and Sparrow, R. (2015), “Poverty, Labour Markets and Trade Liberalization in Indonesia,” Journal of Development Economics,117, pp. 94–106.

Krugman, P. (1987) “The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher,” Journal of Development Economics, 27:1–2, pp. 41–55.

Kubitza, C., and Gehrke, E. (2018) Why does a labor-saving technology decrease fertility rates? Evidence from the oil palm boom in Indonesia”, unpublished working paper

Marden, S. (2018) “The Agricultural Roots of Industrial Development: Rural Savings and Industrialisation in Reform Era China,” working paper.

Matsuyama, K. (1992), “Agricultural Productivity, Comparative Advantage, and Economic Growth,” Journal of Economic Theory, 58: 2, pp. 317–34.

McCaig, B. (2011), “Exporting out of Poverty: Provincial poverty in Vietnam and U.S. market access”, Journal of International Economics, 85:1, pp. 102113.

Meijaard, E. et al. (eds.) (2018) “Oil palm and biodiversity: a situation analysis by the IUCN Oil Palm Task Force”,

Mokyr, J. (1976) “Industrialization in in the low countries, 1795–1850” New Haven CT, Yale University Press.

Naylor, R. L., Higgins, M. M., Edwards, R. B., and Falcon, W. P. (2018), “Decentralization and the Environment: Assessing Smallholder Oil Palm Development in Indonesia,” working paper.

Nunn, N., and Qian, N. (2011) “The Potato’s Contribution to Population and Urbanization: Evidence from a Historical Experiment”, Quarterly Journal of Economics, 126, pp. 593–650

Nunn, N., and Wantchekon, L. (2011) “The Slave Trade and the Origins of Mistrust in Africa,” American Economic Review, 101:7, pp. 3321–52.

Pramudya, E. P., Hospes, O., and Termeer, C. (2016), “Governing the palm oil sector through finance: the changing roles of the Indonesian state,” Bulletin of Indonesian Economic Studies, forthcoming.

Rada, N., Buccola, S.T., and Fuglie, K.O. (2011), "Government policy and agricultural productivity in Indonesia," American Journal of Agricultural Economics, 93:3, pp. 867–884.

Rosalies-Rueda, M., and Triyana, M. (2018), "The persistent effects of early-life exposure to air pollution: evidence from the Indonesian forest fires", Journal of Human Resources, forthcoming.

Suryahadi, A., Suryadarma, D., and Sumarto, S. (2009), "The effects of location and sectoral components of economic growth on poverty: Evidence from Indonesia," Journal of Development Economics, 89:1, pp. 109–117.

Topalova, P. (2007), "Trade Liberalization, Poverty and Inequality: Evidence from Indian Districts", in: Ann Harrison (ed.), Globalization and Poverty, National Bureau of Economic Research and University of Chicago Press: Chicago and London.

Topalova, P. (2010), "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India," American Economic Journal: Applied Economics, 2, pp. 1–41.

United States Department of Agriculture (2016), "Indonesia: Rising Global Demand Fuels Palm Oil Expansion," USDA Commodity Intelligence Report October 8, available online at: http://wwwpecad.fas.usda.gov/

World Bank (2008), "World Development Report: Agriculture for Development," World Bank Washington, DC.

World Bank (2015) "Indonesia Database for Economic and Policy Research," available online: http://data.worldbank.org/

Figures and Tables

Figure 1: Indonesia's Palm Oil Expansion

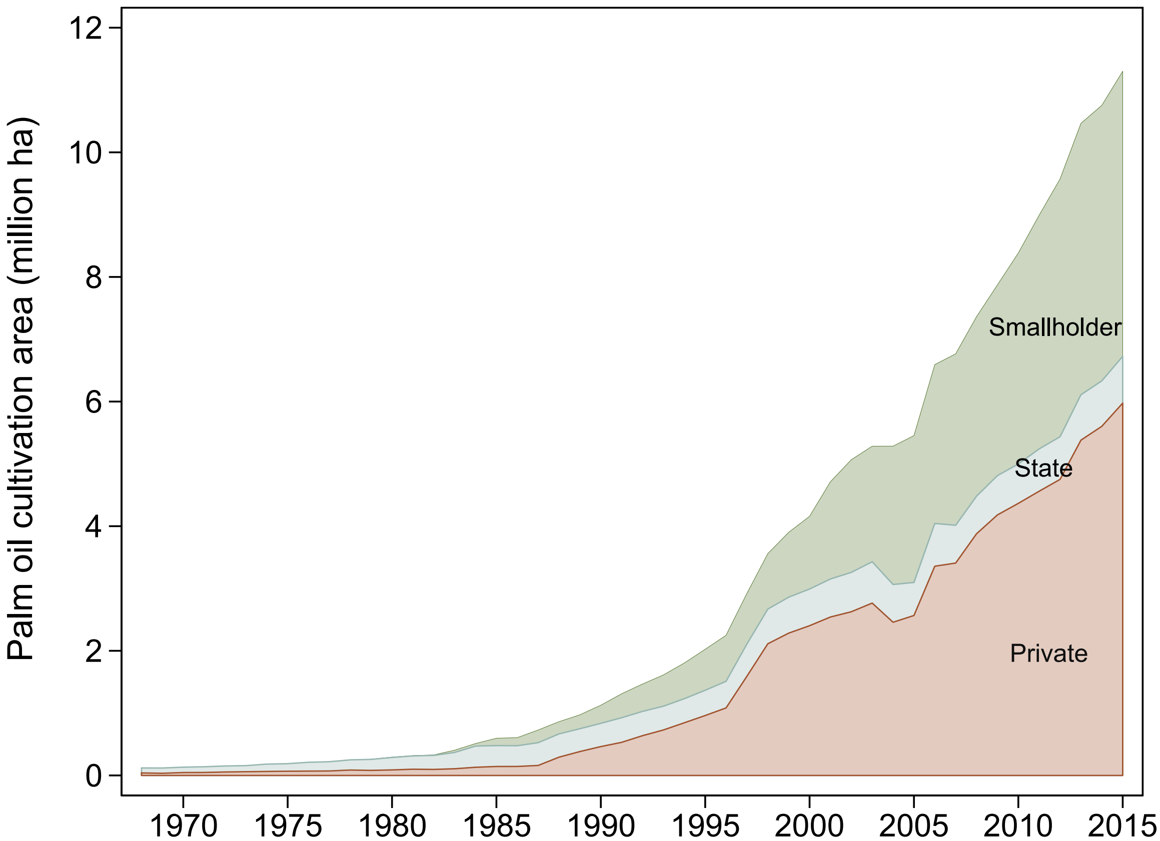

Notes: Data are taken from the Tree Crop Statistics of Indonesia for Oil Palm yearbooks, produced annually by Badan Pusat Statistik (BPS) and the Department of Agriculture of the Government of Indonesia and digitized by the author.

Figure 2: District Palm Oil Acreage, 2015

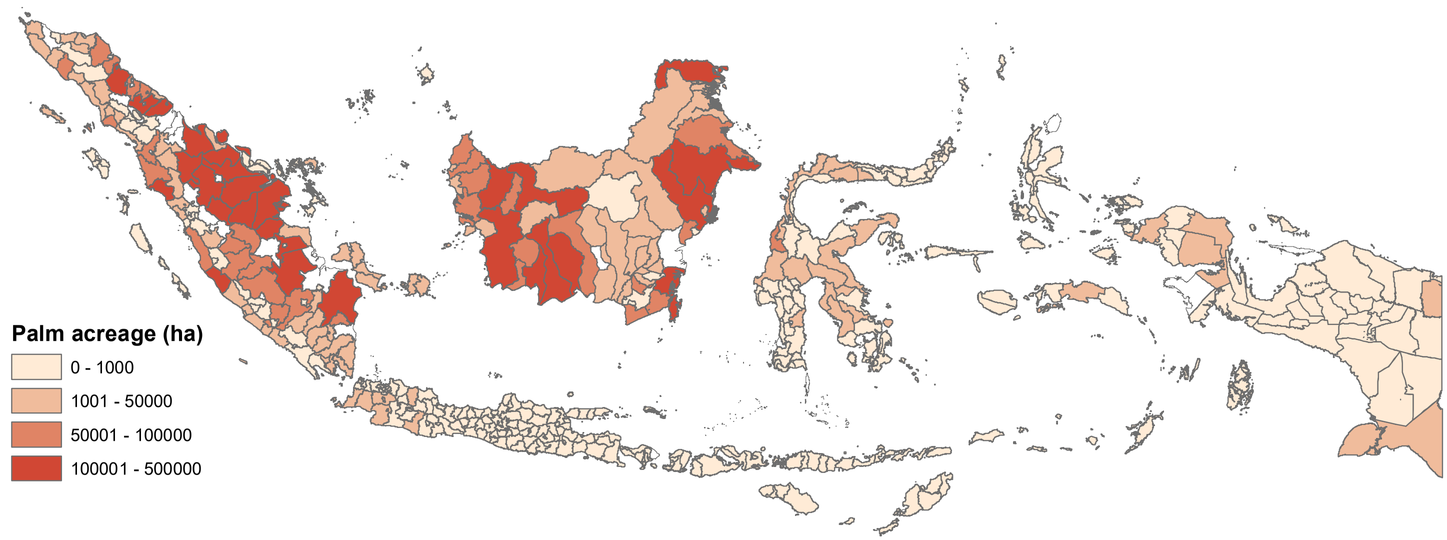

Notes: District palm oil acreage in 2015 is taken from the Tree Crop Estate Statistics of Indonesia, 2014–2016 produced by the Directorate General of State Crops (Statistik Perkebunan Indonesia, 2014–2016 Kelapa Sawit) and digitized by the author. 2010 district boundaries are used.

