

Export agriculture and regional development: evidence from Indonesia

Ryan B. Edwards*
Dartmouth College

August 22, 2018

Abstract

This paper measures the impacts of the world’s largest modern agricultural export expansion—that of Indonesian palm oil since 2000—on regional poverty and consumption. Identification exploits geographic differences in suitability for cultivation and rapid growth in global demand. The median expansion led to 2.7 percentage points faster poverty reduction and 4 percent faster consumption growth. Results can be explained by rising returns to labor and land, and indirect effects through investment, fiscal linkages, and public goods. Each percentage point of palm-driven poverty reduction corresponds to a 1.5–3 percentage point loss of forest area.

*Contact address: Department of Economics, 6106 Rockefeller Hall, Hanover, NH 03755 USA. Email: ryan.b.edwards@dartmouth.edu. I thank Eric Edmonds, Paul Burke, Wally Falcon, Lisa Cameron, Joann de Zegher, Doug Gollin, Grace Hadiwidjaja, Matt Higgins, Roz Naylor, Umbu Raya, Budy Resosudarmo, Na’ama Shenhav, Robert Sparrow, Sudarno Sumarto, Daniel Suryadarma, Yessi Vadila, Matt Wai-Poi, Bill Wallace, Agung Widodo, and seminar participants at the Vatican, Stanford, Duke, ANU, QUB, Monash, CIFOR, Bank Indonesia, World Bank, and the 2017 AEA Annual Meetings for helpful comments and discussions. I am also grateful to the ANU Indonesia Project, Stanford SEED, Wahli, Wilmar, Asian Agri, and WWF-Indonesia for generously hosting me during field visits.

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 commercial farms—is unhelpful for the poor remains widely held (Byerlee, de Janvry, and Sadoulet, 2009; Easterly, 2007; Engerman and Solokoff, 2002). 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 palm oil production since 2000 is the world’s largest modern agricultural expansion. 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 and the European Parliament voted to ban palm oil imports for biofuels in 2017.

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. Indonesia’s recent

expansion is useful here, as 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 demand 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 regions 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 10 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. The magnitude of the effects 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 increased palm cultivation.

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. Thus, 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 and higher agricultural wages confirm this interpretation, as do higher cultivation-poverty elasticities for smallholder cultivation.

Three indirect channels reinforce the direct labor income gains. First, rising household expenditures are concentrated on health, education, and durable good expenditures, which in turn correspond to more household assets and floorspace. I interpret these changes as evidence of 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 considerably faster in expanding regions. 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. Districts that expanded their oil palm acreage have improved electrification, increased use of modern cooking fuels, better roads, and more marketplaces—all often publicly provided. Similar results when controlling for local government revenue and expenditure suggest that at least some these changes are due to the new infrastructure needed to process, transport, and export palm oil in undeveloped rural areas without any prior industry. Such complementary infrastructure could allow economic returns to ratchet up over time, as Donaldson (2018) finds for colonial India and Dell and Olken (2018) for Dutch sugar processing on Java.

This study contributes to three main streams of economics. First, I add to the trade literature new evidence on the distributional impacts of agricultural export growth (Topalova, 2010; Kis-Katos and Sparrow, 2015; Dix-Carneiro and Kovak, 2016; see Goldberg and Pavcnik (2007) for a review). The traditional view is that development strategies based on agricultural exports and commercial farming tend to be harmful for the poor (Engerman and Solokoff, 2002; Easterly, 2007; World Bank, 2008). Far from adverse effects, the poverty reduction and consumption growth I find here is more consistent with Costa, Garret, and Pessoa's (2016) findings on Brazilian export growth and Dell and Olken's (2018) historical findings for Indonesian sugar.

My study also contributes new casual evidence to the large body of work on agriculture and economic development (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 its 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, they 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.

The third 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). 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, and around ten percent more fire. These findings imply that communities—broadly defined as governments, firms, and farmers—prefer the income gains from converting forest to agriculture. Compensation to do otherwise, for example through payments for ecosystem service and other market-based mechanisms, would likely need to offset the income gains and poverty reduction to succeed.

The next section describes Indonesia’s palm oil expansion and Section 3 develops the empirical strategy. Section 4 discusses the main findings on poverty and household consumption. Section 5 explores three potential explanations for the main findings, and Section 6 quantifies environmental-economic trade-offs. Section 7 concludes.

2 Indonesia’s palm oil expansion

Palm oil is derived from the pulp of the fruit of the oil palm, a labor-intensive tree 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 cultivation is one of the most economically attractive uses for land in the tropics. Due to sustained growth in emerging economies, global demand increased from less than 5 million metric tonnes per year in 1970 to over 70 million in 2015, and is expected to double 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 and palm oil has been Indonesia’s largest agricultural export for the last two decades. Rising global demand, the devalued rupiah from the Asian financial crisis in 1998, and the subsequent “big bang” democratization and decentralization reform provided the ideal macroeconomic and institutional conditions for the expansion (Rada, Buccola, and Fuglie, 2011; Naylor, Higgins, Edwards, and Falcon, 2018).

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, and (b) new farmland from scrub, degraded land, or forest—that is, expanding the agricultural frontier.¹ The total area under cultivation for oil palm increased from 2.9 million

¹Gaskell (2015) estimates that 92% of the increase in production is due to land expansion and the remaining 8% due to yield improvements.

hectares in 1997 to over 12.5 million today. 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 studied here thus involves parallel growth in industrial and smallholder farms.² Smallholder farms are usually 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.³ The scale of area expansion and rapid smallholder growth was in no small part enabled by the devolution of power, resources, and responsibilities to regional governments. Decentralization liberalized land use, allowing regional 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).

The process from planting to exporting is characterized by long lead times. Smallholders need time to switch livelihood, prepare land, plant trees, and wait for the first harvest around two and a half years later. Production on industrial estates is characterized by similar lags. Trees take five to seven years to reach a productive state. Replanting occurs after around 25 years, when yields begin to decline and fruit becomes difficult to reach.⁴ Expansion is thus mostly determined by future demand (i.e., over the medium to long term) and alternative rural livelihood opportunities, rather than short-term changes in socioeconomic conditions or commodity prices.

²Large and small farms are usually geographically close. Smallholders need a mill close by to process and market their fruit and firms depend on smallholder supply.

³In 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). A large share plasma farmers were relocated from Java as part of the national transmigration program. Bazzi, Gaduh, Rothenberg, and Wong (2016) study the transmigration program and its modern-day impacts.

⁴The 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). All estimates include district fixed effects to capture any systematic differences across markets.

Geographic differences in growing conditions have 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. Reasonably 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. The median expansion from 2000–15 was around five percent of district area, or 42,000 hectares.⁵

3 Empirical strategy

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 in my analysis, 2000 because it predates the expansion and 2015 as the present, and estimate:

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

$P_{d,t}$ is the share of district area used for oil palm farming in 2000 and 2015.⁶ The temporal bandwidth of 15 years allows lags from planting, to harvesting, to exporting.

⁵This 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 in area planted is 1 percent of district area (6,500 ha). 60 of the 179 rural districts (at 2000 district boundaries, excluding Java) did not expand their area under cultivation from 2000–2015.

⁶Total 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 at the Department of Agriculture. District palm acreage is divided by total district area to scale cultivation intensity by district size. Results are qualitatively similar if not scaled by district area (available on author request), consider in terms of production instead of area (available on author request), or if only smallholder cultivation is considered (Tables A8 and A9).

$y_{d,t}$ 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. Appendix One details variable construction and data sources.

District fixed effects (FEs) δ_d absorb district-specific heterogeneities affecting the local extent of adoption, including geography and climate; historical, cultural, and political institutions; and government policies. δ_t is a 2015 dummy capturing common changes. To account for different development trajectories and convergence dynamics across regions, $X_{d,2000}$ includes initial rural population shares, literacy rates, agricultural and manufacturing employment shares, and poverty rates, all interacted with a post period indicator. Results thus depend upon comparisons between districts with the same initial initial level of development, urbanization, and distribution of employment across sectors.⁷ Robust standard errors are clustered at the district level.

I modify the two-period district panel in three ways to 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.⁸ Second, I remove cities, where little palm is grown but palm oil companies are often headquartered. Third, I eliminate Java which grows little palm to compare only rural districts across outer islands.⁹ Finally, I stress that Equation 1 identifies the impacts of agricultural expansion in districts that expanded cultivation relative to those that expanded less or not at all. It 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.

⁷This set of controls is standard in the literature, following, for example, Topalova (2010), Bustos et al (2016), and Dix-Carneiro and Kovak (2016). Here, they also reflect observable differences between high and low suitability districts (Table 1). Results are similar without these controls and with a wider range of controls (Tables A1 and A2).

⁸Decentralization 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, highlighting how districts splits followed subdistrict boundaries and did not affect neighboring borders.

⁹Tables A1, A2, A6, and A7 show that my main results are qualitatively similar with Java, with cities, and with island-by-year FEs.

3.1 Identification

I exploit geographic differences in suitability and the external demand shock to address the endogenous nature of the area planted with oil palm. Suitability is measured as average district agro-climatically attainable palm oil yield, calculated from the Food and Agriculture Organization’s (FAO) Global Agro-Ecological Zones (GAEZ) dataset.¹⁰ To construct my instrument, 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.

The first stage intuition is that higher potential yields increase the likelihood of developing palm processing infrastructure and planting trees.¹¹ Panel A of Figure 3 shows this graphically, with a binned scatter plot of potential palm yields against the share of 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 potential palm yields do not affect poverty through any channel other than palm cultivation. Clearly, the primary channel for potential palm yields to affect economic outcomes must be through growing palms. However, one might still be concerned that highly suitable districts differ in other ways potentially correlated with adoption and development trajectories. Table 1 presents baseline—i.e., 2000 or as close as feasible to 2000—descriptive statistics, in levels (cf., trends), for sample districts above and below the median palm yield.¹² Highly suitable districts are observably different in terms of poverty, literacy, and rural population already included in $X_{d,2000}$, as well as in other ways.

¹⁰GAEZ 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. Instead, it 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 (1.7 million grid cells across the Earth). See Costinot, Donaldson, and Smith (2016) and Fischer, van Nelthuisen, Shah, and Nachtergaele (2002) for details.

¹¹Qualitative evidence gathered from interviews and focus group discussions with firms suggests that suitability is the first-order concern when developing plantation infrastructure. Farmers are also highly attuned to the relative profitability of adopting, usually from observing neighbors.

¹²Variables shown are those standard in the literature (Bustos et al, 2016; Dix-Carneiro and Kovak, 2016), plus measures of initial palm cultivation intensity, ethnicity, and district splits (Burgess, Hansen, Olken, Potapov, and Sieber, 2012).

