Market access and deforestation

Ryan Abman[1] and Clark Lundberg[2]

[1]Department of Economics, San Diego State University: rabman@sdsu.edu.
[2]Economic Research Service, USDA: clark.lundberg@usda.gov.

Selected Paper prepared for presentation at the 2022 Agricultural & Applied Economics Association Annual Meeting, Anaheim, CA; July 31–August 2

This research was supported by the U.S. Department of Agriculture, Economic Research Service. The findings and conclusions in this research paper are those of the authors and should not be construed to represent any official USDA or U.S. Government determination or policy.

Market access and deforestation*

Ryan Abman† and Clark Lundberg‡

Abstract

Improving market access for rural smallholders holds tremendous potential to improve the livelihoods of local producers and promote rural growth. Despite potential benefits, improved market access may negatively impact the environment as smallholder land-use decisions become coupled with market conditions. Studying the impacts of market access on the environment is difficult, however, as changes to market access are often accompanied or driven by changing macroeconomic conditions that challenge causal identification. We study the causal impacts of market access on forest loss in Central Ghana through a program introducing village access to output markets and input credit markets in oil palm production. Market access is facilitated through production contracts with a large palm oil mill in which smallholders receive credit access to establish oil palm production, a guaranteed price for the entirety of their oil palm output, and output delivery at the village. Using a variety of difference-in-differences approaches, we find substantial increases in forest loss in targeted villages immediately following the program roll-out. Elevated forest loss generally persisted over the following decade in treated villages. This paper provides evidence that the ecological impacts of reforms to market access and contracting may be sizable.

Keywords: Deforestation; Ghana; oil palm; market access; credit constraints; contracting
JEL Classification Codes: Q13; Q23; O13

*The findings and conclusions in this article are those of the authors and should not be construed to represent any official USDA or US Government determination or policy. We are extremely grateful to Anette Ruml and Matin Qaim for sharing their survey data with us and making this project possible. We thank Jeffrey Castanon for excellent research assistance. All errors are our own.

†Department of Economics, San Diego State University. E-mail: rabman@sdsu.edu.
‡Economic Research Service, United States Department of Agriculture. E-mail: clark.lundberg@usda.gov.

1 Introduction

Rural agricultural markets often face important barriers to entry and competition that prevent smallholder farmers from raising their incomes. Credit market constraints, prohibitive transportation costs, and/or output price risk are often present in lower income countries where much of the local population is dedicated to agriculture. Improving market access for rural smallholders offers tremendous potential to improve their livelihoods and promote rural growth.

Despite the potential economic benefits, improved market access may have negative impacts on the local environment. In the short run, changes to relative value of land in agriculture may lead to agricultural extensification at the expense of natural land. In the medium to long run, general equilibrium effects of improved market access may well be ambiguous. Improved incomes may increase the demand for for land-intensive goods, but increased incomes may also lead towards preferences for environmental protection and sustainable agricultural practices. Furthermore, improved credit markets my limit the need for farmers to liquidate natural capital to purchase inputs at the start of a growing season or to cover expenditures during a negative economic shock.

In this paper, we study the causal impacts of market access on forest loss in Central Ghana through a program introducing village access to output markets and input credit markets in oil palm production. Market access is facilitated through production contracts with a large palm oil mill in which smallholders receive credit access to establish oil palm production, a guaranteed price for the entirety of their oil palm output, and output delivery at the village. This program was established in villages where few farmers were marketing their crops. We link a variety of geospatial data to local Ghanaian villages, including high resolution annual estimates of forest loss from 2001 to 2019. We compare changes in forest loss surrounding villages that participated in the contracting and credit program in 2008 to those that did not via a variety of difference-in-differences approaches. We find substantial increases in forest loss in targeted villages immediately following the program roll-out that coincide with the timing of farmer enrollment in the program. Furthermore, we find that higher average annual forest loss among participating villages continued over the following decade. These results are robust to a variety of different control group specifications, including a matched sample of control villages. Fully specified event studies provide strong evidence against spurious findings driven by pre-existing differences in forest loss trends.

This paper contributes causal evidence to the relationship between market access and the environment. Variation in market access is often driven by large, macroeconomic changes, country-level governance institutions or even a local community's ability to solve coordination problems and engage in collective action (Markelova *et al.*, 2009). All of these drivers of market access also have their own direct implications for changes in land use and forest loss that undermine causal interpretation of observed relationships between forest loss and barriers to market. Our setting allows us to study the ecological impacts of improved market access in isolation from other macroeconomic changes that would otherwise confound the analysis. In particular, the market access program had three important aspects that all have implications for land use. First, the program provided direct access to commodity markets. In essence, it reduced previously prohibitive transportation costs to nearly nothing. This aspect of the program relates to the literature on transportation costs and deforestation which has generally focused on changes in transportation costs induced by transportation infrastructure. Generally, this literature has found that blanket decreases in transportation costs for all goods lead to increased forest loss – e.g. Pfaff (1999); Pfaff *et al.* (2007)– however, new-road induced reductions in transportation costs reduce costs for both agricultural goods and timber which effectively decreases clearing costs while also increasing net returns for agricultural products.[1] Our setting allows us to examine the agricultural marketing part in isolation from blanket changes in transportation costs.

Second, this paper contributes to the literature on credit constraints and deforestation. Ex ante, it is unclear whether access to credit increases or decreases forest loss. If liquidity constraints force smallholders to extract forest resources to purchase agricultural inputs because they do not have credit, then improving credit markets may reduce deforestation (Jayachandran, 2013). However, if credit allows landowners to extensify agriculture or pasture, then they may chose to clear more land under credit and use the increased production to make payments. Indeed, Assunção *et al.* (2020) find that credit constraints for cattle ranchers in Brazil substantially reduced deforestation.

[1] Recent work like Asher *et al.* (2020) and Baehr *et al.* (2021) has demonstrated that the relationship between roads and deforestation is considerably more nuanced than some of the earlier research had concluded.

We bring additional evidence to bear on this question but in a new smallholder setting.

Finally, this paper contributes to the work on producer price uncertainty and agricultural land use decisions. Many have argued marketing contracts that enumerate a price for agricultural output in advance reduce the price risk otherwise borne by producers and improve productivity. This logic has motivated a range of different development policies including price stabilization programs, export bans, and a variety of contracting schemes. Despite their promise, reducing price risk may increase deforestation through agricultural extensification. Lundberg and Abman (2022) find decreases in maize price volatility increase subsequent maize production and deforestation. To date there is little well-identified empirical evidence on the role that contracting programs in particular may have in forest loss, a gap in the literature that the present papers helps to fill.

