

# The Development-Environment Tradeoff from Cash Crops: Evidence from Benin\*

Raahil Madhok<sup>†</sup>  
University of Minnesota

Leikun Yin<sup>‡</sup>  
University of Minnesota

Zhenong Jin<sup>§</sup>  
University of Minnesota

October 9, 2024

## Abstract

This paper investigates the development-environment tradeoff from cultivating cash crops. We classify cashew plantations between 2015-2021 across Benin, one of West Africa's largest cashew producers, using a deep learning algorithm validated with field visits. We document large income gains from exposure to these cashew plantations at the expense of nearby forest cover. We identify this tradeoff with cross-sectional comparisons on household survey data, two-way fixed effects with panel data, and a shift-share instrumental variables design using global cashew price shifts to instrument local cultivation. A 10 percentage point increase in land share under cashews increases local GDP by 1.3%, but reduces forest cover by 2.6%. Cost-benefit calculations show that doubling cultivation would generate \$USD 66 million in aggregate income gains but cost \$USD 147 million in terms of forest loss.

**Keywords:** Agriculture, development, deforestation, Benin

**JEL Codes:** Q01, Q56, Q15, Q20, O13

---

\*We thank Marc Bellemare for helpful feedback. All errors are our own.

<sup>†</sup>Department of Applied Economics. Email: rmadhok@umn.edu

<sup>‡</sup>Department of Bioproducts and Biosystems Engineering. Email: yin00406@umn.edu

<sup>§</sup>Department of Bioproducts and Biosystems Engineering. Email: jinzn@umn.edu

# 1 Introduction

Nearly 80% of the rural poor are involved in agriculture, and 50% are smallholder farmers ([The World Bank, 2007](#)). Smallholder agriculture can transform rural livelihoods, boost local economies, and improve food security, yet land scarcity often necessitates conversion of natural landscapes for agriculture. Over the past two decades, agriculture was responsible for 90% of global deforestation ([FAO, 2020](#)).

The goal of this paper is to quantify the tradeoff between smallholder agriculture and forest conservation in a developing country context. An extensive literature by economists, ecologists, and practitioners studies solutions for balancing the tradeoff, including improving agricultural productivity ([Abman et al., 2024](#); [Caunedo and Kala, 2021](#); [Assunção et al., 2017](#)), market access ([Abman and Lundberg, 2024](#); [Bellemare and Barrett, 2006](#)), and payments for ecosystem services ([Jayachandran et al., 2017](#); [Alix-Garcia et al., 2015](#)). Yet few papers have quantified the size of the tradeoff in the first place.

Our first objective is to quantify income effects of agricultural expansion for individual farmers and whole villages. The second objective is to quantify forest loss for each unit of agricultural expansion. We achieve the research goal and objectives with several data sources and empirical designs, ranging from cross-sectional comparisons to an instrumental variables approach. Benchmarking the development-environment tradeoff from agriculture is important for helping developing country governments understand the ecological effects of structural transformation and agricultural modernization.

We focus our analysis on cashews, a widely-farmed tree crop in many developing countries. Tree crops are a stable income source for smallholders due to high cash values and predictable yields ([Lin et al., 2021](#)). Africa produces over half of global cashew volume, mainly through smallholder farming. Since cashews are mainly cultivated by smallholders, cashew industry growth and rural incomes are tightly connected ([Yin et al., 2023](#)), making cashews an ideal case study for this paper.

Our geographic setting is Benin, a small tropical country in West Africa. This is an

ideal setting for several reasons. First, Africa is a major cashew producer, and about 80% of its output is grown by smallholders in West Africa (UNCTAD, 2021). Second, Benin is among the top 10 cashew producers globally and the third largest in West Africa (Duguma et al., 2021). Third, the Benin government spearheaded a national plan to double cashew production between 2016-2021 (Ministere de l’Agriculture de l’Elevage et de la Peche, 2017), a window that covers our study period. Production can increase by converting other land types for agriculture (extensive margin), or by increasing productivity of existing farmland through better technology and extension services (intensive margin). We primarily investigate the extensive margin, although we provide some evidence that extension access exacerbates the development-environment tradeoff.