Figure 3: First Stage

(a) District cultivated area in 2000 and 2015

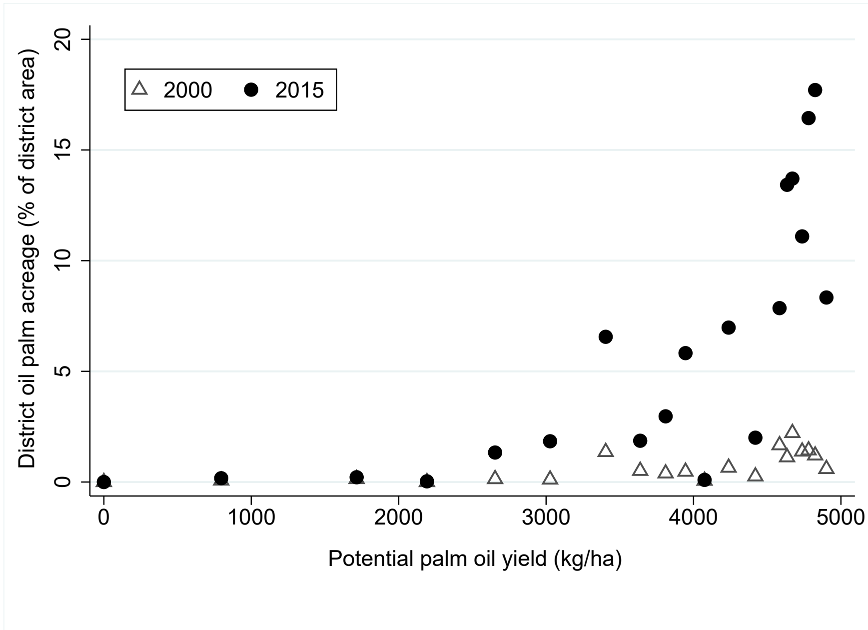

(b) Potential yields and area expansion

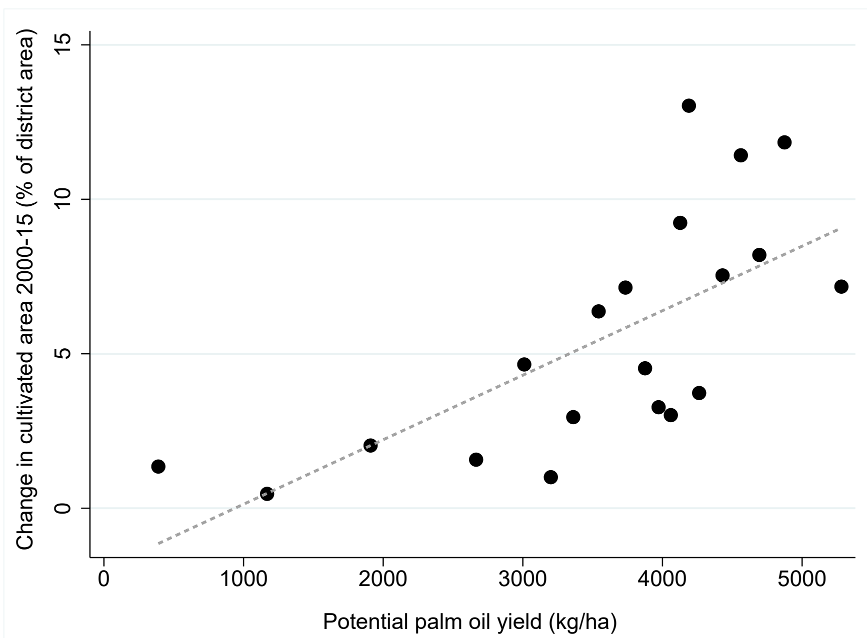

Notes: Panel A presents a binned scatter plot of district potential palm oil yield against the share of each district under cultivation for oil palm, split by year, to illustrate the increasing salience of the instrument after the demand shock. Panel B uses the change from 2000 to 2015 on the Y axis and includes the baseline initial conditions controls, showing the main first stage regression visually. Data are taken from the Tree Crop Estate Statistics of Indonesia and FAO-GAEZ.

Figure 4: Crop-By-Year Effects of Suitability

(a) Palm Oil

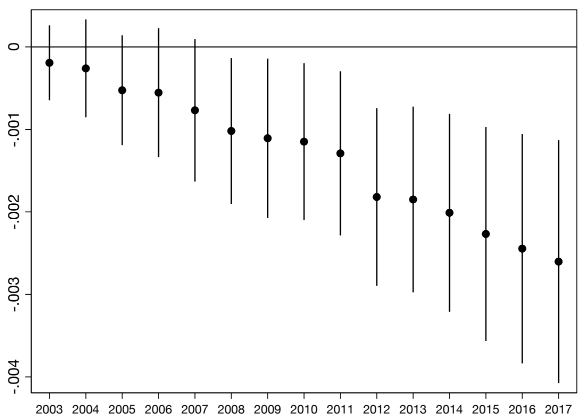

(b) Coffee

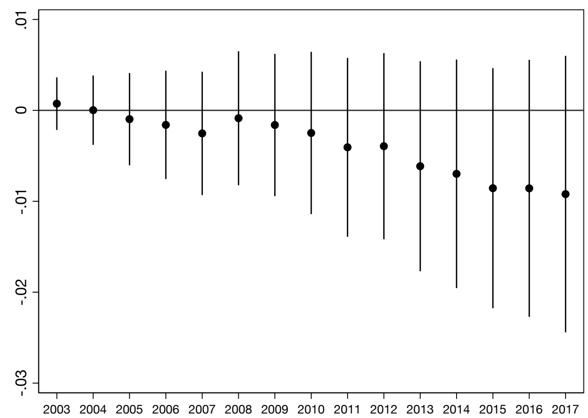

(c) Cocoa

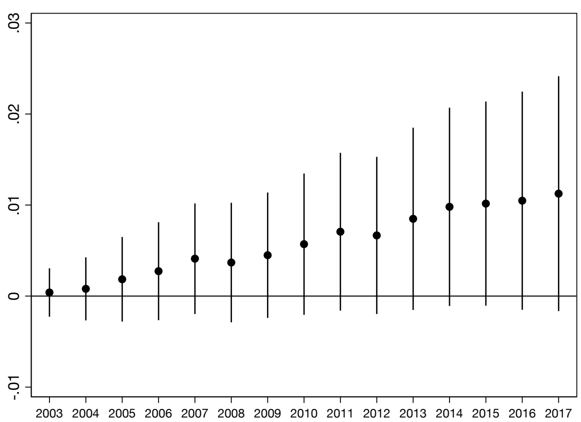

(d) Teas

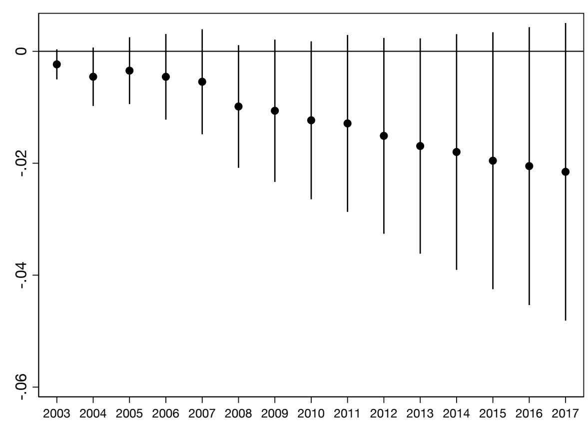

Notes: These figures plot crop-by-year specific effects on poverty from a saturated linear panel model from 2002–2017. BPS district poverty rates are the dependent variable, with the suitability for palm and other major cash crops all interacted with year dummies to trace out the reduced form impacts of suitability over time. The model also includes district fixed effects, district-specific linear trends, and time dummies. Robust standard errors are clustered at the district level, and the vertical lines indicate 95% confidence intervals. Although agricultural suitability for other cash crops appears to grow in importance in explaining variation in poverty over time, these effects are most pronounced for palm oil, where impacts from the mid-2000s are all statistically significant.

Figure 5: Expansion Districts Reduced Poverty Faster

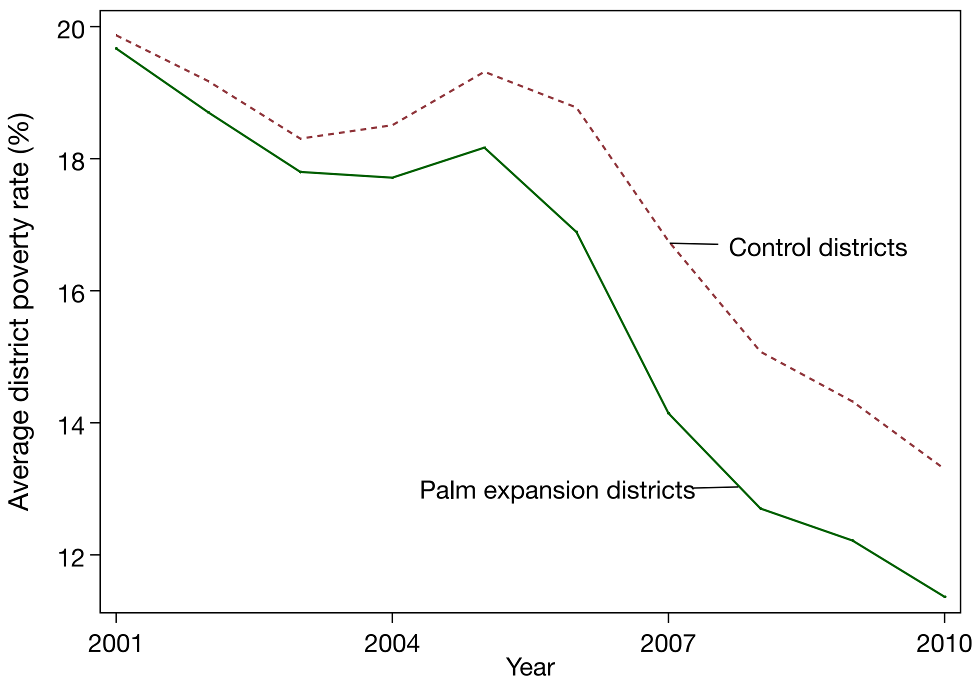

Notes: This figure is constructed using the World Bank's DAPOER database, available through its databank. All cities (*kotas*) and rural districts outside major palm oil cultivating regions are excluded. Expansions are those with the largest expansion— specifically, the top quarter of "expanders" increasing the share of district under cultivation by more than 17.5% from 2000–15.