I provide three types of evidence supporting the exclusion restriction. First, I show that estimates are not particularly sensitive to the inclusion of additional trends (see Table 2, A1, and A2). For example, a key input to the palm oil GAEZ productivity model (e.g., rainfall) could affect productivity of similar tropical crops and therefore welfare through unrelated changes in production in other agricultural sectors. However, estimates are similar if I include potential yields for other key cash crops with similar growing conditions.¹³ For remaining differences and geographically-distributed unobservables not captured by $X_{d,2000}$, I show that estimates are similar if I include a host of additional trends, including a full polynomial in latitude and longitude interacted with the post period indicator.¹⁴

The second set of checks I provide are falsification tests asking whether pre-period trends in poverty and consumption are statistically related to subsequent oil palm expansion and my instrument (Tables A3 and A4). The absence of any statistically significant “placebo” effects suggests that my main estimates are not picking up unobserved preexisting trends. I also follow Nunn and Wantchekon (2011) and provide two additional falsification tests exploiting the reduced form (Table A5). Specifically, I show that (a) the reduced form relationship between suitability and poverty only exists in palm growing regions, and (b) economically and statistically significant effects only emerge for palm suitability in a model saturated with the potential yields of many other crops.

Table 2 presents first stage results. Column 1 includes the instrument, district and year FEs, and five 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

¹³Note that for potential yields of other crops to pose a threat for identification, they must also “switch on” over the 2000s and be correlated with the increasing salience of potential palm yields over the 2000s. Agro-climatic suitability data for rubber, replaced in many areas by oil palm, is unavailable in the GAEZ dataset.

¹⁴Results are similar including island-by-year FEs (see Tables A1 and A2). However, my default specification is the minimal specification first described. Given the exogeneity of the GAEZ data and the evidence presented throughout this section and the appendix, additional covariates do not appear important for identification. The minimal specification provides more conservative estimates than several of these more saturated models and the strongest first stage identification, which seems crucial to maintain as I move to the household and individual level and the excluded F statistics become less informative due the district-level identifying variation.

shows this graphically. Column 2 of Table 2 adds differential trends related to cocoa, coffee, and tea. The coefficient on palm yield is statistically indistinguishable from that in Column 1, suggesting that potential productivity of other agriculture is not a major concern for identification. Column 3 saturates the model with additional trends related to the remaining differences in Table 1—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. Column 4 adds the change in the district poverty rate over the 1990s to explicitly factor in pre-trends, with the smaller sample reflecting 1993 district definitions. Column 5 replaces these trends with a complete polynomial in latitude and longitude interacted with a post indicator, to purge remaining geographic confounding. Across these demanding specifications, the point estimate is stable, standard errors small, and first stage robust.¹⁵

Exploiting the variation in expansion arising from crop-specific agro-climatic suitability isolates the effects of developing oil palm on land where natural agro-climatic characteristics are best, not other sources of profitability like market access, trade costs, or input (e.g., land and labor) costs. This local average treatment effect (LATE) may be different to those relating to these other sources of profitability, adopting in places less suitable, the average treatment effect (ATE), and reduced-form estimates, which understate adoption effects by including suitable districts that do not cultivate much palm.¹⁶

4 Regional 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 4 shows a simplified

¹⁵Analogous second-stage estimates are in Tables A1 and A2.

¹⁶Consistent LATE estimation requires expansion to increase monotonically with suitability. Panel B of Figure 3 shows this. Higher potential yields are unlikely to push districts to cut back production, at least not during the period of this study while palm is a productive use of land relative to alternatives.

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 and the national district average. Rural districts had similar poverty levels in the early 2000s, but districts more intensively increasing palm oil production diverged as the decade progressed.

Table 3 presents the main regression estimates of the impacts of agricultural expansion 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 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. All include district and year FEs and separate trends for different initial levels of economic development, urbanization, and labor market structures. The OLS point estimate on oil palm land in Column 1 is -0.081, but increases in magnitude to -0.536 when instrumented with $\text{post} \times \text{suitability}$. This means that 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. The increase from Column 1 to Column 2 indicates a downward biased least squares estimate (e.g., due to planting in areas with lower land costs or weaker institutions), the LATE being larger than the ATE, or some combination of both. The reduced form relationship between suitability 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. These effects are not trivial compared to the overall decline in national poverty from 18.2% to 11.2% from 2002–2015, and echo Suryahadi, Suryadarma, and Sumarto (2009) on the continued importance of agricultural growth for rural poverty reduction in Indonesia.

Columns 4–6 of Table 3 present estimates on average per capita household expenditure. Although incomparable, the OLS coefficient of 0.001 again illustrates the biases that OLS might introduce relative to the 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 now delve deeper into the household surveys to understand which key groups are most likely to be driving 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.¹⁷ Since cities are dropped, urban households refer to those those living in urban villages—that is, small towns in rural districts.

Figure 5 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 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, rising agricultural incomes is the most plausible explanation for the poverty reduction. I find no evidence of any effect on total expenditures for the average urban household, although this masks a shift from food to non-food spending.

Despite positive impacts for the average household, my main poverty findings could be due to people near the poverty line being lifted just above, with little effect on the extreme

¹⁷Note that 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.

poor. Figure 6 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 7 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 and each is used in the same manner as in Figure 5. The dotted line indicates the point estimate for the average household (Table 3, Column 5).¹⁸ Panel A of Figure 7 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 assisting 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%. 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—roughly four days more consumption above the poverty line.

5 Mechanisms

This section explores three main explanations for the main findings. First, 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. However, in a setting of relatively abundant labor and reliance on land as a factor of production, increasing

¹⁸This approach is analogous 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.

farmland alone (cf., raising productivity) could increase agricultural output and reduce poverty. Hence, I first clarify the sources of the direct income gains by exploring 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 regions.

Indirect effects could see the gains increase over time. I explore three channels through which agricultural surpluses could reinforce poverty reduction. First, households could invest in productive assets and human capital. Second, revenue-flush local governments could do the same, increasing regional disparities. Third, export orientation and immediate processing requirements mean that local infrastructure development may be a necessary condition to expand production and that I may be capturing returns to that effect. Finally, I check for migration responses.

5.1 Direct effects—frontier expansion and productivity

Indonesia’s oil palm expansion is characterized by land use flux. My first mechanism analysis thus seeks to clarify whether the main results can be explained by expansion of the agricultural frontier, rather than a more efficient use of agricultural resources and rising returns to labor and land. Table 4 reports estimated impacts on district poverty incorporating information on the area under cultivation for all types of agriculture. Total district farmland is calculated as the sum of village farmland reported in the 2003 and 2008 villages censuses (PODES).¹⁹ For comparison, Columns 1 and 2 report the main OLS and IV results from Table 3 estimated for 2003—2008 (cf., 2000—2015). Both marginal effects are larger than the main results in Table 3, perhaps reflecting the high palm oil price during the 2008 food price crisis. To clarify the poverty elasticity of crop-switching relative to these aggregate effects including frontier expansion, Columns 3 and 4 denominate palm oil acreage

¹⁹Concordance 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.

with total district farmland. The OLS point estimate is not statistically different from zero and the IV estimate is indiscernible from that using total district area in Column 2. This implies that the main results are not only due to new farmland and expansion of the frontier, but rather its particular use—in other words, changes in crop mix and rising productivity within the agricultural sector. Columns 5–8 probe this conjecture further 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 that in Column 1, although the confidence intervals still overlap. The final two columns use level explanatory variables to run a “horse race” between an additional hectare of oil palm versus any farmland. Palm wins by a factor of eight.

Table 5 examines labor productivity and wages to see whether labor, land, or both are 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.²⁰ Rural services, by comparison, are often unskilled, unproductive, and informal. Manufacturing labor is typically skilled, mobile, and quantitatively smaller, with wages more likely to equalize across regions.

²⁰Estimating the main results using smallholder cultivation instead of total cultivation also gives considerably larger cultivation-poverty elasticities, as might be expected given the larger number of agricultural jobs per hectare (Tables A8 and A9).

5.2 Indirect effects—savings, investments, and public goods

This section explores three indirect mechanisms which could reinforce the poverty reduction. I focus on (a) household savings and investment, (b) government revenue and expenditures, and (c) rural public goods, and begin by examining what households are doing with their rising incomes. The first three panels of Figure 8 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 of Figure 8 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.²¹

Panel A of Table 6 reports effects on local government revenue and expenditure. Columns 1 and 2 find that total district government revenue and spending are almost twenty percent higher in the median expansion district. Columns 3 and 4 turn to villages, the key organizing unit in the Indonesian countryside. The median district agricultural expansion has allowed the average village in that district to generate 35% more own source revenue and increase expenditure by 25%.²²

Against a background of rising household incomes, increased fiscal capacity, and the need for new supply chain infrastructure, Panel B of Table 6 presents estimates on a broad range of public goods.²³ Column 5 of Table 6 finds an economically large improvement in village access to clean cooking fuel—that is, using gas or kerosene provided through utilities

²¹Home 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 is slightly beyond the scope of this paper.

²²Estimates disaggregating revenue and expenditure by type and using transfers as placebos are in Tables A11 and A12. Own source revenue is the smallest revenue stream for villages. Most comes in a grant from the central government known as the Dana Desa. Districts provide additional transfers, often in-kind in the form of health clinics, schools, and other infrastructure.

²³I pick up only average effects across all villages in treated districts. Unless districts increase public good provision across their entire jurisdiction, impacts will be concentrated around factories; in companion work, I find this is the case (Edwards, 2018).

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. The point estimate for palm land in Column 6 is small, but precisely estimated.²⁴ 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 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 A14). 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 area under cultivation for palm 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.

Can the fiscal windfalls for local governments explain the improved public good provision? Panel C of Table 6 presents the same estimates as in Panel B, controlling for log total district government revenue and expenditure and log village expenditure. With the exception of roads, results are quantitatively similar to those in Panel B. At least some of the new public goods in palm producing regions thus appear to be privately provided, possibly through in-kind transfers or new infrastructure to process, transport, and export palm oil.

²⁴Road upgrades mostly use hardened gravel rather than higher-quality asphalt. In fact, villages in expansion districts are less likely to have asphalt roads. That rural villages are not undertaking the “last mile” of road development is unsurprising. Villages in palm oil producing regions are usually extremely remote, with poor quality roads, little capital machinery apart from that around the factory, and intense rainy seasons. Firms focus investments around their estates and factories, not throughout their broader “supply shed” or district. Estimates on all categories of road quality are in Table A13.

5.3 Migration

Three types of population changes could contaminate my findings: (a) differential population growth (i.e., altering compositions); (b) inward migration of non-poor people from non-producing districts (cf., wealthy beneficiaries of natural resource sectors residing in cities and non-poor preferring to migrate to cities); and (c) outward migration of poor people. I explore each in turn.

Table 7 presents estimates of the impacts of palm expansion on different population outcomes. The explanatory variable is the share of district area under cultivation for oil palm in 2000 and 2010. All else is the same. Column 1 estimates impacts on total district population. I cannot reject the null hypothesis of no effect on population. Estimating the main results weighted by population, so small districts where a given level of migration poses more of a threat are weighted less, are also similar.²⁵ Columns 2 and 3 examine inter-district and inter-province migration rates. Expansion districts have slightly less inward migration.