To measure cashew production, we built a first-of-its-kind ultra high resolution data product that uses remote sensing, deep learning, and field visits to classify cashew trees at 3 meter resolution between 2015-2021. Our mapping procedure is the subject of a companion paper (Yin et al., 2023). We also aggregate these tree-level rasters to 200 meter “plantations”. While other tree crops such as rubber and oil palm grow on large commercial farms, which enables mapping by low-resolution satellites (Putra and Wijayanto, 2023), cashew trees grow on small fields with crown sizes  $< 5m$ , preventing classification by most satellites. Our novel classification algorithm overcomes this issue, enabling one of the first studies of cashew farming in the environmental economics literature.

To measure income and forest cover, we use both household survey and satellite data. Survey data are from the 2018 Demographic and Health Survey (DHS), which surveys 3000 households in our study area. The DHS reports household wealth scores, which we complement with our own index based on ownership of various assets. Next, we measure “exposure” to cashew cultivation by the proximity from each survey cluster to the nearest cashew plantation. The survey also reports forest cover around households, which enables a comparison among cashew-exposed and non-exposed households.

For our panel analysis, we use a new gridded GDP product from Chen et al. (2022)

which converts 1km resolution nightlights to GDP. We also use nightlight data directly, which is a common proxy for economic activity (Henderson et al., 2012) and bypasses concerns about the nightlights-to-GDP conversion. Lastly, we measure time-varying forest cover from a complementary satellite product to the one used in DHS. To address concerns that it may misclassify cashew trees as forests, we conduct a validation exercise which shows that over 90% of cashew plantations are outside of pixels classified as forest.

The unique combination of cashew maps, household surveys, and satellite data enable multiple research designs to quantify the development-environment tradeoff from cashew cultivation. We set the stage by comparing income and forest cover among households exposed and non-exposed to cashew cultivation. While cross-sectional comparisons deliver important descriptive evidence, in our second design, we estimate two-way (village and year) fixed effects regressions on panel data to account for unobservables such as geography, crop suitability, and changing agricultural demand. While two-way fixed effects (TWFE) are more credible than a cross-sectional design, reverse causality between income and cultivation remains an important identification concern.

For the third research design, we instrument cashew cultivation with a shift-share style instrument that interacts global cashew prices with local production at baseline. This instrument exploits farmer responses to plausibly exogenous fluctuations in global prices, and incorporates heightened price sensitivity for farmers more involved in cashew cultivation. A potential violation of the exclusion restriction arises if global cashew prices covary with other commodity prices. In a robustness check, our estimates are unchanged when controlling for other agricultural commodity prices, excluding tree nuts.

Our analysis yields two key results. First, cashew cultivation increases household incomes and boosts local economic activity. Cross-sectional evidence shows that household exposure to cashew cultivation is associated with 34% higher wealth. Extension access, measured by proximity to cashew seedling nurseries, accentuates the wealth impact by 30%. Village-level panel estimates show that a 10 percentage point (p.p.) increase in land

area under cashew cultivation increases local GDP by 1.3%. Each marginal plantation increases GDP by 0.7%. Instrumental variable (IV) estimates corroborate the TWFE estimates: a 10 p.p. increase in cashew land share increases economic activity (nightlights) by 2.6%<sup>1</sup>, and by 1.5% for the marginal plantation.

The second result is that economic benefits are at the expense of forests. Cross-sectional, TWFE, and IV estimates all point to forest loss. Cross-sectional estimates show that forest cover is 2% lower around households exposed to cashew cultivation. Panel estimates show that doubling cashew cultivation in a village prompts a 36% reduction in forest cover. Lastly, IV estimates show that doubling cashew cultivation causes forest cover to decline by 26%. One-quarter of cashew expansion is at the expense of forests.