Figure 6: Consumption Impacts, By Type and Sector

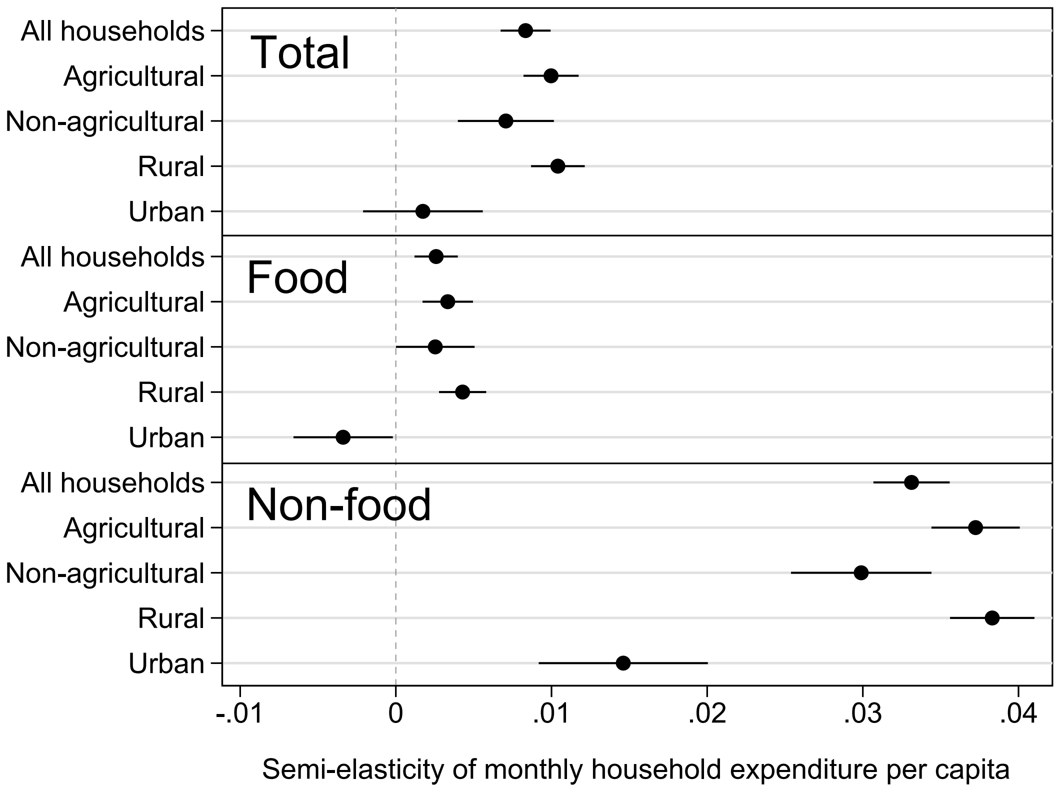

Notes: This graph plots the estimated coefficients on oil palm land from my primary IV estimator using log per capita monthly household expenditure as a dependent variable for the full sample of SUSENAS households ("All households") and for sub-groups listed on the Y axis. Black lines indicate 95% confidence intervals. The full sample is repeat cross-section of all households in SUSENAS 2002 and 2015 linked to two-period balanced panel of all rural districts at 2000 boundaries excluding Java. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator. District and year fixed effects, initial district conditions trends separately interacting 2000 log poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads with a post period dummy, and additional controls for household size, an urban/rural dummy, and sector fixed effects related to where households' primary income source are included throughout. Urban/rural (sector) fixed effects are dropped when I examine effects by urban-rural households (across sectors).

Figure 7: Consumption Distribution, 2015

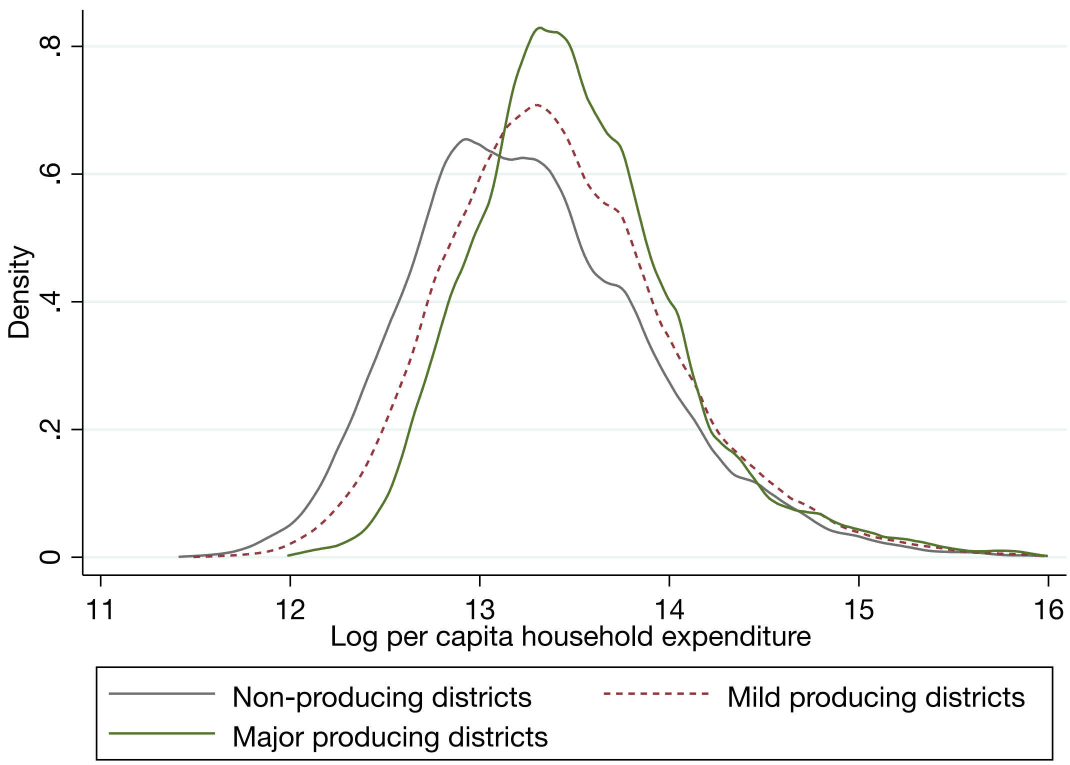

Notes: This graph plots kernel density estimates of log per capita household consumption in 2015 for households in rural districts not on Java that do not produce palm oil (gray solid), those that product only a little (red dash), and those that are major producers (green solid), defined as over 20% of the area under cultivation for oil palm. Data are taken from SUSENAS 2015.

Figure 8: Consumption Impacts By Decile

(a) Relative Gains

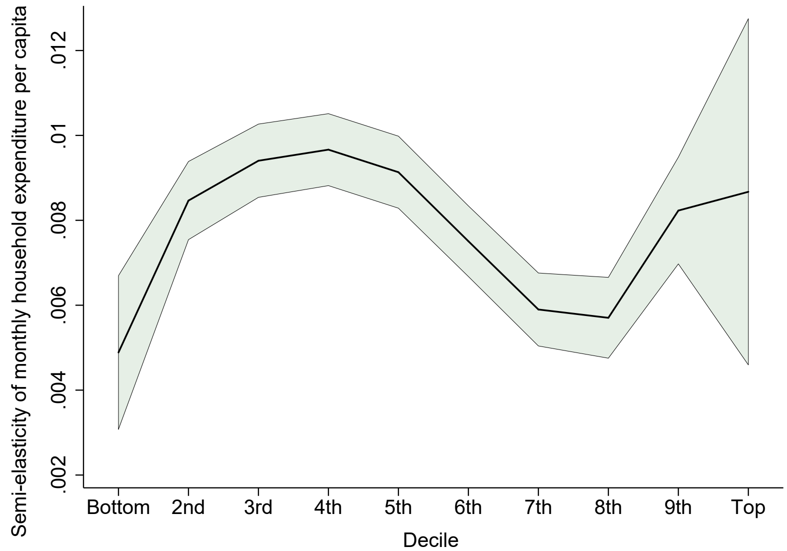

(b) Absolute Gains

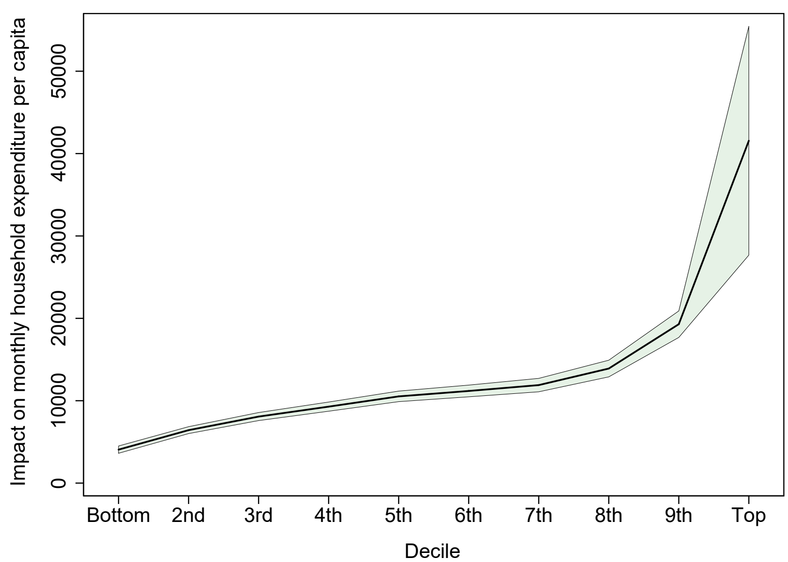

Notes: These graphs plot the estimated coefficients on oil palm land from my primary IV estimator using [log] per capita monthly household expenditure as a dependent variable after dividing each district-year group of households up by decile of the consumption distribution. The green bands indicate 95% confidence intervals. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator. District and year fixed effects, initial district conditions trends separately interacting 2000 log poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads with a post period dummy, and additional controls for household size, an urban/rural dummy, and sector fixed effects are included throughout. The full sample is repeat cross-section of all households in SUSENAS 2002 and 2015 linked to two-period balanced panel of all rural districts at 2000 boundaries, excluding Java.