The remainder of the paper is as follows. We begin by providing background in oil palm farming in Ghana and details on the oil palm contract program we study. We then describe the myriad of data sources we use for the analysis and outline our empirical strategy while providing evidence for our identifying assumptions. We present our findings across a variety of specifications and conclude with some final remarks.

2 Background

Oil palms are native to Ghana with a long history of smallholder cultivation in Ghana. Smallholder farmers produce about 60% of palm oil output in Ghana and account for 85% of the planted area (Khatun *et al.*, 2020). Despite familiarity with palm oil as a commodity crop, many Ghanaian smallholders have failed to take advantage of the global boom in palm oil demand in large part due to market access frictions (Khatun *et al.*, 2020). However, improving access to palm oil markets is likely to drive local land use change at the expense of sensitive ecosystems as oil palm monoculture displaces native forest.

2.1 Oil Palm Contract Program

Our quasi-experimental market access shocks come from a palm oil contracting program introduced in 2008 by the Twifo Oil Palm Plantation (TOPP) in the Central Region of Ghana. TOPP is a large palm oil plantation with a high-capacity palm oil mill. Despite growing very large quantities of oil palm fruit on the plantation, the processing capacity of the mill is high enough that TOPP is unable to fully meet mill capacity with their own output. As a result, TOPP developed a resource-providing contract program to secure additional palm oil output from smallholder farmers in the region. The program was introduced at the village level and TOPP offers contracts to all households in program villages. The program has two unique features that introduce market access to treated villages. First, TOPP arranges transportation to pick up oil palm output in treated villages, eliminating prohibitively high marketing and transportation costs for remote rural smallholders. This effectively brings the commodity market to the smallholder and allows them to participate in palm oil commodity markets that they otherwise would not have access to. Notably, contract households are prohibited under the contract from selling palm oil output to any other buyers. Hence, the contract uniquely establishes the output market location for the entire oil palm output. Second, TOPP contracts provide access to input and credit markets. Contract households can purchase inputs on credit, including planting materials, tools, machinery, and chemical inputs. Access to credit markets allows smallholder farmers to defray the considerable startup costs associated with oil palm production—average costs of smallholder palm plantations from establishment to first harvest are approximately 1600 US dollars per hectare (Ruml and Qaim, 2020). Farmers that purchase inputs on credit are charged an annual interest rate of 11.5% with in-kind repayment through future oil palm output. Smallholders only receive payment for three quarters of their output with the remaining one quarter used to pay down any credit balance until the balance is paid in full. Credit market access is not limited to defraying startup costs; contract participants can purchase additional inputs on credit throughout the contract duration, including labor, tools, fertilizer and pesticides. This contracting arrangement is notably different than many other production contracts which specify particular production methods and input quantities—TOPP contract participants maintain agency over production decisions. Because the contract does not dictate production input decisions, we characterize the contract as providing *access* to credit and inputs with use decisions remaining in the hands of smallholder producers. Finally, like many agricultural contracts regimes, the TOPP program reduces output price risk for smallholder farmers by establishing fixed prices for palm oil output.

3 Data

We use a variety of different data sources in the analysis. For data on participating villages in the market access contract program, we use data from Ruml and Qaim (2020).[2] These data include geo-coordinates for surveyed villages and allow us to identify the locations of 13 different villages that participated in the contract program. The data also include household- and plot-level survey response data for a sample of households participating in the program. Survey villages and households were selected from randomly from participating households The survey was conducted in 2018 but asked a number of retrospective questions with regards to contracting and previous use of oil palm plots.

The household and plot level data allow us to confirm underlying assumptions regarding oil palm expansion in response to the contracting program. First, we see that contracting participation increases rapidly around the the 2008 start of the program with more than 80% of all surveyed households reporting their first oil palm contract was between 2008 and 2009 with small subsequent increases in 2010 and beyond (see Figure 1). Second, the vast majority of surveyed plots (84%) were converted to oil palm after the farmer entered into a contract. One fifth of these plots were forest land prior to conversion (though tree cover loss – as measured in this paper – may also arise from pasture or agricultural land conversion so long as standing trees exist in the area) and the plots are sizeable, approximately 6.25 acres on average.[3] Finally, smallholder survey respondents clearly faced palm oil market access frictions—only 2% of households were engaged in the commercial

[2]The data set was graciously provided by the authors.

[3]We present the kernel density estimates of oil palm plot size by previous land use. Plots of previously forested land are slightly larger than those previously in pasture or agriculture/crops.

sale of oil palm output prior to the introduction of the market access program. The important limitation of the survey data is that it does not identify non-participating households or villages in the surrounding area and, thus, limits our ability to compare plot or household level differences in outcomes in the observed, participating villages to a set of control villages.

We identify a set of other villages in the surrounding area using the Geo-Referenced Infrastructure and Demographic Data for Development (GRID3) data set on settlement extents. This data set identifies the boundary of cities (referred to as built-up areas or BUAs) and the location of small settlement areas (SSAs) via satellite imagery. We use SSAs as a proxy for villages and allow the SSAs that do not overlap with those in the program data set to serve as our control villages. Visual inspection of satellite imagery as well as the consistency with with our surveyed village-level geo-coordinates align with SSA observations suggest the GRID3 SSAs are a suitable source for non-surveyed villages in this region.

The GRID3 data set provides the locations of villages throughout the country. To attribute land area to village locations (both program and control villages), we create Thiessen polygons (a.k.a Voronoi polygons) around SSA centroids. These polygons provide a way to attribute area between village locations to one village or another based on the closest village to any point in space. These polygons are commonly used in settings in which researchers have the universe of village locations but no formal, demarcated boundaries exist (e.g. Alix-Garcia *et al.* (2013); Heß *et al.* (2021)). To ensure that no urban land is attributed to a particular village, we remove any overlapping area of our Thiessen polygons with the extent of the BUAs. This ensures that all area attributed to a village falls outside of any urban or semi-urban area.

With both the coordinates of the villages and the boundaries of the Thiessen polygons, we are able to attribute a number of different spatially-explicit characteristics for each village. We construct average elevation and ruggedness for each village at the polygon level. Elevation is the average elevation of all land within the polygon and ruggedness is the standard deviation of elevation in the same area. We collect historical agro-climatic estimates of potential oil palm yield under rain-fed, high-input production using the Global Agro-Ecological Zoning (GAEZ) dataset (version 4.0). We average the estimated yields within Thiessen polygons to obtain average oil palm potential (in kg oil/ha). We calculate euclidean distance to nearest market center utilizing the market locations from Porteous (2019). We also calculate the distance from village location to nearest primary road and nearest secondary road using spatially-explicit road network data from OpenStreetMap via the World Food Programme's GeoNode data portal. The final distance measures we calculate are the distance between the village center and the Twifo oil palm processing facility at the Twifo Oil Palm Plantaion.