Having documented a clear development-environment tradeoff from cashew cultivation, we next probe mechanisms. We provide suggestive evidence since the DHS is our only source of detailed household data. The DHS does not report whether households grow cashews, meaning we cannot distinguish if our estimated income benefits accrue to cashew growers themselves, or represent broader economic spillovers. With this in mind, we find that our main DHS estimates showing wealth gains for cashew-exposed households are even larger among landowners. It is hard to imagine why income effects are stronger for cashew-exposed landowners compared to non-exposed landowners in the same village unless the former grow cashews themselves. To probe mechanisms underlying forest loss, we show that the farms of cashew-exposed households are 3% larger than farms of further-away households in the same village. Since forest cover is also lower around cashew-proximate households, this suggests that larger farms among these households may materialize through converting adjacent forestland.

The paper concludes by incorporating our estimates into a simple cost-benefit analysis. We use the social cost of carbon to convert forest loss into dollars. We calculate that

---

<sup>1</sup>We use nightlights as the outcome in IV regressions because GDP is available for only two years. Since 2SLS restricts identifying variation in the explanatory variable by design, there is insufficient variation to detect effects. We thus use nightlights as the outcome since it is available for the full study period.

doubling cashew cultivation in Benin would cost \$USD 147 million in terms of forest loss and generate \$USD 66 million in income gains. For every dollar earned from cashew cultivation, the ecological cost is more than twice as large. This represents a lower bound since the value of many forest ecosystem services are excluded from the cost calculation.

**Literature Contributions:** This paper adds to a seminal literature on tradeoffs between economic development and environmental quality ([Grossman and Krueger, 1995](#); [Dasgupta et al., 2002](#); [Stern et al., 1996](#)). Much of this work confronts the controversial Environmental Kuznets Curve in theory or with national data. Instead, we study the development-environment tradeoff within one country with highly disaggregated data, which enables a characterization of the tradeoff for individual households or villages.

We also join a broader literature on structural transformation and agricultural development ([Bustos et al., 2016](#); [Emerick, 2018](#); [Moscona, 2019](#); [Madhok et al., 2024](#)). This work largely studies how agricultural development affect affects on- and off-farm labor at coarser geographic scales, whereas we focus on estimating income benefits from expanded household cultivation within small administrative units.

Lastly, we extend a new literature on agriculture and deforestation ([Abman and Lundberg, 2024](#); [Abman et al., 2024](#); [Brewer et al., 2024](#); [Green et al., 2005](#)). This work mainly focuses on the intensive margin, showing that improvements in agricultural productivity can spare nature. In contrast, we focus on the extensive margin and show that expanded cultivation displaces forests. An exception is [Brewer et al. \(2024\)](#), who show that agricultural labor loss leads to farm size contraction and, separately, that labor loss reduces deforestation. We extend this by directly connecting agricultural expansion to forest loss.

The paper proceeds as follows. The next section provides background on the economy and environment in Benin. Section 3 describes the cashew maps and other data. Section 4 outlines three empirical strategies for quantifying the development-environment tradeoff from agriculture. Section 5 presents main results and Section 6 concludes.