Figure 9: Within-District Village Distance Band Estimates

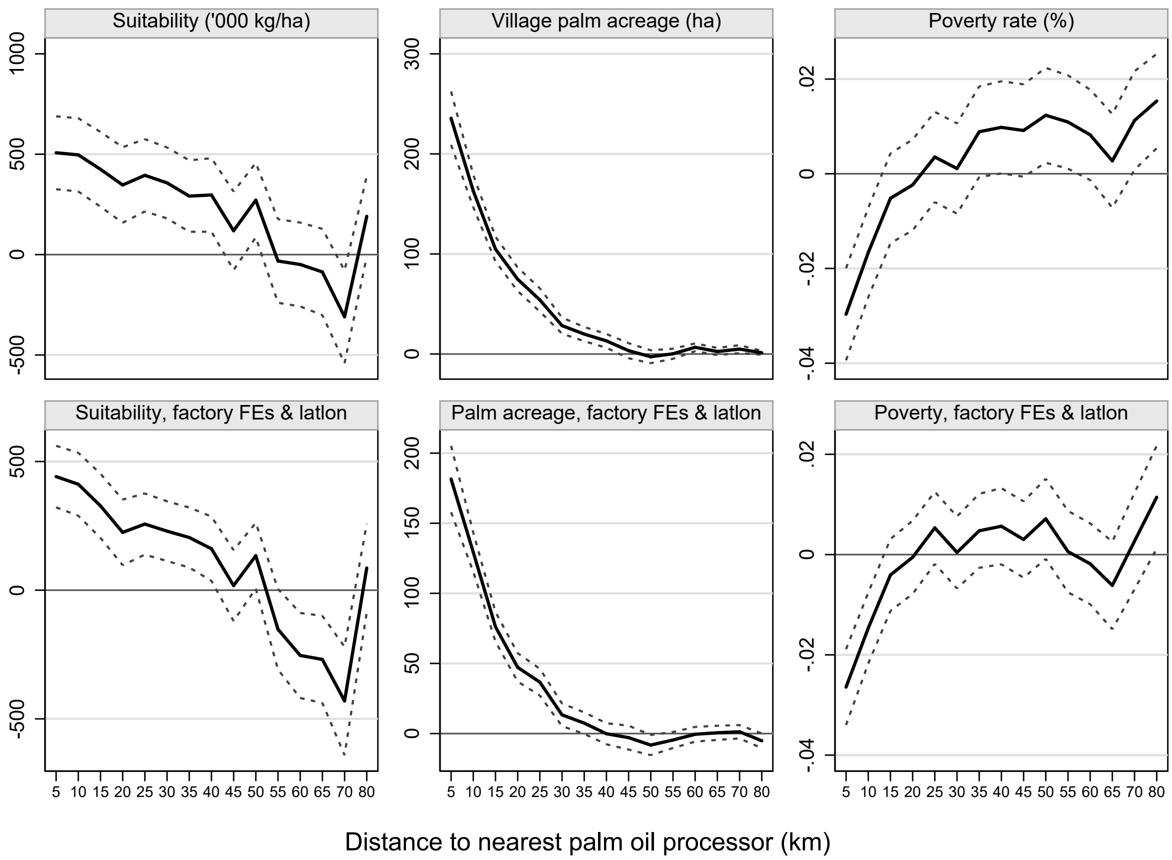

Notes: These figures plot the coefficients from distance band estimates at the village level, relating a village outcome to its proximity to the nearest palm oil processing factory. The top row is a baseline model with a separate dummy indicator every 5 kilometers from the factory, a host of exogenous geographic controls capturing relative suitability, and district fixed effects. Beyond 80 km is the excluded bin and villages more than 100 km away, in cities, and on Java are discarded from the estimation sample. The bottom row adds nearest factory fixed effects and a complete polynomial in latitude and longitude. Palm oil factories are identified in the 2016 Economic Census. The geocoding procedure is described in Edwards (2019). Palm oil suitability is the main GAEZ agro-climatically attainable yield data used to construct my instrument except calculated for every village. Village palm oil acreage is observed in the 2013 Agricultural Census of agricultural households, thus excluding industrial estates. Village poverty is estimated using standard poverty mapping techniques based on the 2010 Population Census and 2015 SUSENAS by the SMERU Research Institute, who kindly shared this data.

Figure 10: Impacts on Non-Food Expenditures and Assets

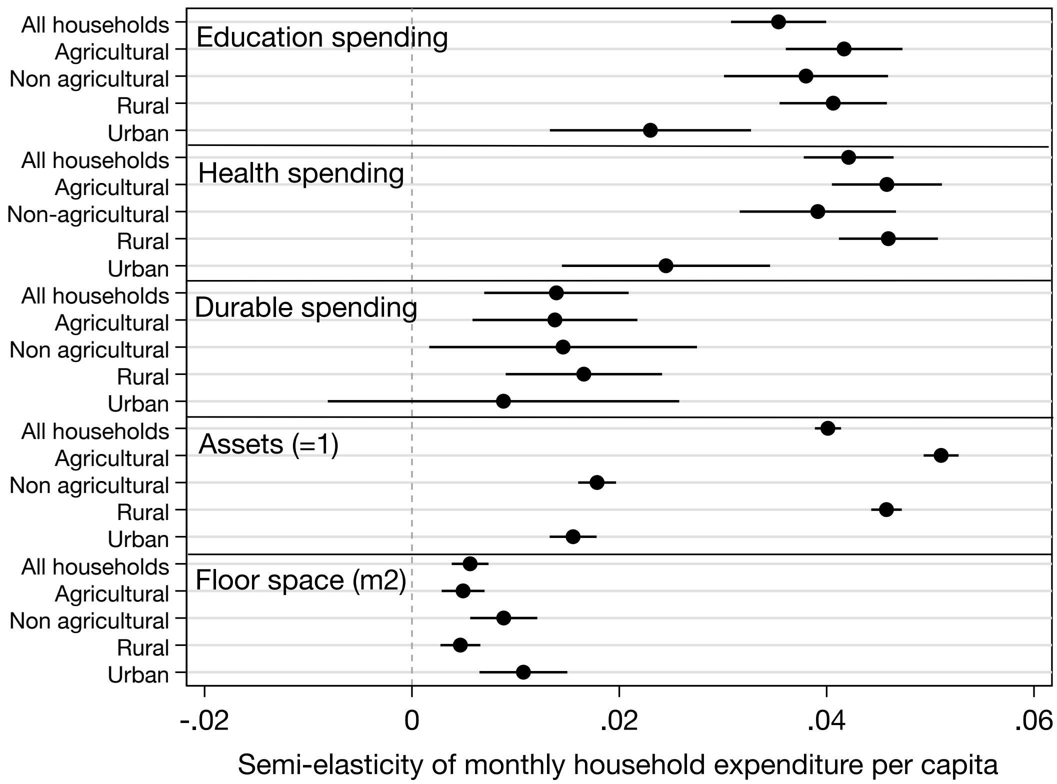

Notes: This graph plots the estimated coefficients on oil palm land from my primary IV estimator using log per capita monthly household expenditure as a dependent variable for the full sample of SUSENAS households ("All households") and for sub-groups listed on the Y axis. Black lines indicate 95% confidence intervals. The full sample is repeat cross-section of all households in SUSENAS 2002 and 2015 linked to two-period balanced panel of all rural districts at 2000 boundaries excluding Java. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator. District and year fixed effects, initial district conditions trends separately interacting 2000 log poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads with a post period dummy, and additional controls for household size, an urban/rural dummy, and sector fixed effects are included throughout. Urban/rural (sector) fixed effects are dropped when I examine effects by urban-rural households (across sectors). Floor space is in logs.

Table 1: First-stage—Suitability and Area Expansion, 2000–15

Dependent variable	Oil palm land/district area (%)				
Column	1	2	3	4	5
Post*suitability (kg/ha)	0.0019***	0.0018***	0.0014***	0.0017***	0.0018***
	(0.0003)	(0.0004)	(0.0003)	(0.0005)	(0.0006)
District and year FEs	Y	Y	Y	Y	Y
Baseline trends	Y	Y	Y	Y	Y
Poverty pre-trends		Y			Y
Additional trends			Y		Y
Lat-long polynomial trends				Y	Y
Observations	334	288	326	334	282

Notes: Sample is a balanced panel of all rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. Baseline trends separately interact 2000 log poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads with a post period indicator variable. Additional trends include differential trends related to initial levels of ethnolinguistic fractionalization, the share of villages in each district with palm farmers, district production in tons, population density, and the percentage of households with access to electricity. Lat-long polynominal interacts each district's latitude and longitude, taken at its centroid, and the squared term of each, with the post period. Column 2 controls for log-changes in district poverty, where pre-period poverty is calculated from 1993–2002, when SUSENAS became district-representative and the Asian Financial Crisis had subsided. Robust standard errors are in parentheses and clustered at the district level.

TABLE 2: MAIN RESULTS—POVERTY AND HOUSEHOLD CONSUMPTION, 2000–15

Dependent variable	District poverty rate (%)			Log expenditure (IDR)		
Estimator	OLS	IV	Reduced form	OLS	IV	Reduced form
Column	1	2	3	4	5	6
Oil palm land/district area (%)	-0.081*	-0.594***		0.001	0.009**	
	(0.042)	(0.182)		(0.001)	(0.004)	
Post*suitability ('000 kg/ha)			-1.112***			0.016*
			(0.283)			(0.008)
Excluded F statistic		33			24	
Observations	334	334	334	237,887	237,887	237,887

Notes: Sample in Columns 1–3 is a two-period balanced panel of rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability. Sample in columns 4–6 are the household observations for the same districts, with identifying variation in oil palm expansion and suitability measured at the district level. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads are included throughout. Household expenditure is measured in average, monthly, per capita terms. Household level estimates in columns 4–6 also include household size, an urban/rural dummy, and primary sector income fixed effects. Robust standard errors are in parentheses and clustered at the district level.

Table 3: Expansion onto Marginal Lands, 2000–08

Dependent variable	District poverty rate (%)						
Estimator	OLS	IV	OLS	IV	OLS	OLS	OLS
Column	1	2	3	4	5	6	7
Oil palm land/district area (%)	-0.109**	-0.993***					
	(0.054)	(0.318)					
Oil palm area / farmland (%)			-0.007	-0.907***			
			(0.051)	(0.308)			
Farmland / district area (%)					-0.034		
					(0.022)		
Oil palm area (000 ha)						-0.009*	
						(0.005)	
Farmland (000 ha)							-0.001***
							(0.000)
Excluded F statistic		23		20			
Observations	334	334	334	334	334	334	334

Notes: This table reports results from palm (SUSENAS) variation from 2000–2008 (2002–2010), half the period of my main results. It shows how point estimates are similar whether total district area or farmland is used as the denominator. Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability. Data on farmland are calculated by aggregating village farmland reported in the 2000 and 2008 village censuses (PODES) up to the district level. IV estimates instrument the district oil palm land variable of interest with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads are included throughout. Robust standard errors are in parentheses and clustered at the district level.

Table 4: Labor Productivity and Wages—IV Results

Dependent variable	Output per worker		Log wages			
Sector	Agriculture	Manufacturing	All	Agriculture	Manufacturing	Services
Column	1	2	3	4	5	6
Oil palm land/district area (%)	1.600***	6.846***	0.039***	0.076**	0.026	-0.006
	(0.578)	(1.564)	(0.013)	(0.037)	(0.023)	(0.009)
Excluded F statistic	35	32	36	20	27	35
Observations	328	298	324	234	242	322

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability in the national labor market survey SAKERNAS, which is used to calculate the outcomes. District oil palm land share is instrumented with district potential palm oil yield interacted with a post period indicator throughout. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads. Robust standard errors are in parentheses and clustered at the district level.

TABLE 5: FISCAL LINKAGES AND INFRASTRUCTURE—IV RESULTS

Panel A: Fiscal outcomes (in logs)				
Level of government	District		Village	
Dependent variable	Revenue	Expenditure	Own source revenue	Expenditure
Column	1	2	3	4
Oil palm land/district area (%)	0.039***	0.043***	0.074**	0.054**
	(0.015)	(0.014)	(0.033)	(0.032)
Excluded F statistic	29	28	26	27
Observations	266	264	44,699	70,977
Panel B: Village economic infrastructure (=1)				
Dependent variable	Clean cooking fuel	Improved road	Street light	Physical market
Column	5	6	7	8
Oil palm land/district area (%)	0.029***	0.002	0.013**	0.004***
	(0.001)	(0.511)	(0.049)	(0.004)
Excluded F statistic	27	27	27	27
Observations	82,349	82,349	82,349	82,349
Panel C: Village social infrastructure (n)				
Dependent variable	Private schools	Government schools	Health clinics	Places of worship
Column	9	10	11	12
Oil palm land/district area (%)	0.079***	0.009	0.067***	0.103**
	(0.019)	(0.013)	(0.024)	(0.045)
Excluded F statistic	27	27	27	27
Observations	82,349	82,349	82,349	82,349

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability. Identifying variation in oil palm expansion and suitability is measured at the district level, and observations and outcomes at either the district (Ministry of Finance via DAPOER) or village level (PODES). District oil palm land share is instrumented with district potential palm oil yield interacted with a post period indicator throughout. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads. Village estimates include additional village level urban, coast, hilly terrain, and primary sector of income dummies. Robust standard errors are in parentheses and clustered at the district level.