Our data on tree cover and forest loss come from the Global Forest Change dataset Hansen *et al.* (2013). The GFC dataset provides high-resolution, spatially-explicit estimates of treecover and annual forest loss across the terrestrial surface of the earth. Version 1.7, the version used in the analysis for this paper, provides annual estimates of loss for 2001 - 2019. We aggregate the 30-square-meter pixels to the village-level using our Thiessen polygons and attribute any forest loss that lies within a particular polygon to the nearby village. We require that land have at least 10% of its area in tree cover to be classified as forest land at baseline. Forest loss in pixels with less than 10% will not be counted in our analysis. In the appendix we present evidence that our findings hold for more stringent baseline thresholds of 30% and 50% pixel tree cover.

The data processing described above yields a final dataset comprised of 930 unique villages in the Central region of Ghana, of which 13 we know to be participants in the Twifo contracting program. Table 2 presents summary statistics for each of these measures for contracting villages and control villages (differentiated by distance to the Twifo processing plant). Figure 2 presents the average annual rate of forest loss by each group from 2001 through 2019.

4 Empirical approach

We estimate the impact of market access on forest cover in the program villages using the contract program introduction as an exogenous shock to market access. Our identification strategy leverages the panel nature of the dataset to estimate two-way fixed-effects difference-in-differences models:

$$Defor_{it} = \beta_1 Contract_i \times Post2008_t + \alpha_i + \gamma_t + \epsilon_{it} \tag{1}$$

where $Defor_{it}$ is an annual measure of deforestation (either the inverse hyperbolic sine of hectares of treecover lost or the rate of forest loss) around village i in year t, $Contract_i$ is a binary treatment variable equal to 1 if village i receives market access via contracts, and $Post2008_t$ is a binary variable equal to 1 after 2008 when the market access contract program enters into effect. α_i and γ_t are village and year fixed effects which control for any time-invariant village characteristics that drive village forest loss (e.g. village population, agricultural suitability, etc.) and common shocks that affect all villages in our sample (e.g. oil palm prices, weather-related productivity shocks, etc.), respectively.

As the timing of the contracting program may also be correlated with other changes in the processing capacity of the Twifo mill, in certain models we augment the above equation to add an additional interaction term to differentiate villages within 50 km of the Twifo facility. Such specification should control for potentially confounding changes that might drive land use change in villages near the mill around the time of the program start.

$$Defor_{it} = \beta_1 Contract_i \times Post2008_t + \beta_2 Near_i \times Post2008_t + \alpha_i + \gamma_t + \epsilon_{it} \quad (2)$$

We also include specifications with district-specific time trends to account for potential changes in regional policies that might alter the trajectory of land conversion. These models should be interpreted with caution, however, as the 13 treated villages all lie within the same district.

The identifying assumption underlying our analysis is that, absent the program, trends in forest loss observed in the 13 villages participating in the program would have continued in parallel to trends in forest loss in our control villages after 2008. While fundamentally untestable, we attempt to validate this underlying assumption a number of different ways. First, we plot trends in average annual rates of forest loss in the treated villages and different groups of control villages (differentiated by distance to the Twifo facility) in Figure 2. Though not perfectly overlaid, the general trends are consistent prior to 2008 with a notable increase in the treated villages not seen in any of the control villages between 2008 and 2010. The absence of clear, diverging pre-trends supports the identifying assumption. Second, we compliment our difference-in-difference estimates with event study models. This allows us to test how forest loss changes in treated villages compares to changes in forest loss in control villages relative to a pre-treatment reference year. We discuss the results further in the following section, but note that leading coefficients are nearly all statistically indistinguishable from zero and demonstrate no clear, systematic trends prior to 2008. Third, we estimate our main difference-in-differences model across a wide variety of different control groups to ensure that our results are not an artifact of peculiarities of a particular control group chosen for the main analysis.

Finally, to address concerns over endogenous treatment selection we supplement the analysis via the use of a matched-sample difference-in-differences approach. For every treated village, we select four control villages that most closely resemble the treated village in question on the observable characteristics described above, specifically; share of tree cover at baseline, area of the Thiessen polygon, average elevation, ruggedness, agro-climatic estimates of potential oil palm yield, distance to nearest market, distance to nearest primary and secondary road, share of land in a game reserve or protected area, and distance to the Twifo plant. We link treated villages to their nearest-neighbors in attribute space by calculating the Mahalanobis distance between each treated village and every untreated village in the Central region of Ghana.[4] We pair the treated village with the four untreated villages that are 'closest' based on this distance measure. We sample with replacement so as to avoid order bias or bias induced by pairing otherwise unlike villages. We then estimate our main difference-in-difference model using the matched villages as the control group. To the degree to which these time-invariant characteristics explain differences in trends in forest loss (rather than simple levels), our matched difference-in-differences approach ensures that both our treated and control village samples look very similar on these dimensions and thus differences in post-treatment trends should not be explained by differences in observable village characteristics. In Figure 3 we plot average annual rate of forest loss in treated villages and the matched control villages. The pre-2008 forest loss trends are even closer to each other than in the unmatched control groups in Figure 2. Note that this similarity arises even though pre-2008 forest loss was not included

[4]The Mahalanobis distance between a treated village t and a control village c is given by $d(t,c) = \sqrt{(X_t - X_c)'\Sigma_{tc}^{-1}(X_t - X_c)}$ where X_t, X_c are vectors of attributes for the treated village and control village, respectively, and Σ_{tc} the variance covariance matrix between these vectors.

in the matching process.

We concede that matching does not eliminate potential treatment endogeneity bias if treatment selection criteria are not included in the set of observable matching variables. However, in our rural Ghanaian experimental setting it is reasonable that, in as much as program villages are strategically selected by TOPP, they are selected primarily based on these observable characteristics. Critically, the program commits the facility to cover transportation of the palm oil output from the village to the mill; potential yields and road accessibility are likely to be the most important criteria for village selection.

Across all models, we calculate standard errors that allow for across-village correlation up to 25 km. By allowing spatial dependence in the error term, we adjust inference to account for location specific shocks that might affect multiple villages at the same time without assuming independence across villages. An alternative would be to cluster at the district level to allow for arbitrary within district cross-correlation of error terms, however the small number of districts observed in the central region prevents us from obtaining consistent estimates of standard errors and is further complicated by the fact that all treated villages lie in the same district.