That migration to expansion districts is less common than elsewhere is reassuring, but these results do not tell us whether low-income people are leaving. Figure A1 shows the probability of migrating by level of education, since population censuses do not have data on income. Mobility increases with education and cross-district migration is twice as common as cross-province migration at all education levels. These patterns are similar in high and low suitability districts (see Figure A2). 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 results.²⁶

²⁵See Tables A15 and A16. Population weighting assigns less weight to expansion regions, so this is not my preferred approach.

²⁶Province-level estimates, which remove the influence of any cross-district migration within provinces, are also similar (Table A17).

6 Environment-poverty trade-offs

An extensive literature highlights environmental costs often associated with poverty alleviation and the public debate surrounding palm oil focuses almost exclusively on environmental concerns. To calculate the local environmental trade-offs arising from an oil palm-driven change in poverty, household expenditures, or any other outcome, I relate district-level environmental degradation to changes in cultivation area since 2000. Since forest loss data cover the whole period since 2000 and fire data are highly seasonal (mostly due to El Niño), I opt for the cross-sectional long-difference analogue of the main panel specification (equation 1). Specifically, I estimate:

$$y_d = \beta(P_{d,2015} - P_{d,2000}) + \gamma X_{d,2000} + \varepsilon_d \quad (2)$$

y_d is either forest loss as a share of total district area or thermal hotspot detections since 2000, since fire is the main way farmers clear land. $P_{d,2015} - P_{d,2000}$ is the change in the share of district area under cultivation, instrumented with suitability. $X_{d,2000}$ includes the same initial conditions controls.

Environmental impacts are presented in Table 8. Columns 1–3 present OLS, IV, and reduced form estimates for district forest loss from palm oil expansion since 2000. The OLS and IV estimates suggest that 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. These large and precisely estimated effects confirm that agricultural growth, forest loss, and fire have—at least over the last fifteen years—gone hand-in-hand in the Indonesian countryside.

What is the poverty-environment trade-off that districts face when considering a development strategy based principally around agricultural growth, at least if pursued in the same manner as the last 15 years? My 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.

7 Conclusion

This paper measured the impacts of Indonesia's rapid increase in palm oil cultivation from 2000 to 2015 on regional poverty. 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 recent work emphasizing the benefits of trade and the importance of agriculture 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. An increasingly outward, market-oriented agricultural sector also appears to be reinforcing these gains through improved local fiscal capacity, infrastructure, and other public goods.

As several major economies turn inwards and invoke trade policies discriminating against particular products from developing countries, this study highlights the potential benefits of export growth and integration into global value chains for producing regions. 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 in producing regions.

References

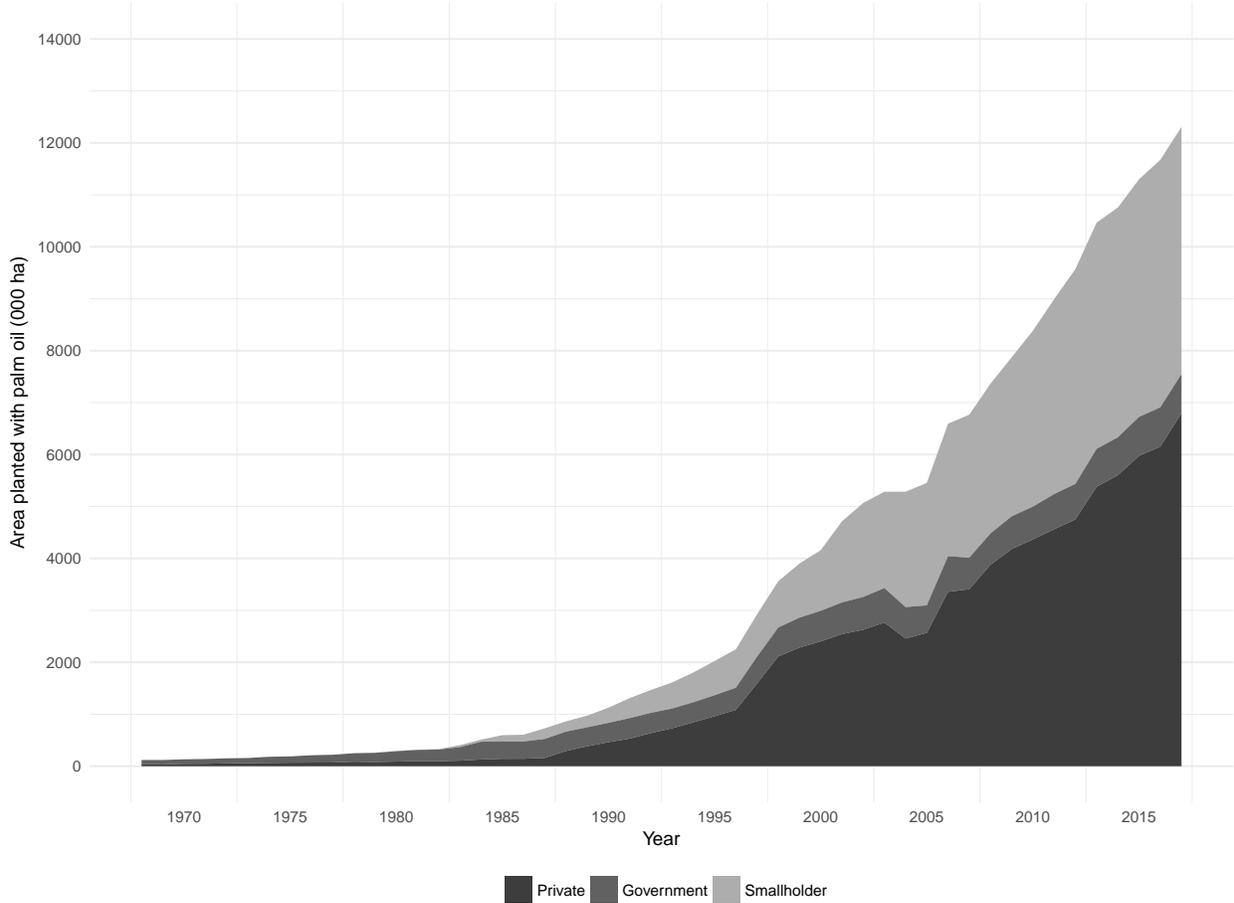
- Badan Pusat Statistik (2013), “Sensus Pertanian 2013”, available online: st2013.bps.go.id/
- Bazzi, S., Gaduh, A., Rothenberg, A., and Wong, M. (2016), “Skill Transferability, Migration, and Development: Evidence from Population Resettlement in Indonesia,” *American Economic Review*, September 2016.
- 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.
- 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.
- 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.
- 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.
- 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.
- 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.
- 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.
- Easterly, W. (2007), “Inequality does cause underdevelopment: Insights from a new instrument,” *Journal of Development Economics*, 85, pp. 755–776.
- Edwards, R. B. (2018), “Agricultural processing spillovers,” working paper.
- Engerman, S., and Sokoloff, K. (2002) “Factor endowments, inequality, and paths of development among new world economies,” NBER Working Paper Series 9259.
- 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.
- 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.
- 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.
- 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.

- Hayami, T. (2010), "Plantations Agriculture," in *Handbook of Agricultural Economics*, pp. 3305–3322.
- 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. (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.
- 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.
- Mokyr, J. (1976) "Industrialization 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 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.
- 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. (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://www.pecad.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/>

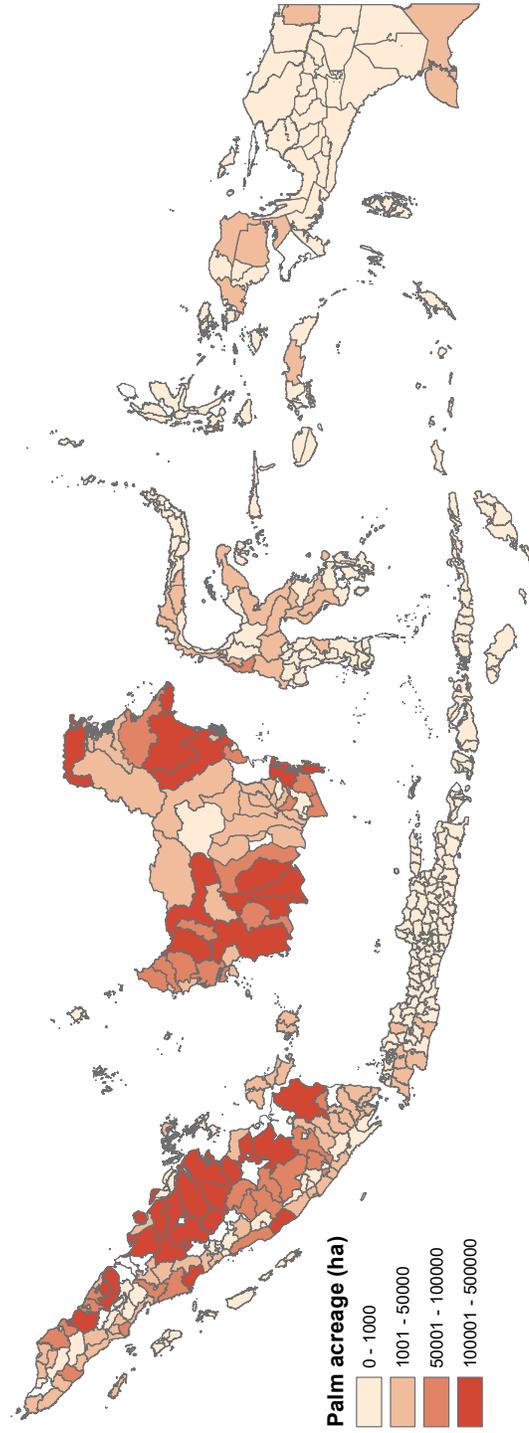
Figures and Tables

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.