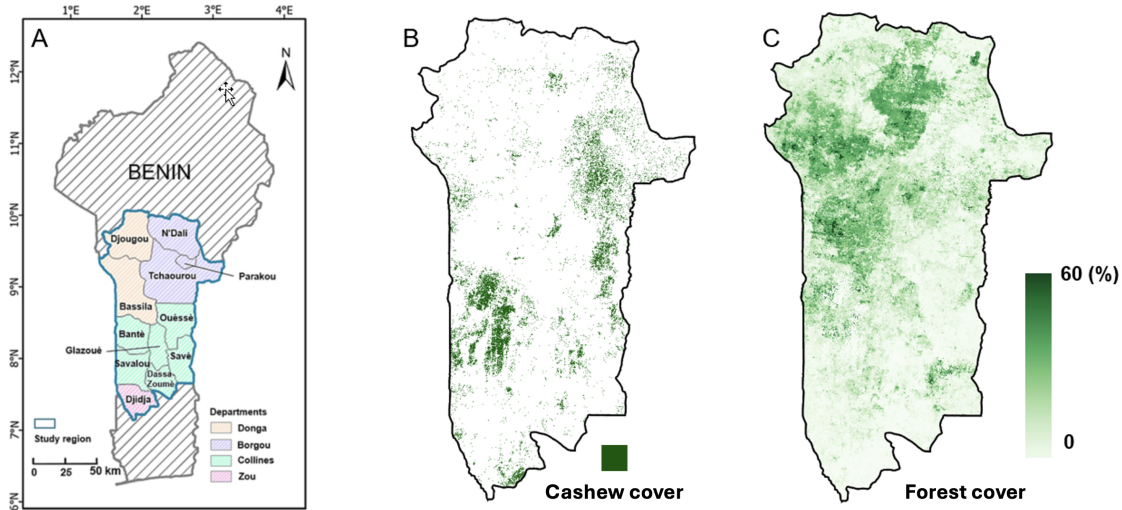


Figure 1: Study Area, Cashew Distribution, and Forest Cover

Note: Panel A shows the study area. Panel B shows the distribution of cashew plantations in the study area at 3m resolution. Panel C plots 2015 forest cover at 200m resolution. Cell values denote % forest cover.

## 2 Background

Benin lies on the West African coastline, bordered by Togo to the west and Nigeria to the east. The administrative structure features 12 departments, subdivided into 77 communes and 546 arrondissements. Arrondissements comprise several villages and form the primary administrative unit for local governance. While part of our analysis (Section 4) is at the household level, the majority is at the arrondissement level.

Agriculture forms the backbone of Benin's economy. Of its 13 million population, 70% rely on agriculture for their livelihoods. The sector accounts for over 30% of GDP and over 80% of export revenue. Cashews are among the main cash crops, in part due to the tropical savanna climate where cashew trees grow well. Nearly 200,000 smallholder cashew growers contribute to the industry, which generates 15% of export revenue (Yin et al., 2023). Recognizing the potential of cashew agroforestry for raising smallholder incomes, the government announced a plan to double cashew production between 2016-2021 (Ministere de l'Agriculture de l'Elevage et de la Peche, 2017). Our study period falls directly into this period of expansion.

Cashew cultivation is mainly practiced in Central Benin. Our study area encompasses the four departments in this region: Donga, Borgou, Collines, and Zou (Figure 1A). These departments comprise 13 communes and 103 arrondissements, collectively home to 25% of the national population. Figure 1B maps the location of cashew plantations within our study area (see Section 3.1 for data details). Although plantations are seen throughout Central Benin, production is concentrated near the eastern and western borders.

Stimulating the cashew sector poses a threat to local ecology since agriculture and forests compete for land in Benin. Nearly 11,000 ha. of forest was lost in Benin during our study period, approximately 7% of the total forest area in 2000 (Global Forest Watch, 2024), with agriculture cited as the main driver (World Bank, 2020). The tension between agriculture and forests can also be observed at a local level: Figure 1C shows that cashew plantations are situated in the regions with high forest cover.

Economic theory states that, as cashew production becomes more lucrative, farmers may intensify farming on existing land (intensive margin), or expand farming by converting adjacent forestland for agriculture (extensive margin). The impact of cashew expansion on forest cover is therefore an empirical question. We focus on the extensive margin since we lack data on inputs such as labor, agrochemicals, and other capital. In any case, we expect the extensive margin response to dominate in our context since farmers are typically factor market constrained in developing countries (Conning and Udry, 2007).

### 3 Data

In the absence of administrative data on cashew cultivation, we develop a first-of-its-kind series of cashew maps using remote sensing, deep learning, and validation data from the field. We complement this with wealth data from household surveys as well as gridded GDP data. Forest cover is measured with satellite data. The final panel covers the study area from 2015 to 2021. This section describes the data and provides summary statistics.