5 Results

We present our difference in difference estimates of equation (1) in Table 3. Columns (1) and (2) correspond to estimates using the inverse hyperbolic sine transformation of hectares lost in each year as the outcome variable and Columns (3) and (4) correspond to estimates using the rate of forest loss (loss in a given year divided by total forested area in 2000). Columns (1) and (3) present estimates of the parsimonious model while Columns (2) and (4) add district-specific trends. Our coefficient estimates in the first columns imply an average annual increase in forest loss of 117% and 40% for models without and with district time trends, respectively.[5] These imply an average additional loss of 4.2 to 1.4 ha of forest area each year per village as a result of the program. Consistent with this result, estimates in Columns (3) and (4) indicate an increase in the rate of

[5] These effects are calculated from the following partial elasticity: $e^{\hat{\beta}} - 1$ as detailed in Bellemare and Wichman (2020). These elasticities themselves are also highly significant.

forest loss by approximately one percentage point (without district time trends) to one half of a percentage point (with district time trends). Our estimates in models with the two-way fixed effects are statistically significant at the 1% level and our estimates in models with district fixed effects are significant at the 10% and 5% levels for the IHS forest loss and rate of forest loss outcomes respectively.

Table 4 builds upon the first specification by also including the interaction of an indicator for villages located within 50 km of the Twifo plant and the post 2008 indicator. This accounts for possible confounding changes to mill operations around 2008 that might have driven land use decisions in villages local to the plant. The coefficient estimates on the program effect remain statistically significant and of similar magnitude as above, assuaging concerns that local spillovers may be driving the findings. In fact, the coefficients on the interacted proximity indicator are consistently negative, indicating that, if anything, forest loss near the Twifo plant was slowing during this time. This is consistent with land use patterns around the Twifo mill. The firm operates large-scale oil palm plantations of their own around the processing plant; hence we expect to see lower rates of land conversion in adjacent areas.

Table 5 presents results analogous to those in Table 3 but uses the matched control group in lieu of all villages in the Central region as controls. Because we allow duplicates in the matching process, our estimation sample consists of our 13 treated villages and 25 matched control villages—a drastic reduction in sample size from the 930 villages in the greater Central region. Matched difference-in-difference results are very similar in magnitude and significance to those from the full sample (with estimates from models with district fixed effects now significant at the 1% level). These results indicate an increase in 3.3 to 2.2 additional ha of forest loss per village per year due to the program.

We present results from our event studies graphically in Figure 4. Panels (a) and (c) correspond to models with the IHS forest loss as outcomes and panels (b) and (d) correspond to models using the rate of forest loss as outcomes. Panels (a) and (b) present estimates using the full sample and panels (c) and (d) present estimates using the matched sample. All models include village and year fixed effects and include separate indicators for treated villages for each year from 2001 to 2019. We omit the year prior to the program (2007) to serve as our reference year, so all coefficient magnitudes can be interpreted as the change in difference between treated and control groups relative to the 2007 difference.

Three notable patterns consistently emerge in all coefficient plots. First, there is no evidence of diverging trends in leading coefficients—i.e. those year-by-year estimates display no trends, with the majority not statistically differentiable from zero. This provides additional support for the parallel trends assumption that underlies our identification. Second, across all models, there is a dramatic increase immediately following the policy. This increase coincides with the uptake in contracting presented in Figure 1 and likely captures the initial clearing for oil palm crops following the start of the program. The third pattern is a persistent rise in program village forest loss relative to control villages from 2013 through the end of the sample period. This increase is somewhat muted in the matched sample analysis compared to the full sample, but is still present and persistent through 2019.

In the appendix we present additional robustness checks of our main specification. First, we demonstrate that our findings hold for more stringent initial definitions of forest land. When we use a 30% or 50% threshold for defining forested land at baseline, our results are qualitatively similar, in fact magnitudes seem only to increase. This result indicates that our findings are not particular to land only with relatively sparse tree cover at baseline. Secondly, we present results varying the control group by distance to the Twifo facility (50, 100, and 150 km) and Region (Central Region vs any Region). Here too, our results remain remarkably stable indicating that our findings are not particular to a choice of a control group.

6 Final remarks

In this study, we provide causal evidence that contracting programs may increase forest loss. In the setting of the Twifo Oil Palm Plantation resource contracting program, villages that received marketing contracts, credit access, and reduced transportation costs saw a large increase in forest loss following the introduction of the program. In some estimates annual forest loss more than doubled. This finding is robust across a variety of models and control group specifications.

These findings should be interpreted in their appropriate context. That forest loss dramatically increased is of concern to the local ecology and the global community concerned about emissions from forest loss and biodiversity through habitat loss and oil palm monoculture. However, this program likely had a strong, positive impact on local livelihoods - one that we cannot quantify with the data we have. As such, we are unable to conduct a full benefit-cost analysis to evaluate to what degree the benefits of this program were offset by the costs.

References

Alix-Garcia, J., McIntosh, C., Sims, K. R. and Welch, J. R. (2013) The ecological footprint of poverty alleviation: Evidence from Mexico's Oportunidades program, *Review of Economics and Statistics*, **95**, 417–435.

Asher, S., Garg, T. and Novosad, P. (2020) The ecological impact of transportation infrastructure, *The Economic Journal*, **130**, 1173–1199.

Assunção, J., Gandour, C., Rocha, R. and Rocha, R. (2020) The effect of rural credit on deforestation: Evidence from the Brazilian Amazon, *The Economic Journal*, **130**, 290–330.

Baehr, C., BenYishay, A. and Parks, B. (2021) Linking local infrastructure development and deforestation: Evidence from satellite and administrative data, *Journal of the Association of Environmental and Resource Economists*, **8**, 375–409.

Bellemare, M. F. and Wichman, C. J. (2020) Elasticities and the inverse hyperbolic sine transformation, *Oxford Bulletin of Economics and Statistics*, **82**, 50–61.

Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S., Tyukavina, A., Thau, D., Stehman, S., Goetz, S., Loveland, T. *et al.* (2013) High-resolution global maps of 21st-century forest cover change, *Science*, **342**, 850–853.

Heß, S., Jaimovich, D. and Schündeln, M. (2021) Environmental effects of development programs: Experimental evidence from west african dryland forests, *Journal of Development Economics*, **153**, 102737.

Jayachandran, S. (2013) Liquidity constraints and deforestation: The limitations of payments for ecosystem services, *American Economic Review*, **103**, 309–13.

Khatun, K., Maguire-Rajpaul, V. A., Asante, E. A. and McDermott, C. L. (2020) From agroforestry to agroindustry: Smallholder access to benefits from oil palm in Ghana and the implications for sustainability certification, *Frontiers in Sustainable Food Systems*, **4**.