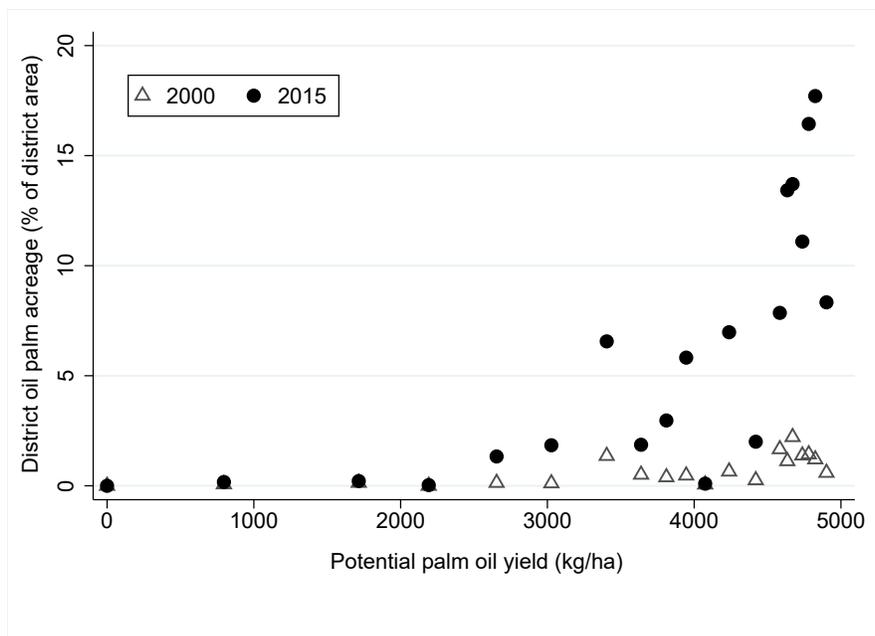
FIGURE 2: DISTRICT PALM OIL ACREAGE, 2015



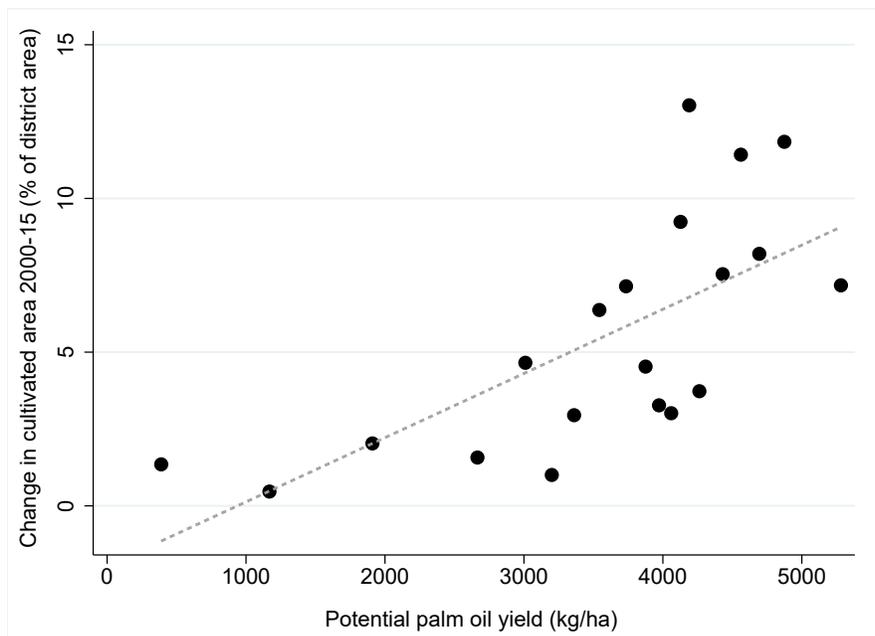
Notes: District palm oil acreage in 2015 is taken from the Tree Crop Estate Statistics of Indonesia, 2014–2016 produced by the Directorate General of State Crops (Statistik Perkebunan Indonesia, 2014–2016 Kelapa Sawit). 2010 district definitions are used.

FIGURE 3: FIRST STAGE

(A) DISTRICT CULTIVATED AREA IN 2000 AND 2015

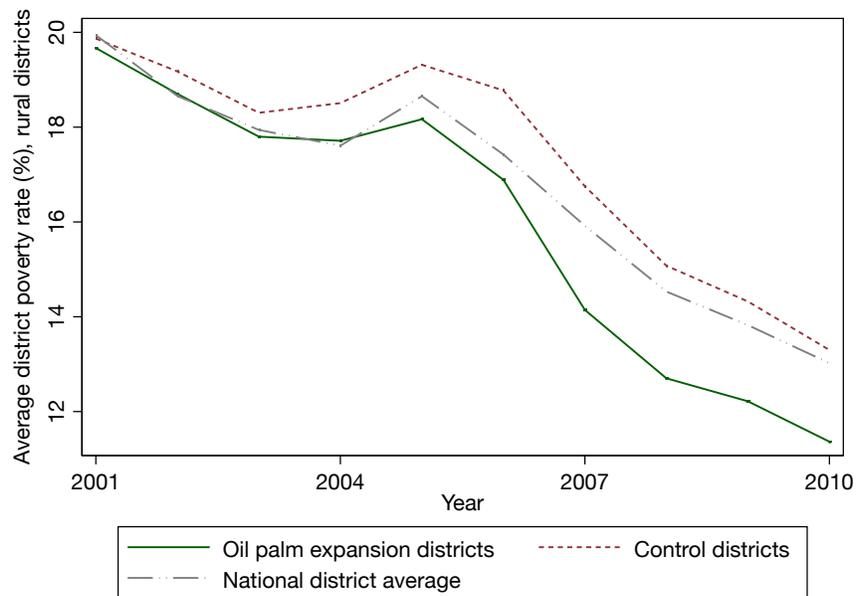


(B) POTENTIAL YIELDS AND AREA EXPANSION



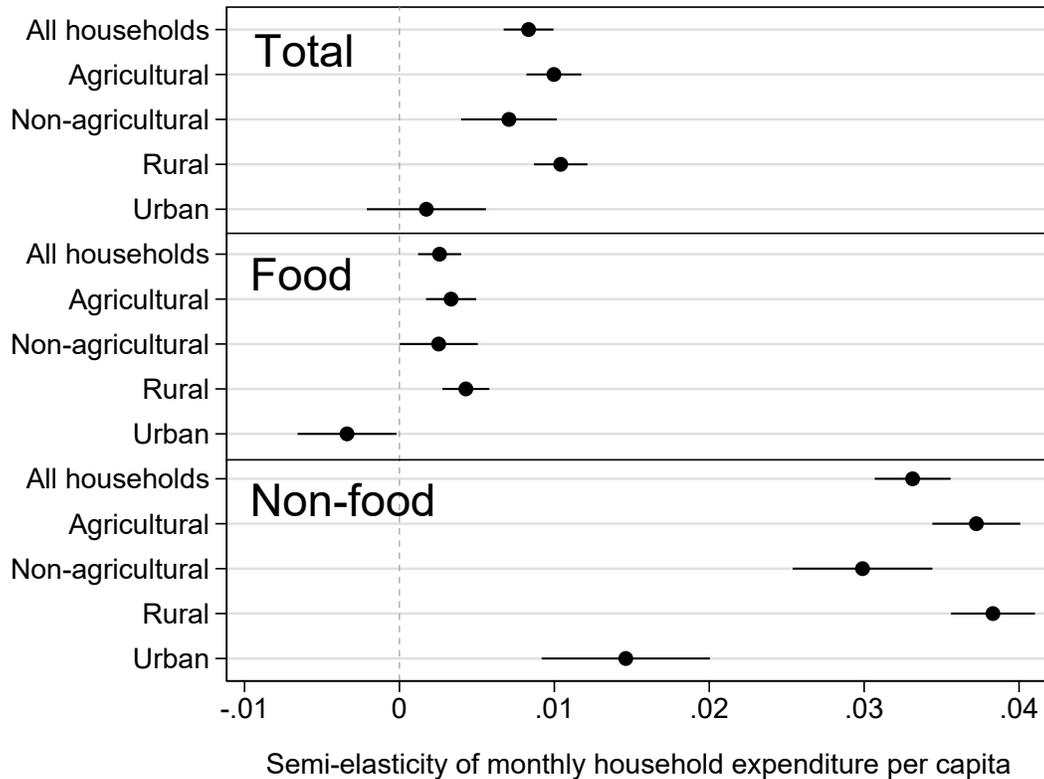
Notes: Panel A presents a binned scatter plot of district potential palm oil yield against the share of each district under cultivation for oil palm, split by year, to illustrate the increasing salience of the instrument after the demand shock. Panel B uses the change from 2000 to 2015 on the Y axis 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: EXPANSION DISTRICTS REDUCED POVERTY FASTER



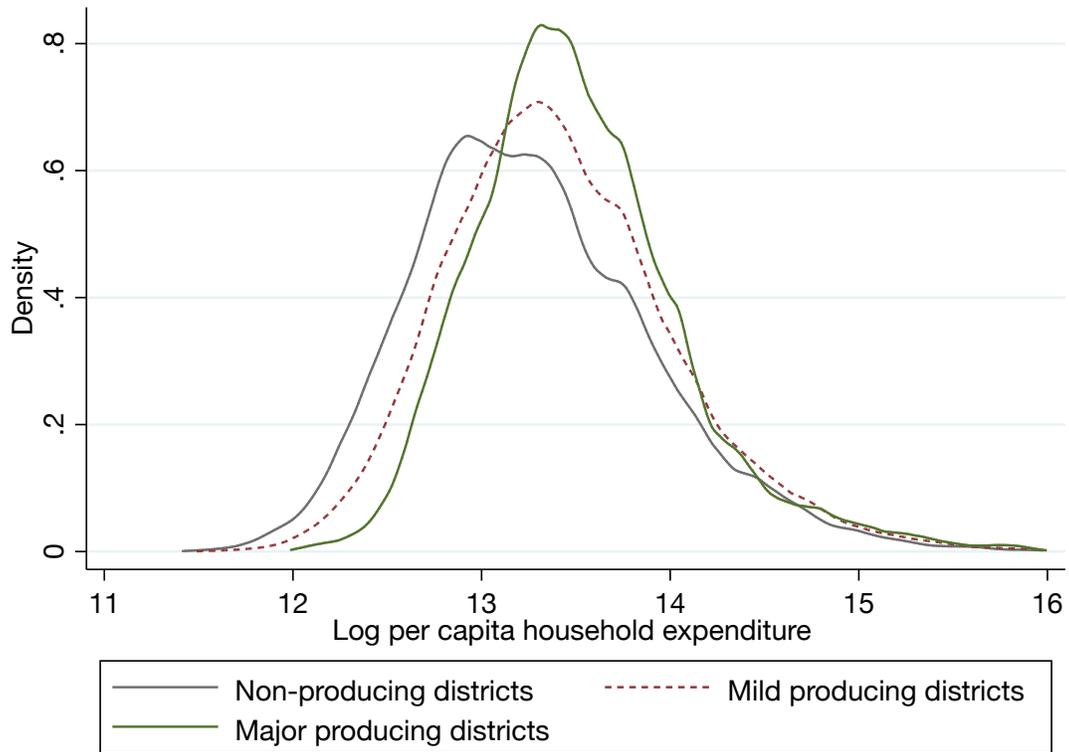
Notes: This figure is constructed using the World Bank’s DAPOER database, available through its databank. All cities (*kotas*) and rural districts outside major palm oil cultivating regions are excluded. The national district average is for all districts nationwide, including cities and regions not producing much palm oil. 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 5: CONSUMPTION IMPACTS, BY TYPE AND SECTOR