Table 6: Population and Migration Impacts—IV Results, 2000–10

Dependent variable	Log population	Lived in a different district..	
		5 years ago (=1)	at birth (=1)
Column	1	2	3
Oil palm land/district area (%)	-0.007	-0.004**	-0.005
	(0.009)	(0.040)	(0.132)
Excluded F statistic	28	10	10
Observations	270	12,703,097	12,719,478

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability. Identifying variation in oil palm expansion and suitability is measured at the district level, and observations and outcomes at household level. Outcomes are taken from the publicly-available IPUMS extracts of the 2000 and 2010 Population Censuses. District oil palm land share is instrumented with district potential palm oil yield interacted with a post period indicator throughout. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads. Robust standard errors are in parentheses and clustered at the district level.

Table 7: Environmental Impacts, 2000–16

Dependent variable	Gross tree cover loss (% district area)			Hotspot detections (N)		
Estimator	OLS	IV	RF	Poisson	IV Poisson	RF Poisson
Column	1	2	3	4	5	6
Δ Oil palm land/district area (%), 2000–15	0.008***	0.017***		0.028**	0.078***	
	(0.002)	(0.002)		(0.011)	(0.016)	
Suitability ('000 kg/ha)			0.038***			0.461***
			(0.005)			(0.155)
Observations	167	167	167	170	167	167

Notes: Sample is a cross-section of all rural districts excluding Java at 2000 district boundaries. Any changes in samples size due to data availability. Forest loss is defined as the total number of pixels of tree cover loss since 2000 as a share of total district pixels, calculated by the author from the University of Maryland Tree Cover Change dataset. Hotspots are detections per district since 2000, calculated by the author from NASA's MODIS Active Fire Products. IV estimates instrument the change in the share of each district planted with oil palm with potential palm oil yield. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages in each district with paved roads. Robust standard errors are in parentheses.

Export agriculture and rural poverty
Supplementary appendix—not for publication

Ryan B. Edwards

March 25, 2019

1 Data Appendix

1.1 Main explanatory variables

- **Share of district area under cultivation for oil palm:** District oil palm acreage is digitized from the Tree Crop Statistics of Indonesia for Oil Palm yearbooks, produced annually by the Directorate General of Estate Crops in the Department of Agriculture and Indonesia's central statistics agency *Badan Pusat Statistik* (BPS). For 2000–2010, this data is available through the World Bank's Indonesia Database for Economic and Policy Research (DAPOER) in their main databank. Bank staff digitized these annual yearbooks completely into the databank. From 2011–present, I downloaded recent yearbooks from the Department of Agriculture website and digitized them myself. Districts with no oil palm land are missing values in the original data. I recode them as zeros to retain the baseline and control districts, after cross-checking against other sources and receiving confirmation from officials that they are nationally exhaustive. The share of each district under cultivation for oil palm is calculated by dividing district acreage by district area in km^2 from BPS.
- **Palm oil suitability:** Calculated from the Food and Agriculture Organization's (FAO) Global Agro-Ecological Zones (GAEZ) dataset by mapping gridded data on crop-specific agro-climatically attainable yields to 2010 Population Census district boundaries and taking district means. See Fischer et al. (2002) for more details.
- **Farmland:** Total district farmland is calculated as the district sum of total village farmland reported in 2003 and 2008 censuses of village heads, Potensi Desa (PODES).

1.2 Outcomes

- **District poverty:** Measured as the share of district population living below an expenditure-based poverty line, roughly equal to $25 United States dollars per person per month. Poverty rates are estimated from the consumption module of BPS' National Socioeconomic Survey (SUSENAS). Implemented at least annually, SUSENAS covers over two million people across all 34 provinces in 2015 and has been district-representative since 1993. Poverty for 2015 is taken from BPS website and for 2002 and 2010 from DAPOER. The method used to calculate poverty changed in 1998 and 2011, but results are similar if 2010 data are used for the final period.

- **Per capita household expenditure:** Calculated for SUSENAS households in 2002 and 2015 as the sum expenditures from the consumption module divided by the number of household members (not adjusted for age). SUSENAS is a repeat cross-section, so the same households are not observed in 2002 and 2015.
- **District output per worker:** Calculated as the district regional gross domestic product (RGDP) for each sector divided by sector employment. RGDP figures are BPS' subnational accounts, taken from DAPOER. Employment is calculated for each district-sector from the national labor market survey SAKERNAS in 2002 and 2013 according to standard industry classifications.
- **Average district wage:** Calculated from SAKERNAS 2002 and 2013 for people over 15 employed in (a) all sectors, (b) agriculture, and (c) manufacturing.
- **District revenue and expenditure:** Taken directly from DAPOER, which in turn are taken from BPS and the Ministry of Finance. All are scaled by BPS' district population (DAPOER), interpolated across years as necessary.
- **District population:** BPS annual population estimates via DAPOER.
- **Village own source revenue and expenditure:** Taken from PODES 2003 and 2014. Some villages do not report either. I treat this data as missing at random.
- **Clean cooking fuel:** a dichotomous indicator equal to one if a village mostly uses electricity, gas, LPG, or kerosene for cooking, calculated from PODES 2003 and 2014.
- **Improved road:** a dichotomous indicator equal to one if a village's main road is made from asphalt or hardened gravel, calculated from PODES 2003 and 2014.
- **Street light:** a dichotomous indicator equal to one if a village's main road is illuminated by a street light, calculated from PODES 2003 and 2014.
- **Marketplace:** a dichotomous indicator equal to one if a village has a permanent built marketplace, calculated from PODES 2003 and 2014.
- **Recent migration status:** a dichotomous indicator equal to one if an individual reports living in a different district/province five years before enumeration in the 2000 and 2010 Population Censuses. I use the publicly-available IPUMS extracts.
- **District forest loss:** Tree cover loss is calculated from Hansen et al (2013) as the number of pixels of tree cover lost from 2000–2015 as a share of total district pixels. Each lost pixel is counted once. Although this is an imperfect measure also picking up forest loss not caused by oil palm expansion (e.g., contemporaneous logging and

the spread of wildfire), it ensures I pick up each initial change in tree cover from land conversion and excludes reforestation through natural regrowth or tree crops.

- **District hotspot detections:** Hotspots are taken from the NASA Moderate Resolution Imaging Spectroradiometer (MODIS) Active Fire Product. I count the total number of detections in each district since 2000. This provides a tractable measure capturing fire duration, scale, and overall intensity by increasing (a) in the times a given fire is observed in the same space, and (b) in multiple detections of single or related fires across pixels.

1.3 Control variables

- **Rural population share:** calculated from the 2000 Population Census, via DAPOER.
- **Over-15 literacy rate:** calculated from 2002 SUSENAS, via DAPOER.
- **Agricultural and manufacturing employment shares:** calculated from SAKERNAS 2002 using standard industry definitions.
- **Ethnic fractionalization:** Index of ethno-linguistic fractionalization calculated from the 2000 Population Census via World Bank.
- **Population density:** District population divided by area, both via DAPOER.
- **Household access to electricity:** calculated from as a share of SUSENAS households in the district, via DAPOER.
- **District palm oil production:** Calculated as the district sum of all village production in PODES 2003.
- **Plantation village share:** The share of villages in each district for which the main source of income is plantation crops, calculated from PODES 2003.
- **Lat-long polynomial:** District latitude and longitude and their squared terms are calculated for each district's centroid, according to BPS' 2010 Population Census shapefile.
- **Poverty pre-trend:** the change in district poverty from 1993 to 2002 is calculated from SUSENAS 1993 and 2002, using BPS' time-varying district poverty lines.

1.4 Constructing the main district panel

The main source of identifying variation used throughout the main paper comes from a two-period balanced panel of district level oil palm cultivation, for rural districts spanning Indonesia's outer islands in 2000 and 2015. Districts are clearly defined legal and geographical units with administrations reflecting local economies and labor markets. By contrast, the level above the district is the province, of which there are only 34. The level below is the subdistrict, not particularly important in economic or political terms. Districts and villages are the key subnational administrative units. District variation is well suited identify aggregate regional economic impacts.

Given my focus on land, I redefine district boundaries to work with constant-area spatial units. Indonesia underwent one of the world's largest reconfigurations of a modern state with the fall of President Suharto in 1997, democratizing and decentralizing power to around 300 district governments. New political and fiscal powers drove the number of districts to proliferate from 292 in 1998 to 514 in 2015, a process known as *pemekaran*. District splits followed sub-district (*kecamatan*) boundaries and did not affect neighboring districts' borders. To obtain a balanced panel of 341 constant geographic units, I apply year-2000 district boundaries from a district crosswalk tracking "parent" and "child" districts over time. Practically, this means summing level variables (e.g., area) and taking the averages of others (e.g., poverty) across proliferated units. Crosswalks can be downloaded from the BPS and World Bank websites.

2 Supplementary Tables and Figures

Table A1: Pre-expansion District Characteristics

Palm oil suitability (above/below median) Variable	Low Mean/SE	High Mean/SE	(1)-(2) t-test Difference
Poverty rate (%)	26.389 [1.315]	21.176 [1.325]	5.213***
Log per capita expenditure (IDR)	11.595 [0.026]	11.744 [0.022]	-0.148***
Over 15 literacy rate (%)	84.089 [1.515]	92.085 [0.493]	-7.996***
Agricultural employment share	0.651 [0.018]	0.632 [0.018]	0.019
Industrial employment share	0.087 [0.009]	0.105 [0.009]	-0.017
Rural population share (%)	85.352 [1.270]	81.099 [1.561]	4.253**
Population density	111.425 [13.856]	64.306 [7.378]	47.119***
Area (km^2)	10,239 [1,864]	14,062 [1,537]	-3,823
Access to electricity (%)	60.245 [2.743]	64.640 [1.741]	-4.396
Oil palm villages share (%)	0.002 [0.001]	0.008 [0.001]	-0.006***
Palm oil production (tons)	7,544 [2,877]	46,897 [13,726]	-39,400**
Ethnolinguistic fractionalization	0.496 [0.036]	0.591 [0.026]	-0.095**
Number of districts in 2015	1.975 [0.147]	1.905 [0.110]	0.069
N districts	79	96	

Notes: This table shows the observable differences between areas with high and low palm oil suitability, defined as being above or below the median agro-climatically attainable yield. Medians are calculated for all districts but difference for the estimation sample, where districts are clearly more suitable, are shown. Observations are districts in 2000 or the nearest feasible period. Variable construction and data sources are detailed in Appendix One.