Lundberg, C. and Abman, R. (2022) Maize price volatility and deforestation, *American Journal of Agricultural Economics*, **104**, 693–716.

Markelova, H., Meinzen-Dick, R., Hellin, J. and Dohrn, S. (2009) Collective action for smallholder market access, *Food policy*, **34**, 1–7.

Pfaff, A., Robalino, J., Walker, R., Aldrich, S., Caldas, M., Reis, E., Perz, S., Bohrer, C., Arima, E., Laurance, W. *et al.* (2007) Road investments, spatial spillovers, and deforestation in the Brazilian Amazon, *Journal of regional Science*, **47**, 109–123.

Pfaff, A. S. (1999) What drives deforestation in the Brazilian Amazon?: Evidence from satellite and socioeconomic data, *Journal of Environmental Economics and Management*, **37**, 26–43.

Porteous, O. (2019) High trade costs and their consequences: An estimated dynamic model of african agricultural storage and trade, *American Economic Journal: Applied Economics*, **11**, 327–66.

Ruml, A. and Qaim, M. (2020) Effects of marketing contracts and resource-providing contracts in the African small farm sector: Insights from oil palm production in Ghana, *World Development*, **136**, 105110.

7 Tables

Table 1: Plot-level characteristics among contracting households

	mean	sd	min	max
Plot established after farmer entered into contract? (Yes = 1)	0.838	0.369	0.000	1.000
Use of plot before oil palm: Forest	0.197	0.398	0.000	1.000
Use of plot before oil palm: Pasture	0.376	0.485	0.000	1.000
Use of plot before oil palm: Agriculture	0.392	0.489	0.000	1.000
Use of plot before oil palm: Other	0.004	0.066	0.000	1.000
Size of oil palm plot (acres)	6.253	5.159	0.000	42.000
Amount of oil palm harvested in previous 12 months (tonnes)	45.290	109.602	0.000	1560.000
Tonnes per acre of oil palm harvested	8.712	19.571	0.000	260.000
Observations	229			

Notes: This table presents summaries of the plot-level characteristics from plots surveyed among farmers participating in the contracting program with Twifo.

Table 2: Summary statistics for Ghana Villages by grouping

	Village Type			
	Resource Contract	Control < 25 km	Control 25 - 50 km	Control > 50km
Average forest loss (ha) before 2008	3.568	9.805	10.30	5.829
	(1.507)	(8.095)	(8.703)	(6.261)
Average rate of forest loss before 2008	0.00467	0.00986	0.00916	0.00731
	(0.00146)	(0.00675)	(0.00564)	(0.00589)
Share with > 10% treecover	0.999	0.994	0.989	0.961
	(0.00246)	(0.00968)	(0.0186)	(0.0879)
Theissen Area (km sq)	8.013	11.57	12.09	8.260
	(3.466)	(8.671)	(9.012)	(5.822)
Average elevation (m)	164.6	105.1	114.2	76.02
	(10.64)	(23.01)	(37.43)	(42.73)
Ruggedness	9.241	9.177	9.149	9.230
	(3.721)	(4.515)	(5.133)	(5.899)
Ave oil palm potential (kg oil/ha)	5061.1	5068.6	4615.3	3539.3
	(13.68)	(47.21)	(619.2)	(878.1)
Distance to nearest market (km)	95.07	79.60	86.17	70.34
	(2.849)	(11.14)	(18.18)	(18.60)
Distance to primary road (km)	15.32	44.96	29.20	8.801
	(2.938)	(6.381)	(10.26)	(9.762)
Distance to secondary road (km)	2.276	2.501	2.639	1.259
	(1.681)	(2.731)	(2.649)	(1.524)
Share of land in reserve/PA	0.0603	0.128	0.101	0.0623
	(0.194)	(0.278)	(0.266)	(0.210)
Dist to Twifo plant (km)	39.66	14.79	39.23	78.39
	(3.038)	(5.926)	(7.019)	(19.98)
Dist to Benso plant (km)	82.43	67.87	85.32	120.2
	(2.780)	(9.228)	(15.25)	(27.66)
Number of villages	13	123	248	546

Notes: This table presents summary statistics across four different groups of villages; treated villages (those participating in the Twifo contracting program), non-participating villages within 25 km of the Twifo facility, non-participating village between 25 and 50 km of the Twifo facility, and non-participating villages further than 50 km from the facility. Our sample is limited to village from the central region only. Means of variables are above with standard deviations below in parenthesis.

Table 3: Estimates of resource contracting program on forest loss

	IHS Forest Loss		Forest Loss Rate	
	(1)	(2)	(3)	(4)
Contract × Post 2008	0.777***	0.324*	0.00917***	0.00532**
	(0.171)	(0.177)	(0.00199)	(0.00209)
District Time Trend		✓		✓
Village FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	17670	17670	17670	17670
R^2	0.003	0.041	0.002	0.031
Effect at mean (ha)	4.19	1.37	–	–

Notes: This table presents difference in difference estimates of the resource contract on forest loss from equation (1). The inverse hyperbolic sine of annual forest loss is the outcome in Columns (1)-(2) and the rate of forest loss (area lost divided by baseline forested area) is the outcome in Columns (3)-(4). All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 4: Estimates of resource contracting and facility proximity on forest loss

	IHS Forest Loss		Forest Loss Rate	
	(1)	(2)	(3)	(4)
Contract × Post 2008	0.906***	0.508***	0.00930***	0.00768***
	(0.176)	(0.181)	(0.00211)	(0.00219)
Near Twifo × Post 2008	-0.216*	-0.296*	-0.000224	-0.00379**
	(0.129)	(0.157)	(0.00159)	(0.00179)
District Time Trend		✓		✓
Village FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
No Obs	17670	17670	17670	17670
R^2	0.008	0.045	0.002	0.034
Effect at mean (ha)	5.26	2.36	–	–

Notes: This table presents difference in difference estimates of the resource contract on forest loss controling for proximity to the Twifo facility - the specification outlined in equation (2). The inverse hyperbolic sine of annual forest loss is the outcome in Columns (1)-(2) and the rate of forest loss (area lost divided by baseline forested area) is the outcome in Columns (3)-(4). All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 5: Matched Difference-in-difference Estimates of resource contracting program on forest loss

	IHS Forest Loss		Forest Loss Rate	
	(1)	(2)	(3)	(4)
Contract x Post 2008	0.653***	0.479***	0.00692***	0.00615***
	(0.143)	(0.145)	(0.00150)	(0.00165)
District Time Trend		✓		✓
Village FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
No Obs	722	722	722	722
R^2	0.055	0.087	0.040	0.053
Effect at mean (ha)	3.29	2.19	–	–