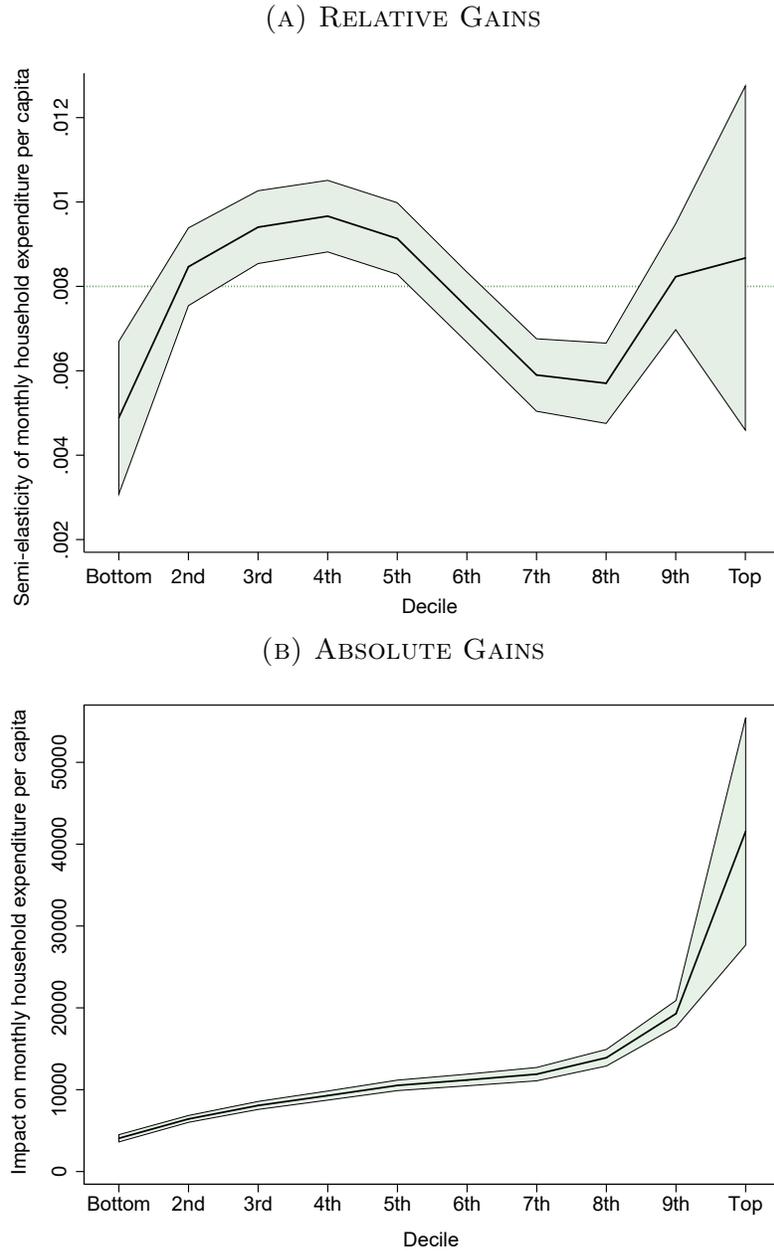
Notes: This graph plots the estimated coefficients on oil palm land from my primary IV estimator using log per capita monthly household expenditure as a dependent variable for the full sample of SUSENAS households (“All households”) and for sub-groups listed on the Y axis. Black lines indicate 95% confidence 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, and sectoral employment shares 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 6: CONSUMPTION DISTRIBUTION, 2015



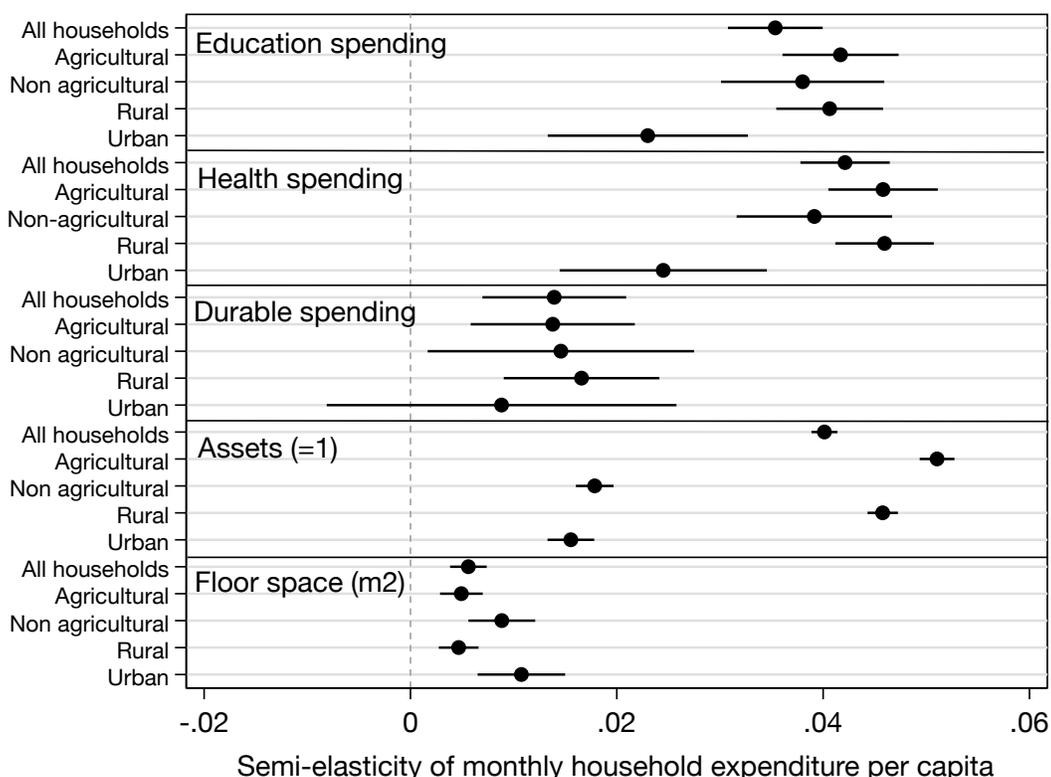
Notes: This graph plots kernel density estimates of log per capita household consumption in 2015 for households in rural districts not on Java that do not produce palm oil (gray solid), those that produce only a little (red dash), and those that are major producers (green solid), defined as over 20%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

FIGURE 7: CONSUMPTION IMPACTS BY DECILE



Notes: These graphs plot the estimated coefficients on oil palm land from my primary IV estimator using \ln 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. The dotted line in Panel A indicates the average effect for everyone. 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, and sectoral employment shares 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 8: 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, and sectoral employment shares 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: 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. Observations are districts in 2000 or the nearest feasible period. Variable construction and data sources are detailed in Appendix One.

TABLE 2: FIRST-STAGE—SUITABILITY AND AREA EXPANSION

Dependent variable	Share of district area under cultivation for oil palm (%)				
	1	2	3	4	5
Column					
Post*suitability (kg/ha)	0.0021*** (0.0004)	0.0022*** (0.0005)	0.0019*** (0.0005)	0.0018*** (0.0005)	0.0023*** (0.0006)
District and year FEs	✓	✓	✓	✓	✓
Baseline trends	✓	✓	✓	✓	✓
Cash crops trends		✓	✓	✓	
Additional trends			✓	✓	
Poverty pre-trend				✓	
Lat-long polynomial trends					✓
Observations	334	334	326	282	334

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, and sectoral employment shares with a post period indicator variable. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial interacts each district's latitude and longitude, taken at its centroid, and the squared term of each, with the post period. Column 5 controls for log-changes in district poverty, and pre-period poverty is calculated from 1993–2002, when SUSENAS became district-representative and the Asian Financial Crisis had subsided. The sample is smaller because I reset the district definitions to 1993 boundaries to ensure comparability. Robust standard errors are in parentheses and clustered at the district level.

TABLE 3: MAIN RESULTS—REGIONAL POVERTY AND HOUSEHOLD CONSUMPTION, 2000–2015

Dependent variable	District poverty rate (%)			Log expenditure (IDR)		
	OLS	IV	Reduced form	OLS	IV	Reduced form
Estimator	1	2	3	4	5	6
Column						
Oil palm land/district area (%)	-0.081** (0.040)	-0.536*** (0.160)		0.001*** (0.000)	0.008*** (0.001)	
Post*suitability ('000 kg/ha)			-1.120*** (0.274)			0.018*** (0.002)
Excluded F statistic		34			19,908	
Observations	340	334	334	241,349	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, and sectoral employment shares 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 4: EXPANSION ONTO MARGINAL LANDS, 2000–2008

Dependent variable	District poverty rate (%)						
	OLS	IV	OLS	IV	OLS	OLS	OLS
Estimator	1	2	3	4	5	6	7
Column							
Oil palm land/district area (%)	-0.108** (0.051)	-0.896*** (0.277)					
Oil palm area / farmland (%)			-0.011 (0.053)	-0.811*** (0.253)			
Farmland / district area (%)					-0.036* (0.020)		
Oil palm area (000 ha)						-0.009** (0.004)	
Farmland (000 ha)							-0.001*** (0.000)
Excluded F statistic		26		26			
Observations	340	334	340	334	340	340	340

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 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, and sectoral employment shares are included throughout. Robust standard errors are in parentheses, clustered at the district level.

TABLE 5: LABOR PRODUCTIVITY AND WAGES—IV RESULTS

Dependent variable	Output per worker			Log wages		
	Agriculture	Manufacturing	All	Agriculture	Manufacturing	Services
Sector	1	2	3	4	5	6
Column						
Oil palm land/district area (%)	1.600*** (0.578)	6.846*** (1.564)	0.039*** (0.013)	0.076** (0.037)	0.026 (0.023)	-0.006 (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. 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. Robust standard errors are in parentheses, clustered at the district level.

TABLE 6: FISCAL LINKAGES AND PUBLIC GOODS—IV RESULTS

<i>Panel A: Fiscal outcomes (in logs)</i>				
Level of government	District		Village	
	Revenue	Expenditure	Own source revenue	
Dependent variable	1	2	3	
Column	4			
Oil palm land/district area (%)	0.039*** (0.015)	0.043*** (0.014)	0.074*** (0.006)	0.054*** (0.003)
Excluded F statistic	29	28	5,104	7,141
Observations	266	264	44,699	70,977
<i>Panel B: Village public good provision (=1)</i>				
Dependent variable	Clean cooking fuel	Improved road	Street light	Market
	5	6	7	8
Oil palm land/district area (%)	0.029*** (0.001)	0.002** (0.001)	0.013*** (0.001)	0.004*** (0.001)
Excluded F statistic	8,447	8,447	8,447	8,447
Observations	82,349	82,349	82,349	82,349
<i>Panel C: Village public good provision, conditional on fiscal indicators</i>				
Dependent variable	Clean cooking fuel	Improved road	Street light	Market
	9	10	11	12
Oil palm land/district area (%)	0.035*** (0.001)	0.001 (0.001)	0.013*** (0.001)	0.004*** (0.001)
Excluded F statistic	8,531	8,531	8,531	8,531
Observations	67,444	67,444	67,444	67,444

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 or village level. 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. Village estimates include additional village level urban, coast, hilly terrain, and primary sector of income dummies. Panel C includes village expenditure, district expenditure, and district revenue as controls. Robust standard errors are in parentheses, clustered at the district level.