Table A2: Alternative Palm Parameterizations, Poverty

LHS variable	District poverty rate							
RHS variable	Palm area per person in 2000 (ha)				Palm production per person in 2000 (tons)			
Functional form (RHS)	Level		Inverse hyperbolic sine		Level		Inverse hyperbolic sine	
Estimator	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Column	1	2	3	4	5	6	7	8
RHS palm oil variable	-2.254**	-16.460***	-2.797**	-18.519***	-0.795***	-5.683***	-1.732***	-8.808***
	(0.877)	(4.937)	(1.203)	(5.423)	(0.181)	(1.903)	(0.509)	(2.574)
District and year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y
Excluded F		32		39		18		40
Observations	340	334	340	334	340	334	340	334

Notes: This table shows that the main findings are qualitatively similar with alternative parameterizations of the main palm explanatory variable. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects are included throughout. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

Table A3: Alternative Palm Parameterizations, Consumption

LHS variable	Per capita household consumption							
RHS variable	Palm area per person in 2000 (ha)				Palm production per person in 2000 (tons)			
Functional form (RHS)	Level		Inverse hyperbolic sine		Level		Inverse hyperbolic sine	
Estimator	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Column	1	2	3	4	5	6	7	8
op_pc	0.061*	0.307*	0.090*	0.333**	0.008	0.105*	0.042*	0.148**
	(0.036)	(0.156)	(0.046)	(0.166)	(0.011)	(0.056)	(0.023)	(0.072)
District and year FEs	Y	Y	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y	Y	Y
Excluded F		21		25		15		29
Observations	237887	237887	237887	237887	237887	237887	237887	237887

Notes: This table shows that the main findings are qualitatively similar with alternative parameterizations of the main palm explanatory variable. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects are included throughout. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, sectoral employment shares, and the share of village with upgraded roads with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

Table A4: Alternative Functional Form, Log Poverty

Dependent variable	Log district poverty rate (%)		
Estimator	OLS	IV	Reduced form
Column	1	2	3
Oil palm land / district area (%)	-0.009**	-0.041***	
	(0.004)	(0.010)	
Post*suitability			-0.086***
			(0.018)
Excluded F statistic		35	
Observations	340	334	334

Notes: This table shows that the main findings are qualitatively similar if the natural logarithm of district poverty is used as the dependent variable. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects are included throughout. All estimates include baseline trends separately interacting 2000 poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

TABLE A5: ALTERNATIVE SAMPLES, POVERTY

Dependent variable	District poverty rate (%)								
Sample	All districts			No Java			No cities		
Estimator	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.058*	-0.240*		-0.052	-0.428***		-0.074**	-0.277*	
	(0.033)	(0.144)		(0.034)	(0.147)		(0.036)	(0.150)	
Post*suitability ('000 kg/ha)			-0.436*			-0.933***			-0.499*
			(0.254)			(0.282)			(0.257)
Excluded F statistic		45			41			42	
Observations	610	578	578	424	394	394	520	514	514

Notes: This table runs the main estimates for all districts, excluding Java and including cities, and including Java and excluding cities. Sample is a two-period balanced panel of districts, at 2000 district boundaries. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads are included throughout. Robust standard errors are in parentheses and clustered at the district level.

Table A6: Alternative Samples, Consumption

Dependent variable	Log per capita household expenditure (Indonesian rupiah)								
Sample	All districts			No Java			No cities		
Estimator	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	0.001***	0.006***		0.001**	0.007***		0.002***	0.006***	
	(0.000)	(0.001)		(0.000)	(0.001)		(0.000)	(0.001)	
Post*suitability			0.010***			0.016***			0.010***
			(0.001)			(0.002)			(0.001)
Excluded F statistic		33,979			24,746			32,163	
Observations	434,841	416,283	416,283	288,178	270,649	270,649	383,716	380,254	380,254

Notes: This table runs the main estimates for all districts, excluding Java and including cities, and including Java and excluding cities. Sample is the household observations for a two-period balanced panel of districts, at 2000 boundaries, with identifying variation in oil palm expansion and suitability measured at the district level. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects, differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgrade roads, and household-level controls for household size, an urban/rural dummy, and primary sector income fixed effects are included throughout. Robust standard errors are in parentheses and clustered at the district level.

Table A7: Pre-period Placebo Tests, 2000 District Definitions

Lagged outcome	Poverty rate (%)			Poverty gap index			Log expenditure		
Estimator	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.010	-0.101		-0.007	-0.053		0.001	0.003	
	(0.017)	(0.079)		(0.006)	(0.037)		(0.002)	(0.008)	
Post*suitability ('000 kg/ha)			-0.206			-0.108			0.006
			(0.158)			(0.074)			(0.016)
Excluded F statistic		29			29			35	
Observations	294	288	288	294	288	288	340	334	334

Notes: This table uses outcomes before the expansion to conduct falsification tests, i.e., in-time placebo tests. Sample is a two-period balanced panel of rural districts excluding Java at 2000 district boundaries, with any changes in samples size due to data availability. Dependent variables are observed in 1993 and 2002 and explanatory variables 2000 and 2015. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads are included throughout. Household expenditure is measured in average, monthly, per capita terms, and RF refers to the reduced form of the IV estimate. Robust standard errors are in parentheses and clustered at the district level.

Table A8: Pre-period Placebo Tests, 1993 District Definitions

Lagged outcome	Poverty rate (%)			Poverty gap index			Log expenditure		
Estimator	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.006	-0.105		-0.002	-0.057		0.001	0.002	
	(0.017)	(0.086)		(0.007)	(0.043)		(0.002)	(0.009)	
Post*suitability ('000 kg/ha)			-0.203			-0.111			0.004
			(0.163)			(0.081)			(0.018)
Excluded F statistic		25			25			31	
Observations	236	230	230	236	230	230	272	266	266

Notes: This table uses outcomes before the expansion to conduct falsification tests, i.e., in-time placebo tests. Sample is a two-period balanced panel of rural districts excluding Java at 1993 district boundaries, with any changes in samples size due to data availability. Dependent variables are observed in 1993 and 2002 and explanatory variables 2000 and 2015. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads are included throughout. Household expenditure is measured in average, monthly, per capita terms, and RF refers to the reduced form of the IV estimate. Robust standard errors are in parentheses and clustered at the district level.

Table A9: Alternative Instrumental Variables, Poverty

Dep. Var.	District poverty rate (%)					
Instrument	Palm suitability		Relative to cash crops		Relative to all crops	
Estimator	IV	RF	IV	RF	IV	RF
Column	1	2	3	4	5	6
Oil palm area	-0.594***		-0.267**		-0.453***	
	(0.182)		(0.107)		(0.150)	
Suitability		-5.773***		-3.381**		-4.926***
		(1.468)		(1.357)		(1.408)
Excluded-F	33		32		34	
Observations	334	334	334	334	334	334

Notes: This table shows that the main findings are relatively insensitive to alternative suitability based instruments, specifically relative to the normalized difference between palm suitability and cash crops, and between palm and all other crops. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district normalized palm yield (or the differences described above) interacted with a post period indicator, and district and year fixed effects are included throughout. All estimates include controls separately interacting 2000 poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

Table A10: Alternative Instrumental Variables, Consumption

Dep var	Log per capita household consumption			
Instrument	Palm suitabiltiy		Relative to cash crops	
Estimator	IV	RF	IV	RF
Column	1	2	3	4
Oil palm area	0.008**		0.010***	
	(0.004)		(0.003)	
Suitability		0.079*		0.125***
		(0.042)		(0.039)
Excluded F	24		27	
Observations	237887	237887	237887	237887

Notes: This table shows that the main consumption findings are relatively insensitive to an alternative suitability based instrument, the normalized difference between palm suitability and cash crops. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district normalized palm yield (or the differences described above) interacted with a post period indicator, and district and year fixed effects are included throughout. All estimates include controls separately interacting 2000 poverty, rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

Table A11: Additional Trends Sensitivity Analysis, Poverty

Dependent variable	District poverty rate (%)						
Column	1	2	3	4	5	6	7
Oil palm land/district area (%)	-0.672***	-0.594***	-0.626***	-0.758***	-0.486**	-0.534*	-0.660**
	(0.219)	(0.182)	(0.206)	(0.243)	(0.236)	(0.273)	(0.281)
District and year fixed effects	Y	Y	Y	Y	Y	Y	Y
Baseline trends		Y	Y	Y	Y	Y	Y
Pre-period poverty trends			Y				
Additional trends				Y		Y	
Lat-long polynomial trends					Y	Y	
Island-by-year fixed effects							Y
Excluded F statistic	39	33	26	22	10	10	13
Observations	350	334	288	326	334	326	334

Notes: This table shows that the main findings are relatively insensitive to additional controls designed to capture potential violations of the exclusion restriction. Sample is a balanced panel of rural districts in 2000 and 2015, at 2000 district boundaries, excluding cities and Java. Changes in samples size are due to data availability. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects and initial poverty trends are included throughout. Baseline trends separately interact 2000 rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Pre-period poverty refers to the change from 1993–2002. Additional trends include differential trends related to initial levels of ethnolinguistic fractionalization, the share of villages in each district with palm farmers, district production in tons, population density, and the percentage of households with access to electricity. Lat-long polynominal interacts each district's latitude and longitude, taken at its centroid, and the squared term of each, with the post period. Island groups are defined as Sumatra, Kalimantan, Sulawesi, and Eastern Indonesia. Robust standard errors are in parentheses and clustered at the district level.

Table A12: Additional Trends Sensitivity analysis, Consumption

Dependent variable	Log per capital household consumption (IDR)						
Column	1	2	3	4	5	6	7
Oil palm land / district area (%)	0.008***	0.008**	0.009**	0.014**	0.003	0.004	0.005
	(0.003)	(0.004)	(0.004)	(0.006)	(0.004)	(0.007)	(0.005)
District at year fixed effects	Y	Y	Y	Y	Y	Y	Y
Baseline trends		Y	Y	Y	Y	Y	Y
Pre-period poverty trends			Y				
Additional trends				Y		Y	
Lat-long polynomial trends					Y	Y	
Island-by-year fixed effects							Y
Excluded F statistic	30	24	23	16	8	9	11
Observations	245375	237887	213605	230999	237887	230999	237887

Notes: This table shows that the main consumption findings are relatively insensitive to additional controls, as long as there is not a weak instrument problem. Sample is household observations in 2002 and 2015 with identifying variation in oil palm expansion and suitability measured at the district level, based on 2000 district boundaries, excluding cities and Java. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects and initial poverty trends are included throughout. Household controls cover household size, an urban/rural dummy, and primary sector of income fixed effects. Baseline trends separately interact 2000 rural population shares, literacy rates, sectoral employment shares, and the share of villages with upgraded roads with a post period indicator indicator. Additional trends include differential trends related to initial levels of ethnolinguistic fractionalization, the share of villages in each district with palm farmers, district production in tons, population density, and the percentage of households with access to electricity. Lat-long polynominal interacts each district's latitude and longitude, taken at its centroid, and the squared term of each, with the post period. Island groups are defined as Sumatra, Kalimanta, Sulawesi, and Eastern Indonesia. Robust standard errors are in parentheses and clustered at the district level.