Notes: This table presents matched difference in difference estimates of the resource contract on forest loss from equation (1). Contract villages are matched with their four most similar control villages on the time-invariant characteristics in table 2 and only matched control villages are included in the analysis. The inverse hyperbolic sine of annual forest loss is the outcome in Columns (1)-(2) and the rate of forest loss (area lost divided by baseline forested area) is the outcome in Columns (3)-(4). All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

8 Figures

Figure 1: Farmer uptake in oil palm contracting - date of first contract by household

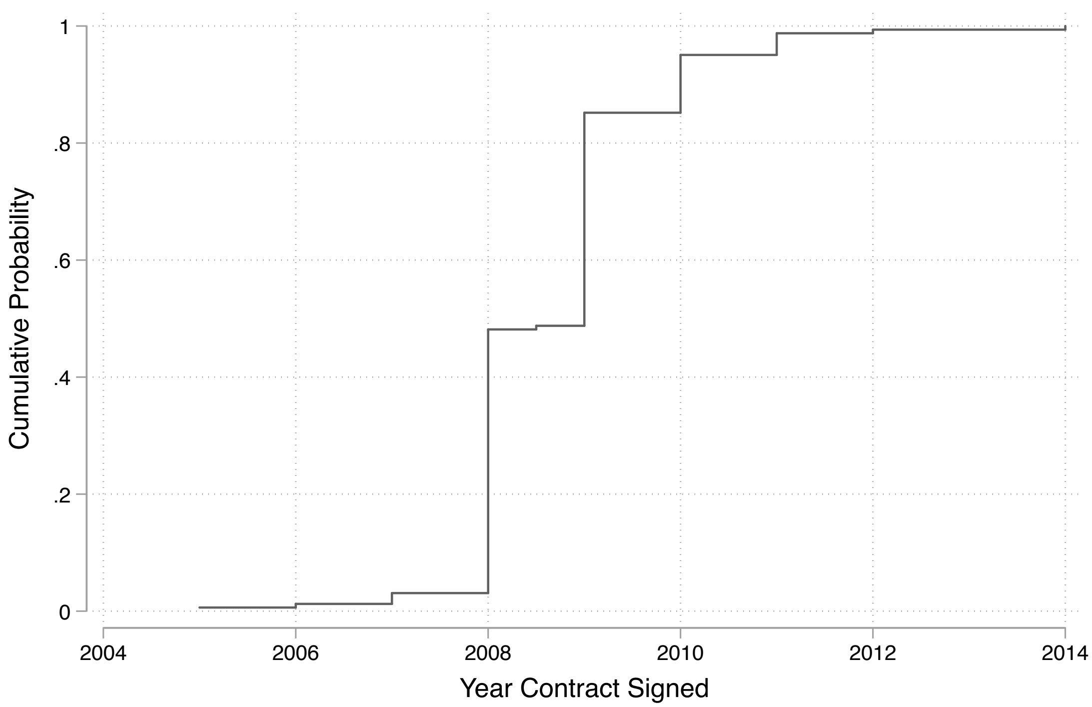

This figure plots the CDF of first contracting date from 165 households sampled across our 13 contracting villages.

Figure 2: Average annual rate of forest loss by group - Ghana

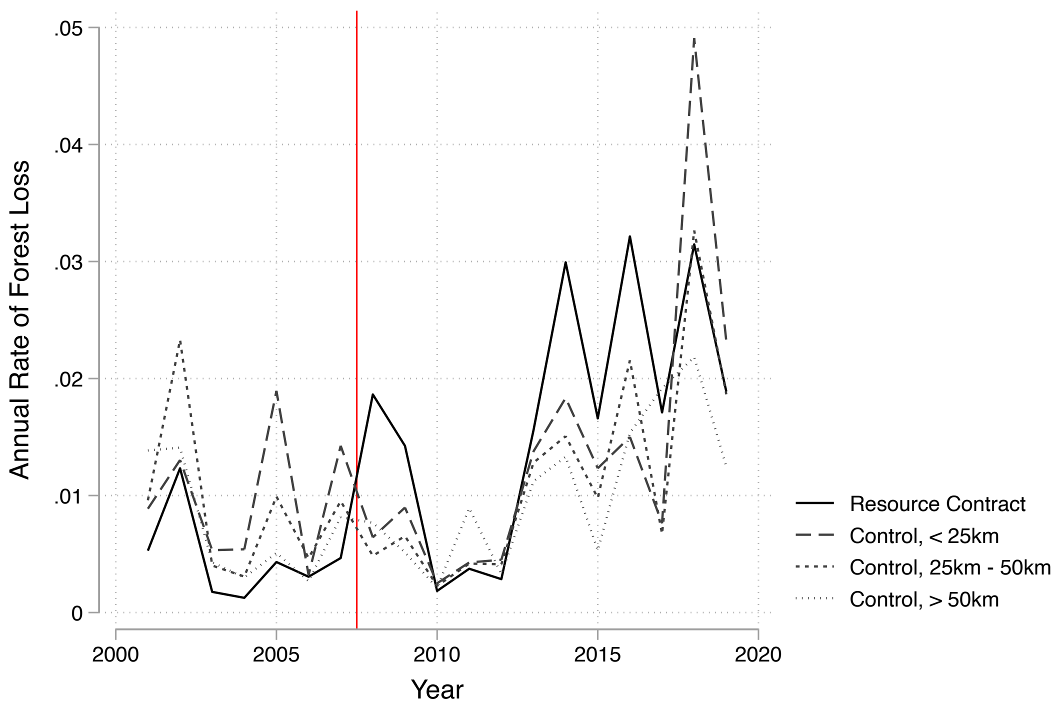

Notes: The above figure presents average annual rates of forest loss over time for four different groups of villages; treated villages (those participating in the Twifo contracting program), non-participating villages within 25 km of the Twifo facility, non-participating village between 25 and 50 km of the Twifo facility, and non-participating villages further than 50 km from the facility. The vertical line denotes the start of the contacting program.