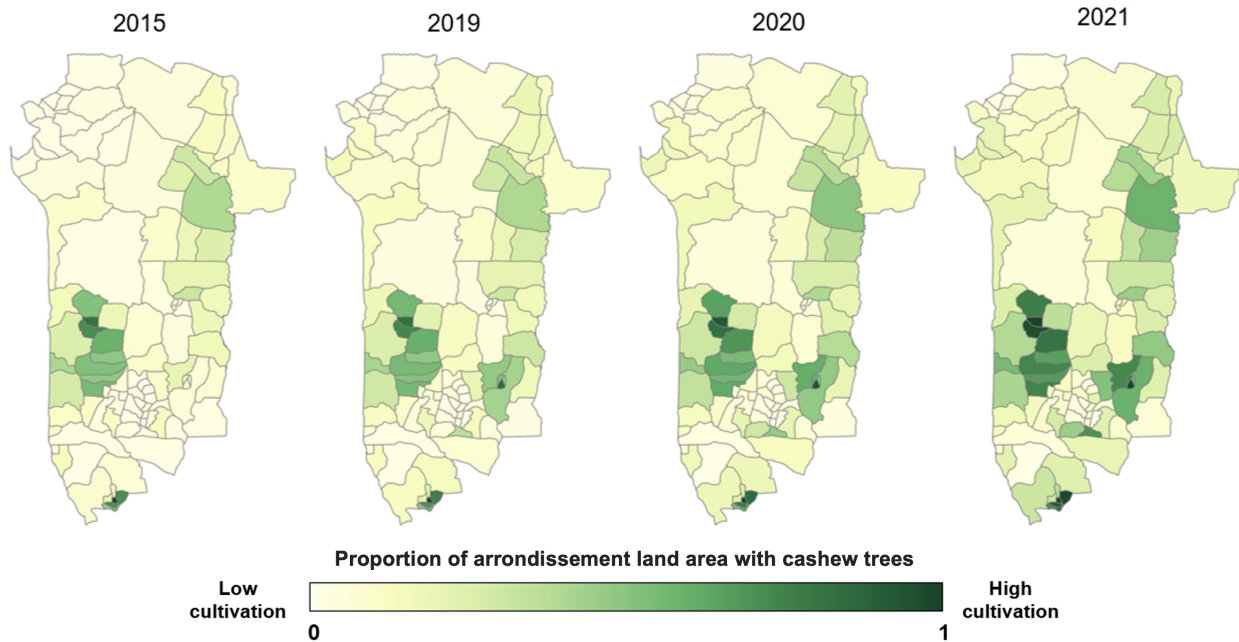


Figure 2: Spatial distribution of cashew coverage (%)

Note: Borders delineate arrondissements in the study area. Shading represents the percentage (%) of grid cells in an arrondissement with cashew crops, as classified by the model.

### 3.1 Smallholder Cashew Plantations

#### 3.1.1 Remotely Sensed Cashew Maps

Large-scale studies of tree crops are often restricted to commercially-grown crops on large farms (e.g. palm oil, rubber), since these can be mapped by lower-resolution satellites. (Putra and Wijayanto, 2023). In contrast, cashews have small crowns ( $< 5m$ ) and are grown on irregular or small fields, leading to a dearth of satellite data on their distribution and, therefore, a knowledge gap about their environmental and economic ramifications.

We overcome this data gap by developing a novel set of remotely-sensed maps of cashew cultivation in Benin for years 2015 and 2019-2021 using a combination of deep-learning algorithms and field data for ground-truthing. Details of the mapping procedure are provided in a companion paper (Yin et al., 2023) and summarized briefly here. The 2015 map was generated by applying the U-Net classification algorithm (Ronneberger et al., 2015) to Benin's national-scale aerial imagery (original resolution is 0.5m and resam-

pled to 3m). Maps for 2019-2021 were produced using the more recently available Planet imagery (original resolution is 2.4m and resampled to 3m) and the SpatioTemporal Classification with Attention (STCA) algorithm, which accounts for phenological changes in cashew trees and accurately delineates the boundaries of cashew plantations.

Accuracy of the cashew classification algorithm was assessed using a ground-truthing procedure based on stratified sampling. Ground labels were collected for 1400 validation points by our local partner, the TechnoServe BeninCajù program, who sent a field survey team to the sampling locations to record land cover types. Approximately 85% were correctly classified by our remote sensing procedure, representing state-of-the-art performance for smallholder tree crop classifications.