TABLE 7: POPULATION AND MIGRATION IMPACTS—IV RESULTS

Dependent variable	Lived in a different district..				
	Log (population)	5 years ago (=1)	at birth (=1)	5 years ago (=1)	Lived in a different province.. at birth (=1)
Column	1	2	3	4	5
Oil palm land/district area (%)	-0.007 (0.009)	-0.006*** (0.000)	-0.016*** (0.000)	-0.001*** (0.000)	-0.003*** (0.000)
Excluded F statistic	28	569,058	570,362	569,678	569,678
Observations	270	9,523,094	9,533,206	9,535,015	9,535,015

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. 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. Robust standard errors are in parentheses, clustered at the district level.

TABLE 8: ENVIRONMENTAL IMPACTS, 2000–2016

Dependent variable	Forest loss			Hotspot detections		
	OLS	IV	RF	Poisson	IV Poisson	RF Poisson
Estimator	1	2	3	4	5	6
Column						
Δ Oil palm land/district area (%), 2000–15	0.008*** (0.002)	0.017*** (0.002)		0.028** (0.011)	0.078*** (0.016)	
Suitability ('000 kg/ha)			0.038*** (0.005)			0.461*** (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. Hotspots are detections per district since 2000. 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, and sectoral employment shares. Robust standard errors are in parentheses, clustered at the district level.

Export agriculture and regional development: evidence from Indonesia

Supplementary appendix—not for publication

This supplementary appendix comprises three sections. Section 1 is the data appendix, explaining variables and the construction of the main district panel. Section 2 provides robustness checks and falsification exercises related the validity of my instrumental variable. Section 3 includes additional robustness checks and remaining tables and figures referenced in the main paper.

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 km² 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 (see Table 4 in the main paper).

- **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. 2013 is used for no reason other than data availability.
- **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.
- **Lifetime migration status:** a dichotomous indicator equal to one if an individual reports being born in a different district/province 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 Identification checks

2.1 Additional covariate sensitivity analyses

Tables A1 and A2 include additional trends intended to capture potential violations to the exclusion restriction. Results are generally stable across specifications.

2.2 Pre-trends falsification tests

Tables A3 and A4 regress pre-period poverty and household consumption trends on subsequent palm oil expansion and suitability, i.e., estimating the main results with the time period for my dependent variable deeply lagged to before the agricultural expansion. Table A3 uses the same 2000 district definitions used throughout the paper. Table A4 resets district definitions to boundaries in 1993, which is the pre-period year used. 1993 was the first year SUSENAS became fully representative at the district level and 2002 skips the spike in poverty associated with the Asian Financial Crisis, which differentially affected poverty across districts.

2.3 Reduced form falsification tests

Columns 1–3 of Table A5 regress suitability interacted with a post indicator on poverty, first with only palm, then with additional cash crops, then with cash and food crops. Across all crops, a strong correlation between suitability and poverty only appears for palm oil in the most saturated model (Column 3). Columns 4 and 5 do the same splitting the sample of all Indonesian districts, into those with a positive change in cultivated area from 2000–15 (Column 5) and those with no change (Column 4). A strong reduced form relationship between suitability and poverty reduction appears only in palm cultivating districts, suggesting this is the main channel at work and the exclusion restriction is likely satisfied. Columns 6 and 7 look at the reduced form relationship in areas that do not cultivate much palm—specifically, Java (Column 6) and Java plus eastern Indonesia (Columns 7). There is no statistically significant relationship between palm suitability and poverty in non-producing regions.

3 Further results and robustness checks

Tables A6 and A7 show that the main results are not particularly sensitive to my sample restrictions—that is, removing cities and Java. The first three columns present OLS, IV, and reduced form estimates for the entire sample of districts, including urban cities and Java. The middle three columns exclude only Java and include cities. The final three columns exclude cities and include districts on Java.

Table A8 and A9 use smallholder oil palm cultivated area, which tends to be more labor intensive, as the explanatory variable (c.f., total palm area). Estimates using only smallholder farms are larger. Importantly, the patterns from including additional covariates are similar to those in Tables A1 and A2.

Table A10 logs the poverty dependent variable, so that districts with higher initial poverty rates are less influential in estimation.

Tables A11–A13 relate to the second part of my mechanism analysis focused on fiscal linkages and rural public goods. Table A11 disaggregates the impacts on total fiscal revenue at the district level by revenue source. Table A12 does the same for expenditure. Table A13 and Table A14 present impacts on village road quality by type of road and household access to electricity.

Figure A1 shows the probability of migrating by education level. Figure A2 shows the probability of migrating by education level after dividing districts by those with above-median suitability and below-median suitability. Tables A15–A17 are additional estimates related to migration and the reallocation of people across regions. Table A15 and A16 weight the main results by population. Table A17 concludes with province level estimates, with district data from DAPOER aggregated up to the province level. Qualitatively similar to the long difference OLS estimates, intra-province migration across districts appears unlikely to be substantially affecting estimates. Figures A3 and A4 map poverty in 2015 and the gridded potential yield data.

TABLE A1: MAIN IV RESULTS, POVERTY—ADDITIONAL TRENDS

Dependent variable	District poverty rate (%)							
Column	1	2	3	4	5	6	7	8
Oil palm land/district area (%)	-0.536*** (0.160)	-0.536** (0.222)	-0.783*** (0.240)	-0.374** (0.156)	-0.600** (0.268)	9.627 (107.655)	-0.533** (0.207)	-0.346** (0.160)
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends		Y			Y	Y		
Additional trends			Y		Y	Y		
Lat-long polynomial				Y		Y		Y
Island-by-year fixed effects							Y	Y
Excluded F statistic	35	18	24	16	14	0.008	16	13
Observations	334	334	326	334	326	326	334	334

Notes: This table shows that the main findings are relatively insensitive to additional controls designed to capture any 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 are included throughout. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A2: MAIN IV RESULTS, CONSUMPTION—ADDITIONAL TRENDS

Dependent variable	Log per capita household expenditure (Indonesian rupiah)							
Column	1	2	3	4	5	6	7	8
Oil palm land (%)	0.008*** (0.001)	0.014*** (0.002)	0.014*** (0.001)	0.003*** (0.001)	0.010*** (0.002)	0.003 (0.021)	0.005*** (0.001)	0.003*** (0.001)
Household controls	Y	Y	Y	Y	Y	Y	Y	Y
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends		Y			Y	Y		
Additional trends			Y		Y	Y		
Lat-long polynomial				Y		Y		Y
Island-by-year FEs							Y	Y
Excluded F statistic	19,908	4,942	17,161	19,793	12,773	86	15,675	20,940
Observations	237,887	237,887	230,999	237,887	230,999	230,999	237,887	237,887

Notes: This table shows that the main findings are relatively insensitive to additional controls designed to capture any potential violations of the exclusion restriction. 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. Household controls cover household size, an urban/rural dummy, and primary sector of income fixed effects. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A3: PRE-PERIOD FALSIFICATION TEST, 2000 DISTRICT DEFINITIONS

Lagged outcome	Poverty rate (%)			Poverty gap index			Log expenditure		
	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.010 (0.017)	-0.101 (0.079)		-0.007 (0.006)	-0.053 (0.037)		0.001 (0.002)	0.003 (0.008)	
Post*suitability ('000 kg/ha)			-0.206 (0.158)			-0.108 (0.074)			0.006 (0.016)
Excluded F statistic	29	288	288	294	288	288	340	334	334
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, and sectoral employment shares 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 A4: PRE-PERIOD FALSIFICATION TEST, 1993 DISTRICT DEFINITIONS

Lagged outcome	Poverty rate (%)			Poverty gap index			Log expenditure		
	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Estimator									
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.006 (0.017)	-0.105 (0.086)		-0.002 (0.007)	-0.057 (0.043)		0.001 (0.002)	0.002 (0.009)	
Post* suitability ('000 kg/ha)			-0.203 (0.163)			-0.111 (0.081)			0.004 (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, and sectoral employment shares 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 A5: REDUCED FORM FALSIFICATION TESTS

Dependent variable Sample Column	District poverty rate (%)						
	1	2	3	4	5	6	7
	Main sample			Non-expanders		Expanders	Java & east
Post*palm suitability ('000 kg/ha)	0.110 (0.367)	-1.570* (0.867)	-2.975*** (1.102)	-0.137 (1.035)	-7.365*** (2.047)	2.120 (1.486)	1.018 (1.368)
Cocoa		0.008* (0.005)	-0.001 (0.006)	0.001 (0.005)	0.021* (0.011)	-0.002 (0.007)	-0.002 (0.006)
Coffee		-0.003 (0.004)	0.007 (0.006)	0.002 (0.007)	-0.015 (0.015)	0.018 (0.011)	-0.000 (0.008)
Teas		-0.009** (0.004)	-0.012** (0.005)	0.001 (0.009)	-0.009 (0.006)	0.010 (0.007)	0.009 (0.009)
Coconut			0.000 (0.001)	-0.000 (0.004)	0.003* (0.001)	-0.002 (0.003)	0.002 (0.004)
Tobacco			-0.035*** (0.013)	-0.013 (0.012)	-0.024 (0.017)	0.018 (0.015)	-0.008 (0.013)
Dryland rice			0.035*** (0.010)	0.011 (0.013)	0.017 (0.021)	-0.016 (0.022)	-0.002 (0.018)
Wetland rice			-0.018*** (0.006)	-0.014** (0.006)	0.005 (0.013)	0.044** (0.018)	-0.012** (0.006)
Casava			0.002 (0.004)	0.008 (0.008)	-0.006** (0.003)	-0.027* (0.014)	0.014 (0.009)
Observations	350	350	350	278	252	186	260

Notes: Sample is a two-period balanced panel of all rural districts excluding Java in 2002 and 2015, at 2000 district boundaries. All agroclimatically attainable yields are separately 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. No expanders is districts that did not increase their area under cultivation for oil palm from 2000 to 2015. Java+east are districts not on Sumatra, Kalimantan, or Sulawesi. Cluster-robust standard errors are in parentheses.

TABLE A6: MAIN POVERTY RESULTS—DIFFERENT SAMPLES

Dependent variable	District poverty rate (%)								
	All districts			No Java			No cities		
Sample	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Estimator	1	2	3	4	5	6	7	8	9
Column									
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*suitalibility ('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, and sectoral employment shares are included throughout. Robust standard errors are in parentheses and clustered at the district level.