Figure 3: Average annual rate of forest loss by group - Ghana

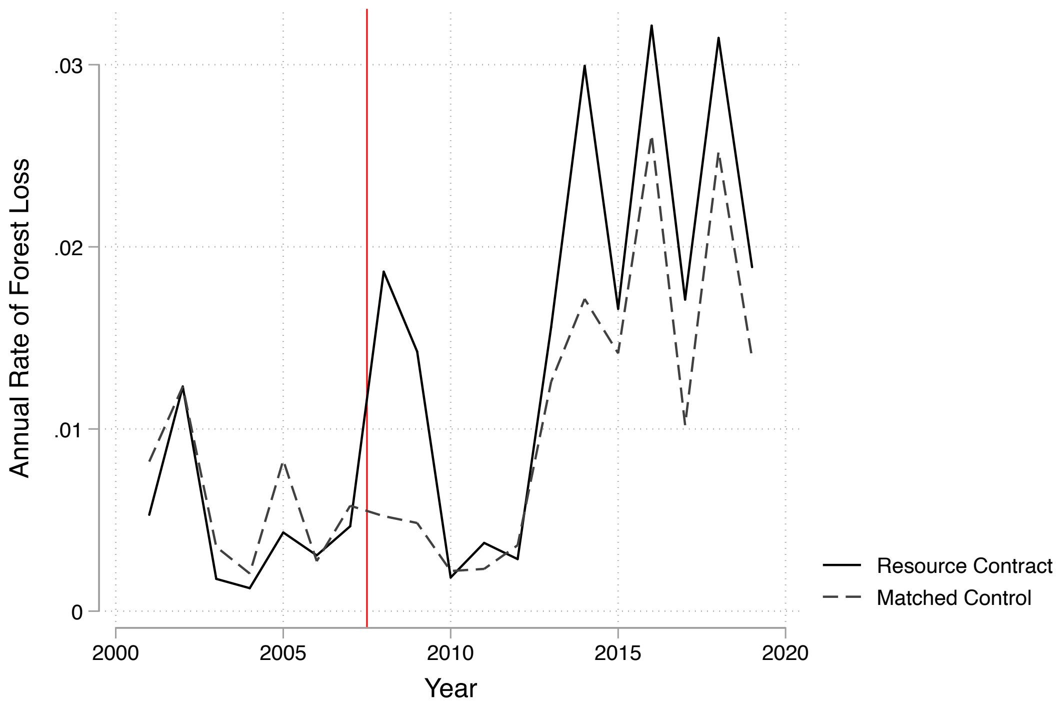

Notes: The above figure presents average annual rates of forest loss over time for the treated villages (those participating in the Twifo contracting program) and the matched, non-participating villages. The vertical line denotes the start of the contacting program.