Appendix Figure B1 shows classification output for four example cashew plantations. Figure 2 plots cashew cultivation at the arrondissement level across the study period. Shading represents the fraction of grid cells growing cashews. An increase in the density of cashew plots is observed as green areas become darker over time. An increase in the extension margin is also observed as some arrondissements with no cashew crops in 2015 come under cultivation in later periods (e.g., the southeast region).

In addition to ground-truth data, TechnoServe also provided data on cashew nurseries and smallholder trainings between 2018-2021. Cashew nurseries are established to support sustained production across Benin. Trainings are conducted by TechnoServe to educate smallholders in plant spacing, cashew tree pruning, and replanting. We obtained locations of 157 cashew nurseries and 2200 training events. We use this data to explore treatment heterogeneity by extension access.

### 3.1.2 Cashew Prices

We use data on global wholesale cashew prices to construct an instrument for local cashew production in Section 4.3. We calculate yearly prices by dividing global supply value (\$USD billions) by production volume (metric tons) reported in the International Nut

and Dried Food Council Statistical Yearbooks ([INDFC, 2023](#)). We convert to 2015 prices to account for inflation. In a robustness check, we separately control for fluctuations in agricultural commodity prices. This data is obtained from the IMF Commodity Database. We use the food price index, which includes cereals and legumes, but not tree nuts.

## 3.2 2018 DHS Survey

**Measuring Household Wealth:** For the cross-sectional analysis, we measure household wealth from the 2018 DHS survey. The survey is nationally representative and covers 14,156 households, of which 20% are in our study area. The closest period of cashew data is from 2015. Our analysis thus estimates medium-term (three-year) impacts.

We use two wealth indicators from the DHS survey, one provided and one constructed. First, we use the wealth index provided in the survey, which ranks household wealth on a scale of 1 to 5. Second, we construct another wealth index based on whether households have a constructed floor, washing machine, internet, cellphone, and television. We follow the [Anderson \(2008\)](#) approach whereby each variable influences the index proportional to the information it adds. We first compute the z-score of each underlying variable and then compute the index as a weighted sum of these standardized values with weights equal to the row sum of the inverse covariance matrix.

Since the DHS are available in 2018, we cannot estimate household wealth impacts of cashew expansion *over time*. Instead, we exploit variation in exposure to cashew cultivation *across space* by matching households to the nearest plantation. We do this by first aggregating the 2.4m tree crop raster to 200m “plantations”. Then, we match households to the nearest plantation based on euclidean distance from the centroid of their DHS sampling cluster<sup>2</sup> to the nearest plantation<sup>3</sup>. This method measures exposure to cashew cultivation even if the household does not cultivate cashews themselves. A positive as-

---

<sup>2</sup>A DHS cluster typically corresponds to a census enumeration area i.e., a neighborhood ([Demographic and Health Surveys Program, 2023](#)).

<sup>3</sup>The DHS randomly displaces the coordinates of each cluster by 0-5km for privacy ([Burgert et al., 2013](#)).

sociation between inverse distance and wealth indicates the presence of spillovers from nearby cashew cultivation into the local economy.

**Measuring Forest Quality:** The DHS reports forest quality in the area surrounding each DHS sample cluster using the Enhanced Vegetation Index (EVI). EVI reflects vegetation “greenness” and is a commonly used measure of forest cover. In contrast to wealth data, which is available three years after the first cashew map, we use the DHS-provided EVI from 2015, the same year as the cashew plantation data. It is calculated as an average within a 2km buffer around each rural survey cluster and 10km buffer for urban clusters. At such small radii, we can elicit partial equilibrium forest impacts of cashew cultivation and more comfortably interpret effects as forest clearing for agriculture.

**DHS Covariates:** We include several geographic and household covariates from the DHS. Geographic covariates are at the level of a DHS cluster (typically, a village) and include: rain, temperature, latitude, longitude, and nightlights. Weather and latitude/longitude control for differential poverty rates and agricultural development by climate and altitude. Nightlights ensure that comparisons are made between households with the same level of urbanization. Household covariates include household size and language, both which may predict wealth and agricultural involvement.