TABLE A7: MAIN CONSUMPTION RESULTS—DIFFERENT SAMPLES

Dependent variable	Log per capita household expenditure (Indonesian rupiah)											
	All districts				No Java				No cities			
	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF	OLS	IV	RF
Sample	1	2	3	4	5	6	7	8	9			
Oil palm land (%)	0.001*** (0.000)	0.006*** (0.001)		0.001** (0.000)	0.007*** (0.001)		0.002*** (0.000)		0.006*** (0.001)			
Post*suitability			0.010*** (0.001)			0.016*** (0.002)					0.010*** (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	380,254	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, and sectoral employment shares, 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 A8: SMALLHOLDER EXPANSION—POVERTY

Dependent variable	District poverty rate (%)							
	1	2	3	4	5	6	7	8
Smallholder palm land (%)	-1.390*** (0.449)	-1.599** (0.721)	-1.796*** (0.621)	-0.999** (0.424)	-1.985* (1.064)	1.247 (1.362)	-1.266** (0.493)	-0.849** (0.385)
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends		Y			Y	Y		
Additional trends			Y		Y	Y		
Lat-long polynomial				Y		Y		Y
Island-by-year fixed effects							Y	Y
Excluded-F statistic	24	11	15	15	6	2.306	16	14
Observations	334	334	326	334	326	326	334	334

Notes: This table replicates Table A1 using smallholder acreage instead of total acreage. 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, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A9: SMALLHOLDER EXPANSION—CONSUMPTION

Dep. variable	Log per capita household expenditure (Indonesia rupiah)								
Estimator	OLS	IV	IV	IV	IV	IV	IV	IV	IV
Column	1	2	3	4	5	6	7	8	9
Smallholder palm	0.003*** (0.001)	0.021*** (0.002)	0.039*** (0.005)	0.029*** (0.003)	0.008*** (0.003)	0.031*** (0.005)	-0.002 (0.016)	0.011*** (0.002)	0.007*** (0.002)
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends			Y			Y	Y		
Additional trends				Y		Y	Y		
Lat-long poly.					Y		Y		Y
Island-year FEs								Y	Y
Excluded F		21,965	4,853	15,617	17,253	4,819	620	21,576	23,192
Observations	241,349	237,887	237,887	230,999	237,887	230,999	230,999	237,887	237,887

Notes: This table replicates Table A2 using smallholder acreage instead of total acreage. 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. Household controls cover household size, an urban/rural dummy, and primary sector of income fixed effects. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A10: MAIN RESULTS USING LOG POVERTY AS DEPENDENT VARIABLE

Dependent variable	Log district poverty rate (%)								
	OLS	IV	IV	IV	IV	IV	IV	IV	IV
Estimator	1	2	3	4	5	6	7	8	9
Oil palm land (%)	-0.009** (0.004)	-0.041*** (0.010)	-0.050*** (0.017)	-0.057*** (0.015)	-0.021** (0.010)	-0.056*** (0.020)	0.471 (5.360)	-0.029** (0.012)	-0.015 (0.010)
Baseline trends		Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends			Y			Y	Y		
Additional trends				Y		Y	Y		
Lat-long polynomial					Y		Y		Y
Island-by-year fixed effects								Y	Y
Excluded F statistic		35	18	24	16	15	0	17	13
Observations	340	334	334	326	334	326	326	334	334

Notes: This table replicates Table A1 with the dependent variable logged. 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, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A11: 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** (0.077)	0.029** (0.014)	-0.001 (0.064)	0.005 (0.022)	0.059 (0.073)	0.039*** (0.015)	0.055*** (0.020)
Excluded F statistic	22.484	28.503	27.227	28.503	21.788	28.503	28.503
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 A12: 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 (0.019)	0.067*** (0.025)	0.140** (0.070)	0.046 (0.029)	0.131** (0.058)	0.037* (0.021)	0.096 (0.079)	0.074*** (0.028)
Excluded F statistic	26.882	28.655	28.655	28.655	27.693	28.655	16.663	28.655
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** (0.047)	0.059 (0.039)	0.122** (0.053)	0.043** (0.021)	-0.033 (0.025)	0.031** (0.013)	0.043*** (0.014)	177155.763** (71338.617)
Excluded F statistic	23.698	27.511	27.225	28.470	28.470	28.470	28.470	27.492
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 A13: IMPACTS ON VILLAGE ROADS

Road outcome	Gravel (=1)		Asphalt (=1)		Gravel or Asphalt (=1)				
	OLS	IV	OLS	IV	OLS	IV			
Estimator	OLS	RF	OLS	RF	OLS	RF			
Column	1	2	3	4	5	6	7	8	9
Oil palm land (%)	0.003*** (0.000)	0.007*** (0.001)	-0.000 (0.000)	-0.004*** (0.001)	0.003*** (0.000)	0.002*** (0.001)			
Post * suitability		0.014*** (0.003)				-0.009*** (0.003)			0.005*** (0.002)
Excluded F statistic		8447.504		8447.504		8447.504		8447.504	
Observations	83551	82349	82349	83551	82349	82349	83551	82349	82349

Notes: Sample is villages observed in 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 village level. IV estimates instrument district oil palm land with district potential palm oil yield interacted with a post period indicator. All estimates include district and year fixed effects and differential trends for initial poverty rates, rural population shares, literacy rates, and sectoral employment shares, and village estimates include additional village level urban, coast, hilly terrain, and primary sector of income dummies. Cluster-robust standard errors are in parentheses.

TABLE A14: IMPACTS ON HOUSEHOLD ACCESS TO ELECTRICITY

Road outcome	Has electricity (=1)			Has PLN electricity (=1)		
	OLS	IV	RF	OLS	IV	RF
Estimator	1	2	3	4	5	6
Oil palm land (%)	-0.0002 (0.182)	0.0086*** (0.000)		0.0006 (0.000)	-0.0095*** (0.000)	
Post * suitability			0.0181*** (0.000)			0.0201*** (0.000)
Excluded F statistic		19,909			19909	
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. Robust standard errors are in parentheses and clustered at the district level.

TABLE A15: MAIN RESULTS, POVERTY—ADDITIONAL CONTROLS, WEIGHTED BY POPULATION

Dependent variable	District poverty rate (%)							
	1	2	3	4	5	6	7	8
Oil palm land/district area(%)	-0.382** (0.162)	-0.631 (0.453)	-0.771** (0.348)	-0.223* (0.122)	-0.496 (0.305)	2.526 (13.182)	-0.339* (0.195)	-0.185* (0.110)
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends		Y			Y	Y		
Additional trends			Y		Y	Y		
Lat-long polynomial				Y		Y		Y
Island-by-year fixed effects							Y	Y
Excluded F statistic	16.277	3.549	7.558	9.072	9.021	0.045	6.218	8.436
Observations	334	334	326	334	326	326	334	334

Notes: 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, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A16: MAIN RESULTS, CONSUMPTION—ADDITIONAL CONTROLS, WEIGHTED BY POPULATION

Dependent variable	Log per capita household expenditure (Indonesian rupiah)								
	OLS	IV	IV	IV	IV	IV	IV	IV	
Estimator	1	2	3	4	5	6	7	8	
Column	1	2	3	4	5	6	7	8	
Oil palm land (%)	0.001*** (0.000)	0.009*** (0.001)	0.014*** (0.003)	0.016*** (0.002)	0.001 (0.001)	0.007*** (0.002)	-0.254** (0.121)	0.003*** (0.001)	0.001* (0.001)
Household controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Baseline trends	Y	Y	Y	Y	Y	Y	Y	Y	Y
Cash crops trends			Y			Y	Y		
Additional trends				Y		Y	Y		
Lat-long polynomial					Y		Y		Y
Island-year FEs								Y	Y
Excluded F		13541.181	1626.114	8695.811	20952.111	8548.664	6.689	12885.850	21529.738
Observations	223942	220693	220693	214600	220693	214600	214600	220693	220693

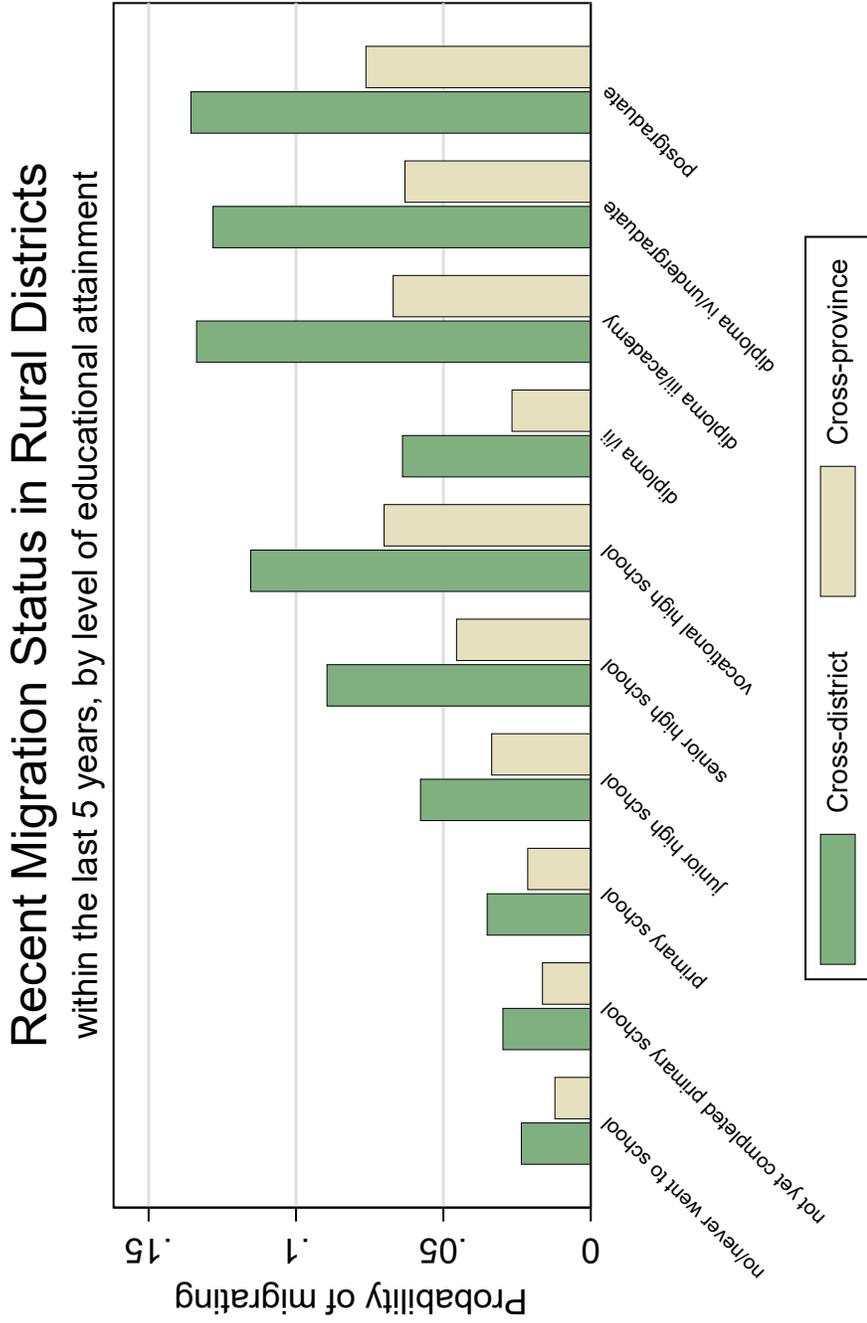
Notes: Sample is household observations in 2002 and 2015 with identifying variation in oil palm expansion and suitability measured at the district level, based on 2000 district boundaries, excluding cities and Java. District oil palm land is instrumented with district potential palm oil yield interacted with a post period indicator, and district and year fixed effects are included throughout. Household controls cover household size, an urban/rural dummy, and primary sector of income fixed effects. Baseline trends separately interact 2000 poverty, rural population shares, literacy rates, and sectoral employment shares with a post period indicator. Cash crops suitability includes similar interactions for cocoa, coffee, and teas. 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 polynomial 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 A17: PROVINCE 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). Heteroskedasticity-robust standard errors are in parentheses. This appendix table is from an earlier draft available online, where additional results can be found: <http://dx.doi.org/10.2139/ssrn.3040400>