Figure 4: Event Study Coefficient Estimates

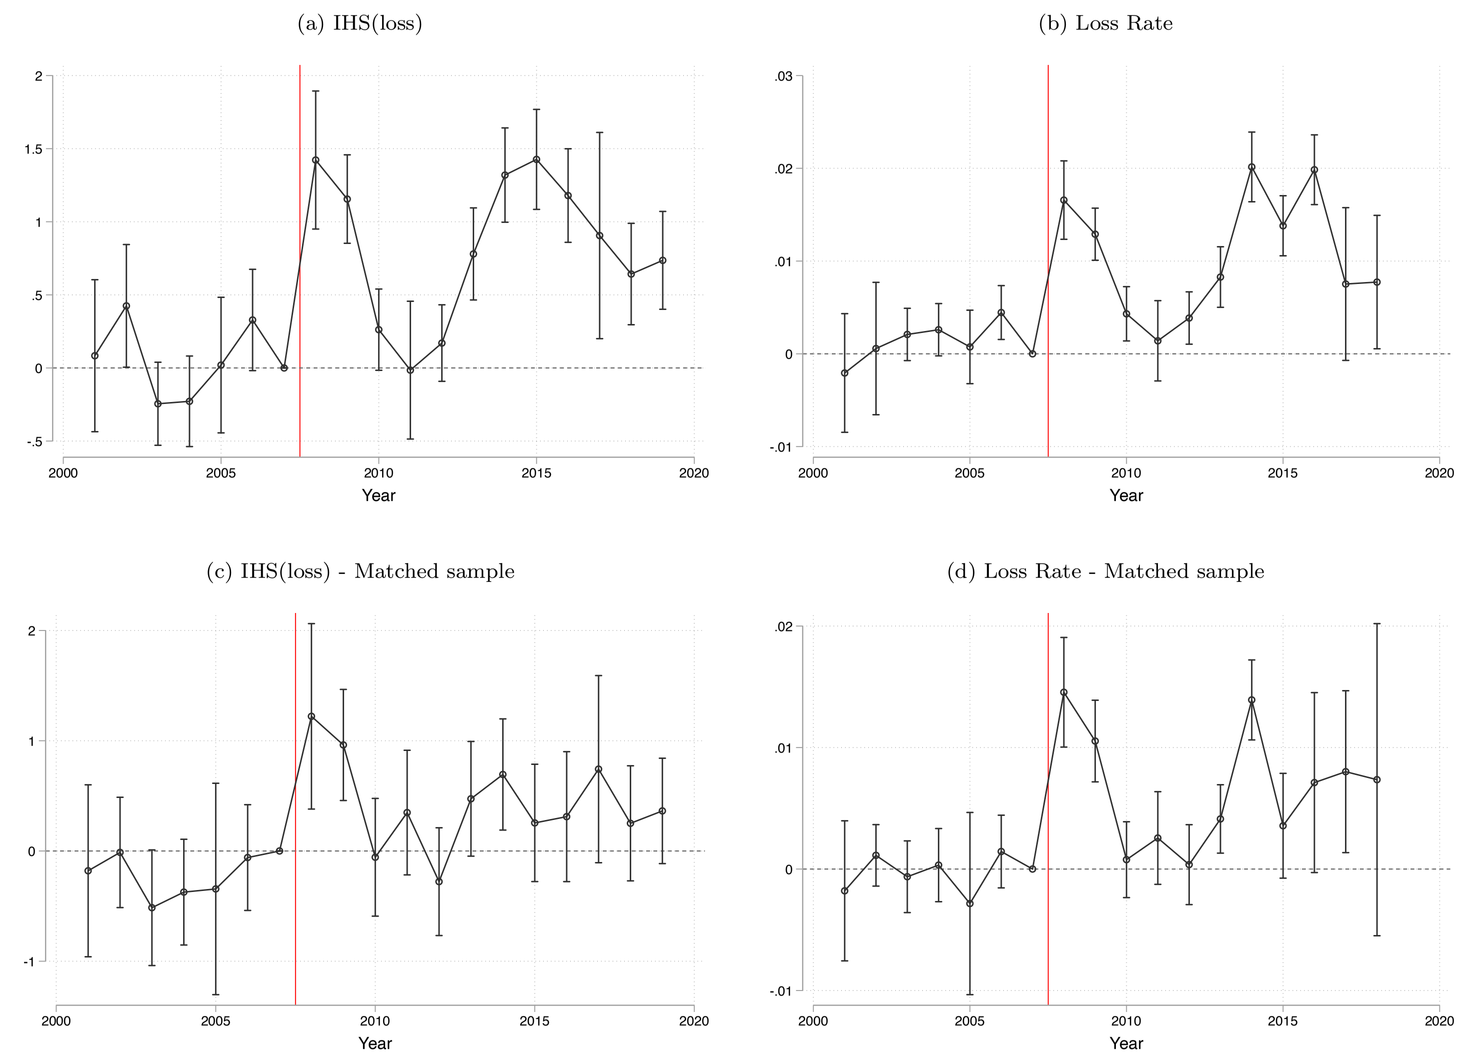

Notes: These figures present estimates from an event-study comparing annual forest loss in contracting villages to that of other Central-Region villages over time. Differences are normalized to 2007, the year prior to Twifo contracting. Panels (a) and (c) present estimates using the inverse hyperbolic sine of annual forest loss as the outcome variable, while panels (b) and (d) present estimates using the ratio of annual loss to forested area as an outcome. Panels (c) and (d) restrict the control group to matched villages only. All models include village and year fixed effects. Error bars correspond to 95% confidence intervals with Conley standard errors allowing for spatial correlation up to 25 km.

A Appendix - Additional Tables and Figures

Figure A.1: Kernel density of oil palm plot size by previous use of plot

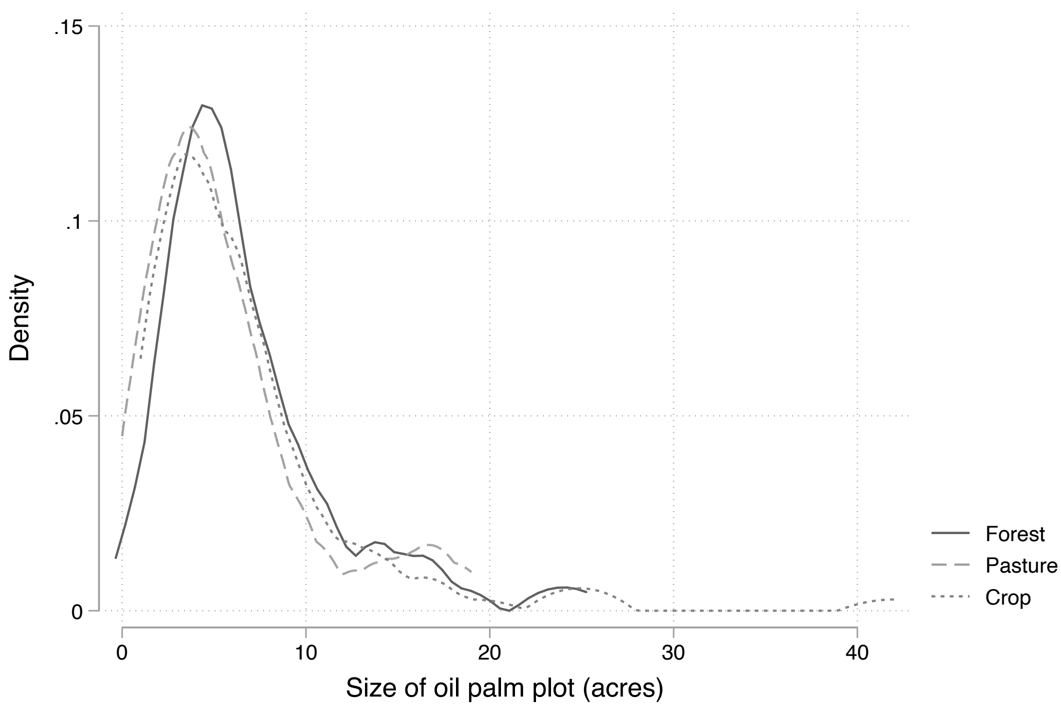

Table A.1: Main specification results across different baseline forest thresholds

	IHS Loss				Loss Rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Contract x Post 2008	0.949***	0.300	1.276***	0.266	0.0112***	0.00430*	0.0263***	0.00139
	(0.181)	(0.186)	(0.225)	(0.205)	(0.00250)	(0.00228)	(0.00727)	(0.00508)
Time Trend	None	District	None	District	None	District	None	District
Forest threshold	30%	30%	50%	50%	30%	30%	50%	50%
No Obs	17670	17670	17670	17670	17613	17613	17461	17461
R^2	0.005	0.053	0.010	0.135	0.001	0.026	0.001	0.067

Notes: This table presents results of the initial specification while imposing a higher threshold for land classified as 'forest' at baseline. The inverse hyperbolic sine of annual forest loss is the outcome in Columns (1) - (4) and the rate of forest loss (area lost divided by baseline forested area) is the outcome in Columns (5) - (8). All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A.2: Difference in differences estimates across varying contol groups - IHS forest loss

	(1)	(2)	(3)	(4)	(5)	(6)
Contract x Post 2008	0.867***	0.906***	0.712***	0.833***	0.553***	0.777***
	(0.173)	(0.178)	(0.166)	(0.171)	(0.171)	(0.171)
Twifo dist	< 50km	< 50km	< 100km	< 100km	< 150km	< 150km
Region	Any	Central	Any	Central	Any	Central
No Obs	11704	7296	36195	15637	56886	17670
R^2	0.008	0.014	0.001	0.004	0.000	0.003

Notes: This table presents difference in difference estimates of the resource contract on forest loss from equation (1) using the inverse hyperbolic sine of annual forest loss is the outcome. Samples in each column vary by two restrictions: the maximum distance from the Twifo facility and the restriction that villages belong to the Central region. All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A.3: Difference in differences estimates across varying contol groups - rate of forest loss

	(1)	(2)	(3)	(4)	(5)	(6)
Contract x Post 2008	0.00890***	0.00930***	0.00782***	0.00966***	0.00698***	0.00917***
	(0.00194)	(0.00210)	(0.00178)	(0.00204)	(0.00167)	(0.00199)
Twifo dist	< 50km	< 50km	< 100km	< 100km	< 150km	< 150km
Region	Any	Central	Any	Central	Any	Central
No Obs	11704	7296	36195	15637	56886	17670
R^2	0.003	0.005	0.001	0.003	0.000	0.002

Notes: This table presents difference in difference estimates of the resource contract on forest loss from equation (1) using the annual rate of forest loss as the outcome. Samples in each column vary by two restrictions: the maximum distance from the Twifo facility and the restriction that villages belong to the Central region. All estimates include village and year fixed effects. Standard errors allow for spatial dependence up to 25 km. Statistical significance denoted by $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$