Table A13: Zero First Stage Falsification Test

Dependent variable	District poverty rate			Per capita household consumption		
Sample	Main sample	Java	Java and east	Main sample	Java	Java and east
Column	1	2	3	4	5	6
Post*suitability	-1.112***	1.080**	-0.346	0.016*	-0.017	-0.018**
	(0.283)	(0.472)	(0.325)	(0.008)	(0.016)	(0.009)
District and year FEs	Y	Y	Y	Y	Y	Y
Baseline controls	Y	Y	Y	Y	Y	Y
Observations	334	180	240	237887	142367	181612

Notes: This table shows that the positive link between palm suitability and poverty reduction is not present in non-producing areas where the first-stage is effectively zero. Column 2 shows that this relationship in fact reverses, suggesting that if anything suitability may attenuated the results, as more palm suitable areas on Java tend to be poorer. Sample is a two-period balanced panel of all rural districts excluding Java in 2002 and 2015, at 2000 district boundaries. Palm oil suitability is interacted with an indicator period for 2015, and district and year fixed effects are included throughout. Any changes in samples size are due to changes in the regions used in the sample, per the row descriptions. Java and eastern Indonesia are districts not on Sumatra, Kalimantan, or Sulawesi. Cluster-robust standard errors are in parentheses.

Table A14: Alternative Temporal Bandwidth, Short-run Poverty Impacts

Dependent variable: log district poverty rate							
Estimator	First difference			Within FE			
Columns	1	2	3	4	5	6	7
Palm land share	-0.003**	-0.004***	-0.004**	-0.007***	-0.005**	-0.007***	-0.006***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
First lag		-0.003	-0.002		-0.003		
		(0.002)	(0.003)		(0.003)		
Second lag		-0.009***	-0.008**		-0.007**		
		(0.003)	(0.004)		(0.003)		
Third lag		-0.006***	-0.005*		-0.005		
		(0.002)	(0.003)		(0.003)		
Σ coefficients		-0.022***	-0.019***		-0.020***		
District fixed effects	N	N	Y	Y	Y	Y	Y
Island-specific time trends	N	N	N	N	N	Y	N
Province-specific time trends	N	N	N	N	N	N	Y
Observations	3040	2371	2371	3386	2717	3386	3386

Notes: Heteroskedasticity-robust standard errors are in parentheses, clustered at the district level. Sample is an annual 341 district panel from 2002–2010, taken directly from the World Bank's DAPOER database. Oil palm land is lagged one period. 2001 district boundaries are used, with new districts collapsed into year 2001 parent districts. Changes in sample size are due to data availability. Island–year fixed effects are included throughout, with island groupings defined Java, Sumatra, Kalimantan, Sulawesi, and Eastern Indonesia. The within FE estimator refers to the mean differenced (within-district) fixed effects estimator. Significance reported for sum of the coefficients relates to the test that the sum of the coefficients on oil palm is equal to zero.

Table A15: Within Sector Heterogeneity, All Palm Vs. Smallholder Palm

Dependent variable	District poverty rate			
Estimator	OLS	IV	OLS	IV
Column	1	2	3	4
Palm / district area	-0.081**	-0.536***		
	(0.040)	(0.160)		
Smallholder palm / district area			-0.156**	-1.390***
			(0.065)	(0.449)
Excluded-F		34.866		24.109
Observations	340	334	340	334

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with changes in samples size due to data availability. IV estimates instrument district oil palm land share with district potential palm oil yield interacted with a post period indicator. District and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, and sectoral employment shares are included throughout. Robust standard errors are in parentheses, clustered at the district level.

Table A16: District Government Revenue Impacts, By Type

Revenue type	DAK	DAU	DBH SDA	OSR	Other	Total	DBH Pajak
Column	1	2	3	4	5	6	7
Oil palm land / district area (%)	-0.165**	0.029**	-0.001	0.005	0.059	0.039***	0.055***
	(0.077)	(0.014)	(0.064)	(0.022)	(0.073)	(0.015)	(0.020)
Observations	198	266	262	266	206	266	266

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size are due to data availability. District oil palm land share is instrumented with district potential palm oil yield interacted with a post period indicator throughout. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, and sectoral employment shares. Revenue data are taken from the World Bank databank (DAPOER), and described further there. Robust standard errors are in parentheses. DAK=Total Special Allocation Grants. DAU=Total General Allocation Grant. DBH SDA=Total Natural Resource Revenue Sharing. OSR=Total Own Source Revenue (PAD). DBH Pajak=Total Tax Revenue Sharing. All are in IDR.

Table A17: District Government Expenditure Impacts, By Type

Log expenditure type	Admin	Agriculture	Economy	Education	Environment	Health	Housing	Infrastructure
Column	1	2	3	4	5	6	7	8
Oil palm area (%)	0.008	0.067***	0.140**	0.046	0.131**	0.037*	0.096	0.074***
	(0.019)	(0.025)	(0.070)	(0.029)	(0.058)	(0.021)	(0.079)	(0.028)
Observations	248	262	262	262	256	262	204	262
Log expenditure type	Social prot.	Law & order	Tourism	Goods & services	Other	Personnel	Total	Total, non log
Column	9	10	11	12	13	14	15	16
Oil palm area (%)	0.099**	0.059	0.122**	0.043**	-0.033	0.031**	0.043***	177155.763**
	(0.047)	(0.039)	(0.053)	(0.021)	(0.025)	(0.013)	(0.014)	(71338.617)
Observations	230	248	246	264	264	264	264	270

Notes: Sample is a two-period balanced panel of all rural districts excluding Java at 2000 district boundaries, with any changes in samples size are due to data availability. District oil palm land share is instrumented with district potential palm oil yield interacted with a post period indicator throughout. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, and sectoral employment shares. Expenditure data are taken from the World Bank databank (DAPOER), and described further there. Specifically, names are shortened to a more general expenditure function (e.g., housing is housing and public facilities functions). Robust standard errors are in parentheses.

TABLE A18: IMPACTS ON HOUSEHOLD ACCESS TO ELECTRICITY, SUSENAS

Road outcome	Has electricity (=1)			Has PLN electricity (=1)		
Estimator	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6
Oil palm land (%)	-0.0009	0.0079		0.0006	0.0107**	
	(0.0012)	(0.0051)		(0.0013)	(0.0054)	
Post * suitability			0.0149			0.0201*
			(0.0098)			(0.0102)
Excluded F statistic		24			24	
Observations	241,350	237,888	237,888	241,350	237,888	237,888

Notes: This table presents impacts on household access to electricity, for any electricity and for electricity provided by the main electricity company PLN. Sample is household observations in 2002 and 2015 with identifying variation in oil palm expansion and suitability measured at the district level, based on 2000 district boundaries, excluding cities and Java. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects are included throughout. Additional household controls cover household size, an urban/rural dummy, and primary sector of income fixed effects are included in all estimates, as are district trends separately interact 2000 poverty, rural population shares, literacy rates, and sectoral employment shares with a post period indicator indicator. Robust standard errors are in parentheses and clustered at the district level.

Table A19: Province-level Results

Dependent variable: log district poverty rate (%)			
Estimator	FE	FE	LD
Column	1	2	3
Oil palm land / district area (%)	-0.014 *	-0.007**	-0.013**
	(0.006)	(0.003)	(0.004)
Linear island trends	Y	N	N
Island-by-year fixed effects	N	Y	N
Island fixed effects	N	Y	Y
Observations	319	319	30

Notes: This table shows that least squares magnitudes are similar when aggregating up to the province level. Sample is an annual balanced panel of Indonesian provinces from 2002–2010, with oil palm land lagged one period. Data are taken from the World Bank's DAPOER databank. Estimates use a within estimator with province fixed effects (FE) and a long difference least squares estimator (LD). Heteroskedasticitiy-robust standard errors are in parentheses.

Figure A1: Migration Status, By Education

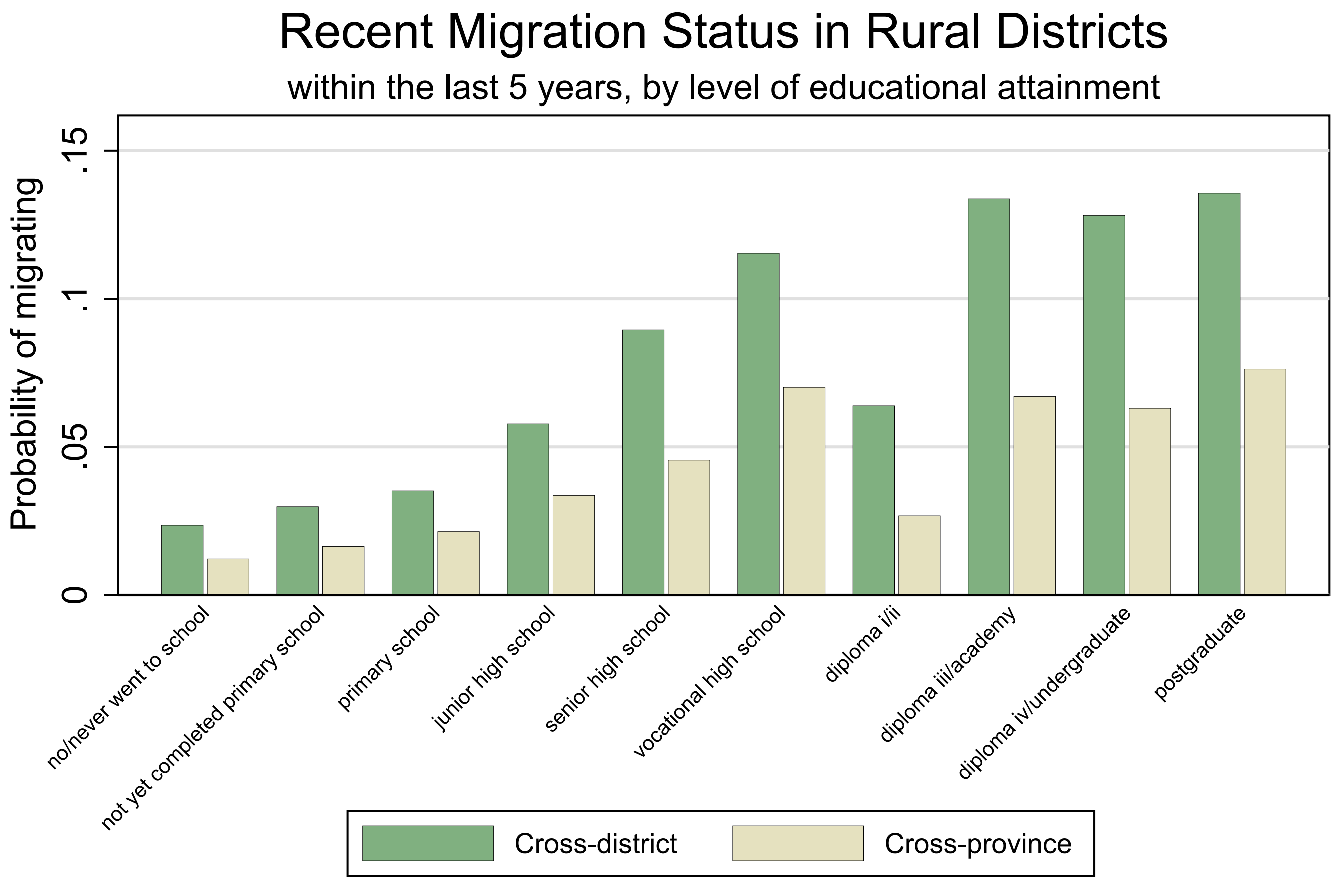

Notes: This graph plots migration status reported in the 2010 Population Census by level of education. Data are for a restricted sample of all rural districts not on the island of Java, from the ten percent sample available publicly via IPUMS.