FIGURE A17: 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 A17: MIGRATION STATUS, BY EDUCATION AND SUITABILITY

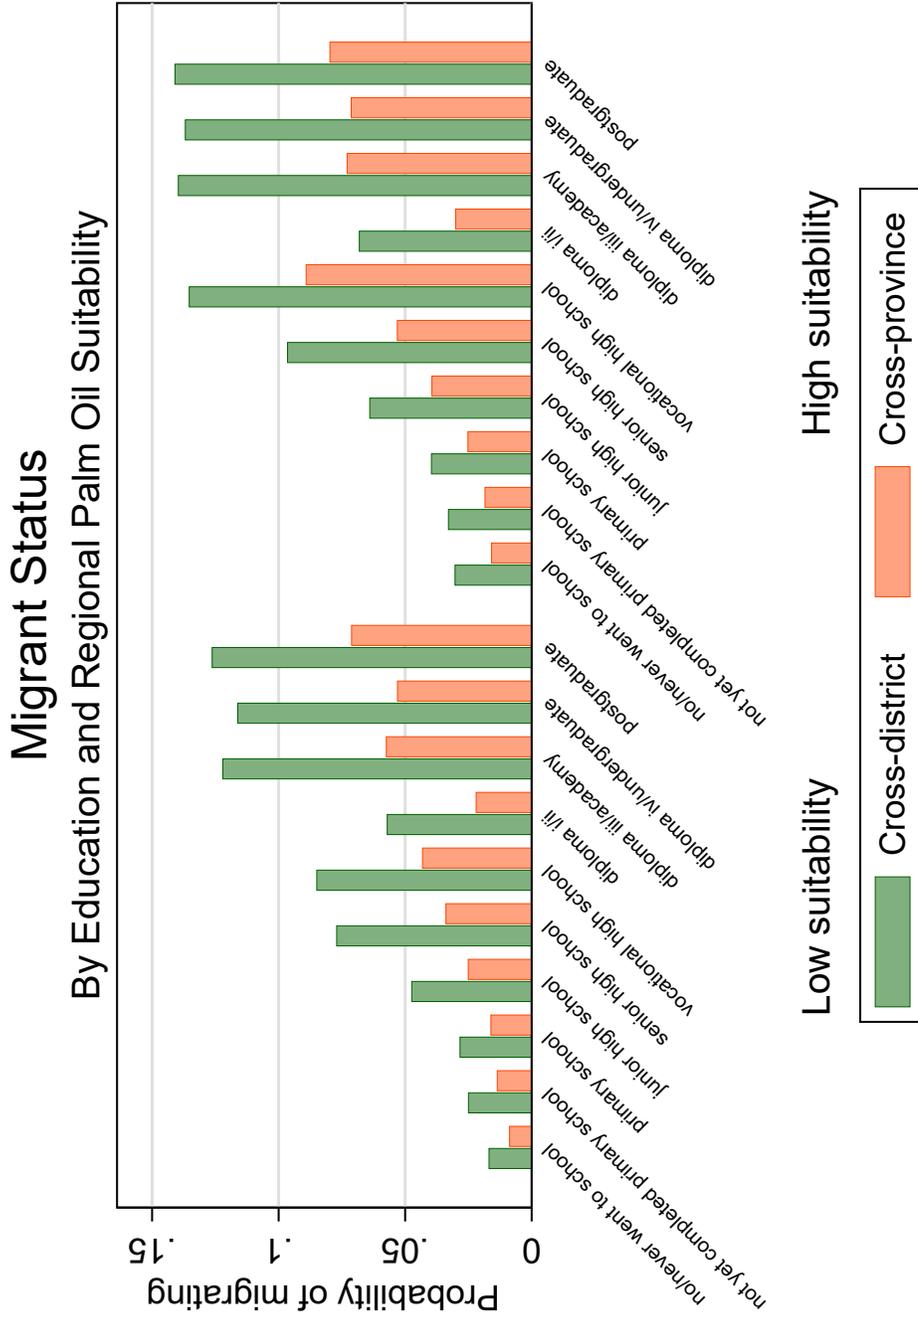
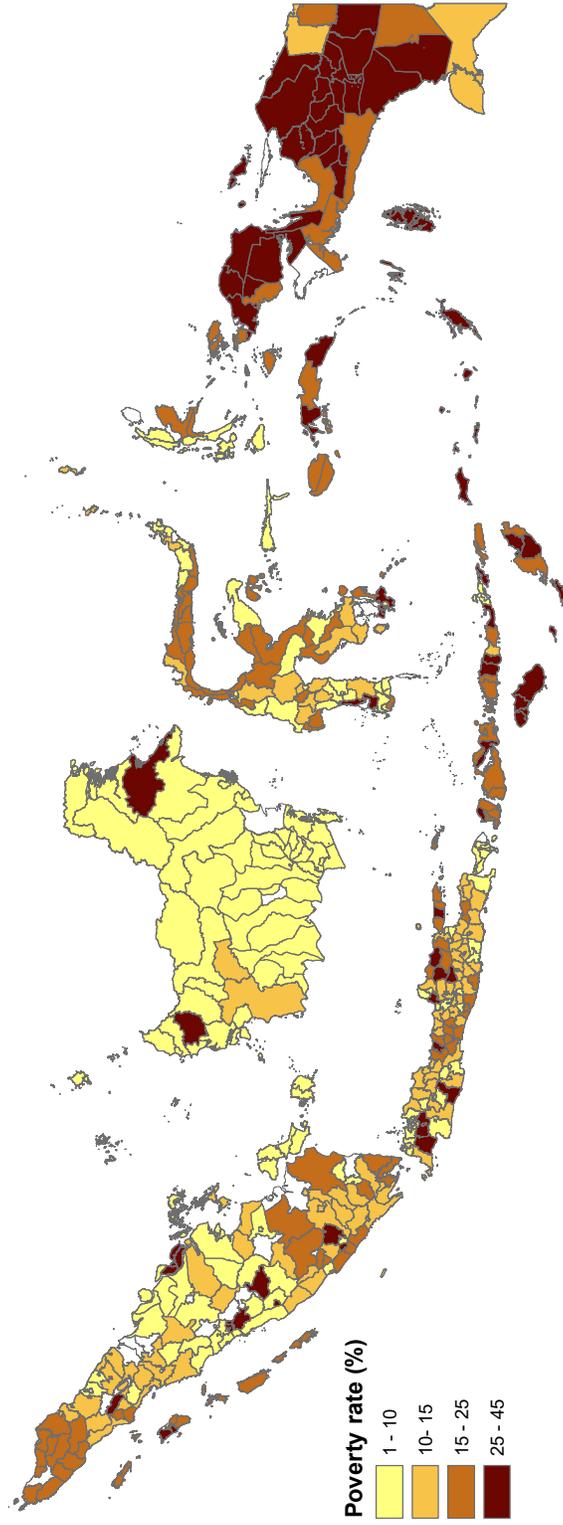
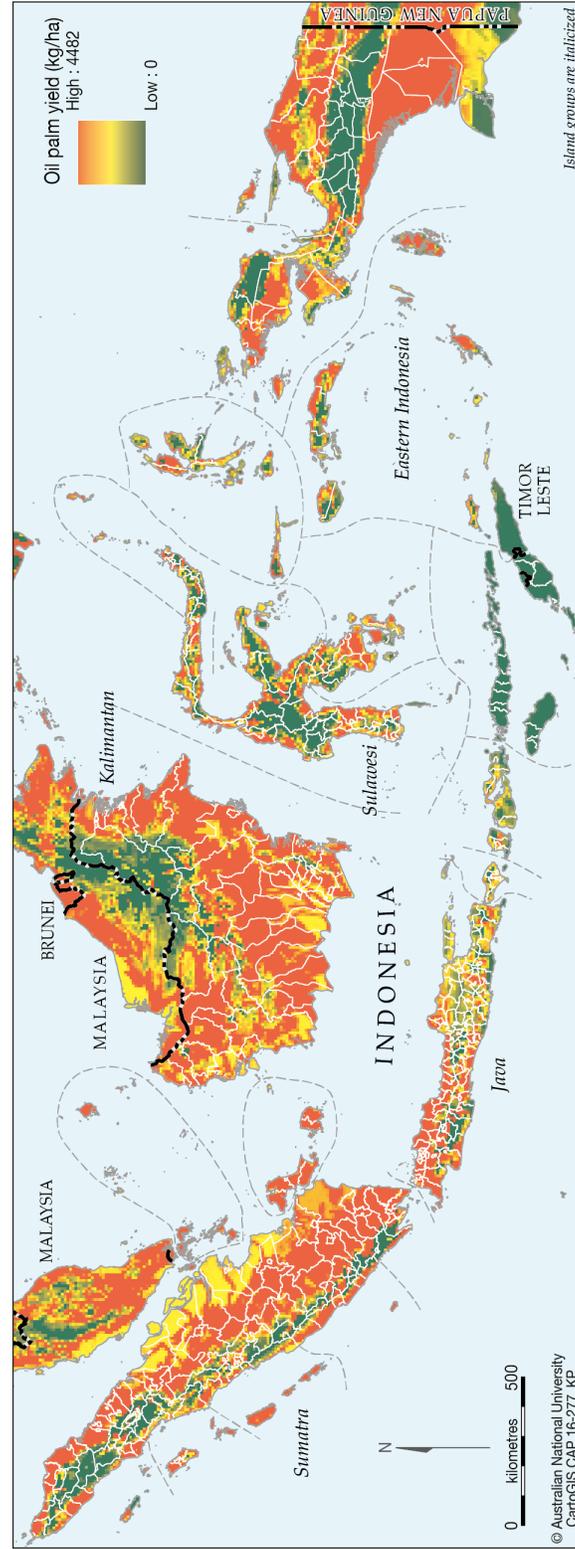


FIGURE A17: DISTRICT POVERTY RATES IN 2015



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

FIGURE A17: AGRO-CLIMATICALLY ATTAINABLE PALM OIL YIELDS, FAO-GAEZ GRIDDED DATA



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