Figure A2: Migration Status, By Education and Suitability

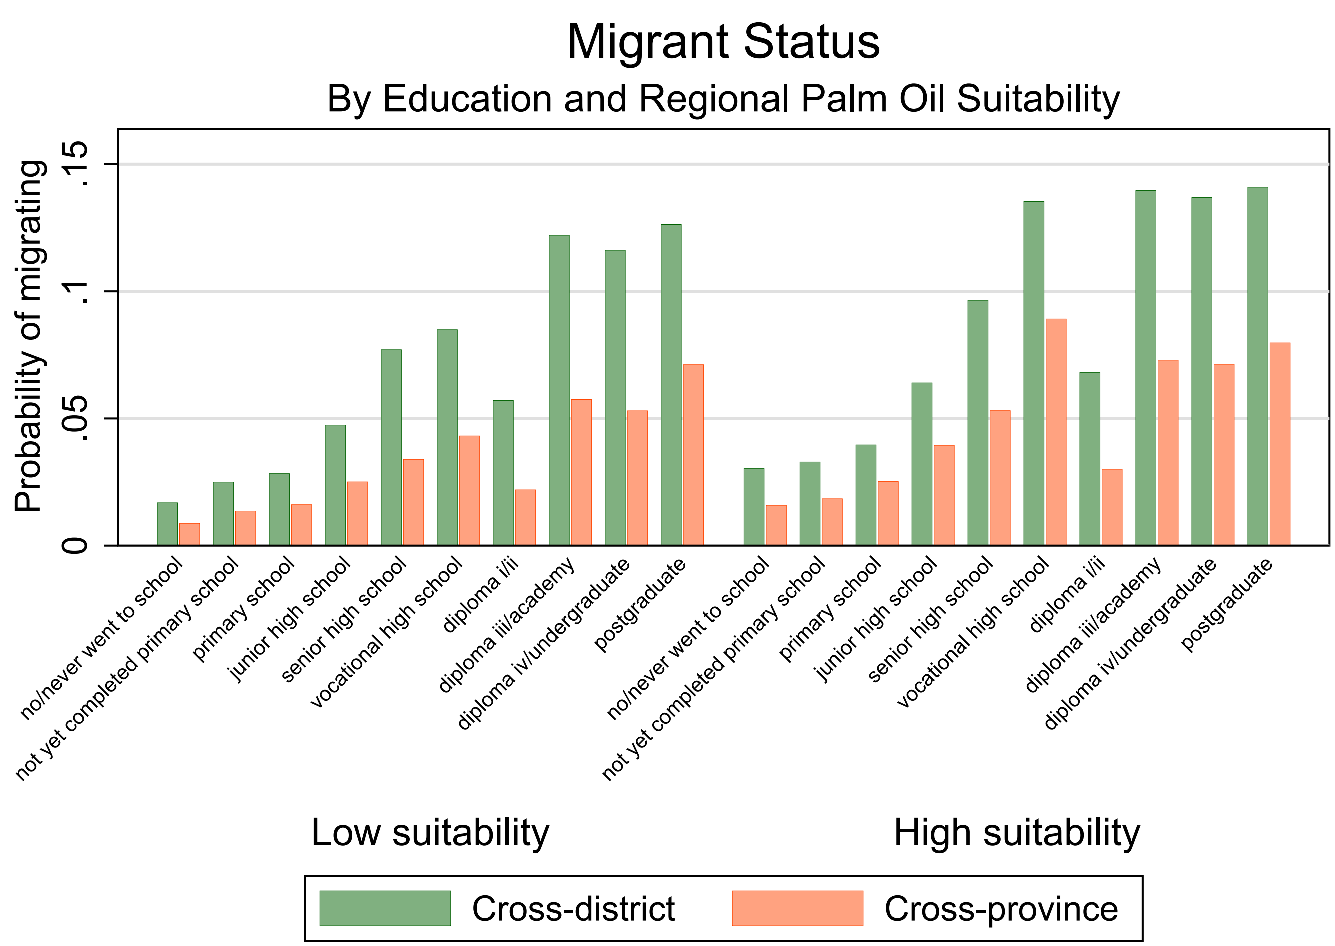

Notes: This graph plots migration status reported in the 2010 Population Census by level of education. Data are for a restricted sample of all rural districts not on the island of Java, from the ten percent sample available publicly via IPUMS.

Figure A3: District Poverty Rates in 2015

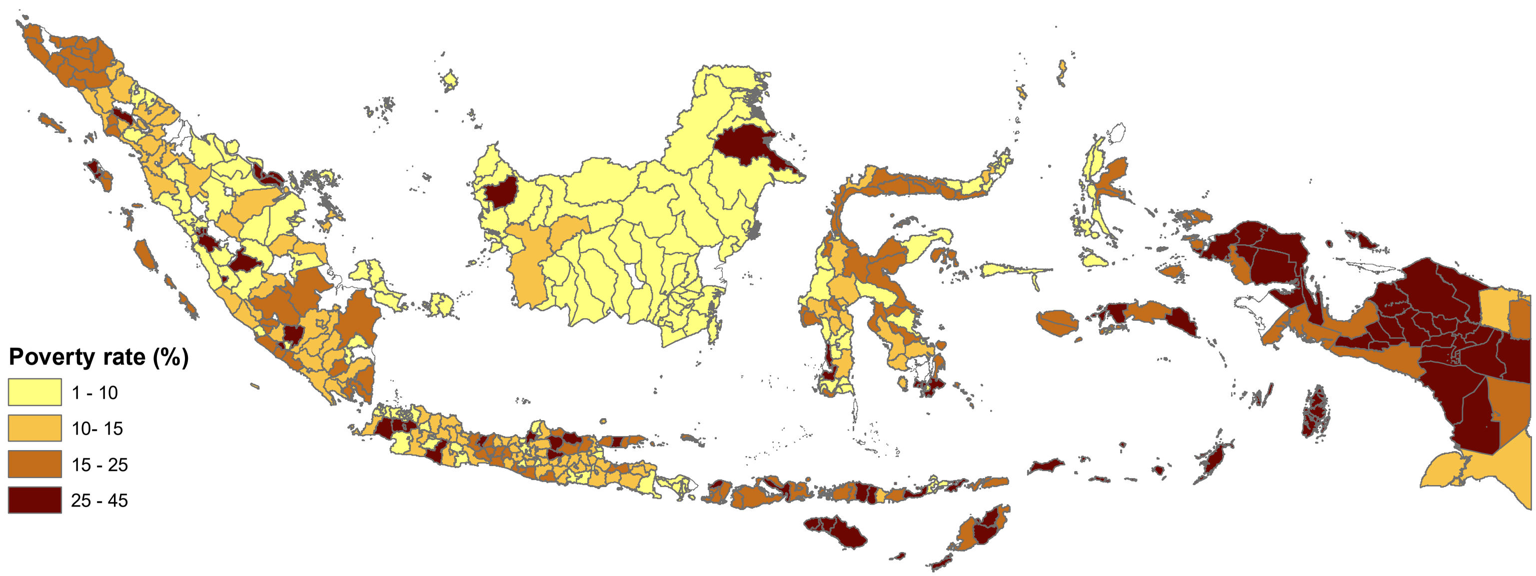

Notes: Data are official district poverty rates in 2015 from BPS.

Figure A4: Agro-climatically Attainable Palm Oil Yields, FAO-GAEZ Gridded Data

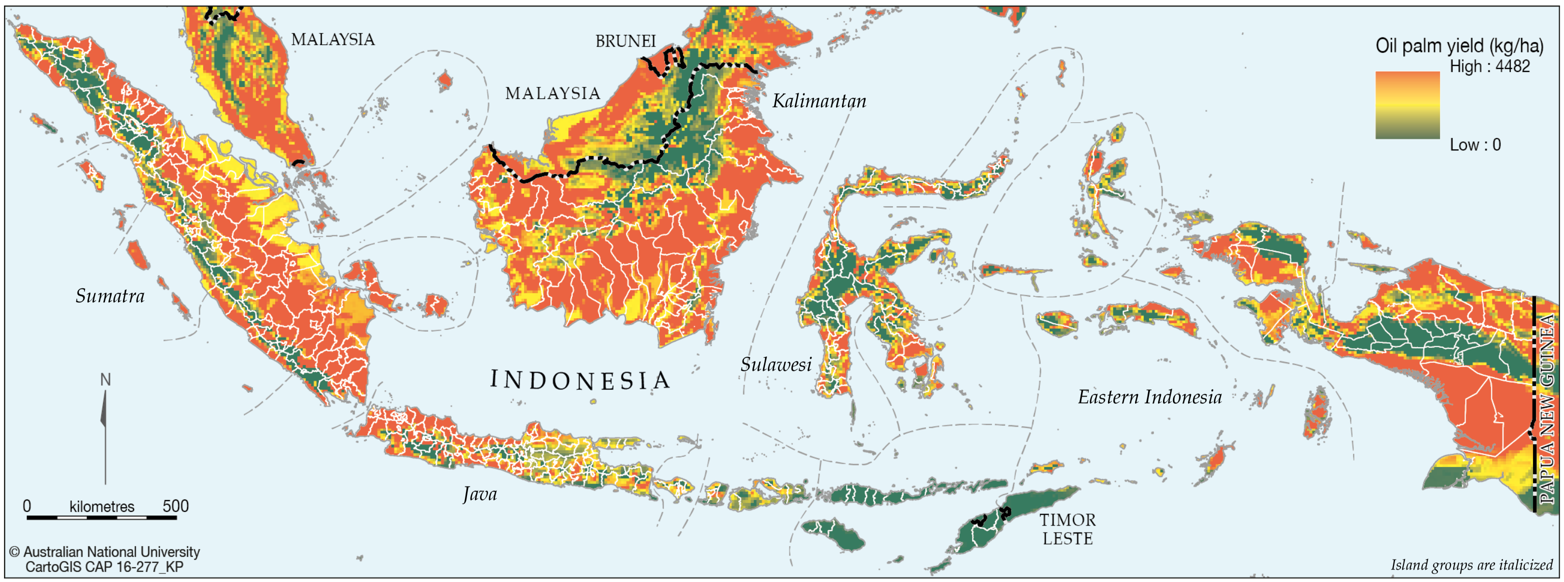

Notes: Data are gridded agro-climatically attainably yields for palm oil from the FAO-GAEZ database.