### 3.3 Other Satellite Data

#### 3.3.1 Gridded GDP

Measuring the dynamics of agricultural development requires a time-varying wealth measure at high resolution. As a starting point, we use nightlight intensity, which is considered a strong proxy for local GDP ([Henderson et al., 2012](#)). Second, we also measure local GDP directly using a new, off-the-shelf gridded GDP product developed by [Chen et al. \(2022\)](#). Their algorithm approximates cell-level GDP at 1km resolution for 2015 and 2019 by scaling light intensity values by real GDP growth rates. We defer the reader

to [Chen et al. \(2022\)](#) for computational details. To calculate the outcome, we first sum GDP across pixels in each arrondissement. We then do the same for gridded population counts<sup>4</sup> and divide GDP by population to obtain GDP per capita.

**Data Limitations:** While the [Chen et al. \(2022\)](#) data is among the first to provide GDP estimates at high resolution, it suffers two important shortcomings. The first is conceptual: GDP is an inherently aggregated measure, reflecting gross value added from all production or consumption activities in a country. Dividing this into 1km cells assumes we can measure the contribution of a handful of individuals to national GDP based on how their economic production is picked up by nightlights. The second issue is econometric: luminosity values feature non-classical measurement error since the satellite has difficulty detecting lights at low levels. This can attenuate GDP values in rural areas where cash injections from cashew farming may not be enough to be detected from outer space.

While we acknowledge these shortcomings, hundreds of papers have found strong correlations between nightlights and GDP (see [Gibson et al. \(2020\)](#) for a review), which lends at least some credibility to its use for measuring economic activity. To address the second issue, we complement our fixed effects approach with an instrumental variables design. Orthogonality between the instrument and luminosity error allow us to circumvent issues of non-classical measurement error in our 2SLS estimates.

### 3.3.2 Forest Cover

Forest cover is obtained from the Vegetation Continuous Fields (VCF) satellite product ([Townshend et al., 2017](#)), which measures percent forest cover in 200m gridcells<sup>5</sup>. Our main forest cover measure in the panel fixed effects and IV analysis is the share of arrondissement land area under forest. To compute this, we first calculate the weighted

---

<sup>4</sup>Gridded population counts are obtained from WorldPop ([www.worldpop.org](http://www.worldpop.org))

<sup>5</sup>We opt not to use the popular Global Forest Change (GFC) satellite product ([Hansen et al., 2013](#)) because GFC classifies pixels only as deforested or not *relative to 2000*. As such, it does not capture forest regrowth, nor does it provide a continuous measure of forest cover. This is important because VCF data shows that forest cover slightly *rose* during our study period, a trend that would not be picked up by GFC.

sum of pixel values in each arrondissement-year, with weights equal to pixel area, and then divide by arrondissement land area.

A crucial measurement concern is whether the VCF satellite classifies newly established cashew plantations as forest. This would attenuate our coefficient for the forest cover impact of cashew plantations, since declines in forest cover from agricultural encroachment would be offset by misclassifying new cashew plantations as forest gain. We rule out this misclassification concern in Appendix C by showing that new plantations are almost always located outside of forest pixels. Using a 15% threshold to classify VCF pixels as forest, we show that over 90% of new cashew plantations in each year are outside of forests (Table C9; Figure C1). This holds true even under stricter thresholds up to 25%. This exercise improves confidence that our estimates of forest cover impacts in Section 4 are devoid of mechanical endogeneity from misclassifying cashew trees as forest.

### 3.3.3 Covariates

We include three covariates in the the panel analysis: rain, temperature, and drought intensity. Controlling for weather is important because changes in climate can affect both agricultural productivity and economic output. Gridded annual temperature ( $^{\circ}\text{C}$ ) and rainfall (mm) are from the ERA5 product on a  $0.125^{\circ} \times 0.125^{\circ}$  grid (Hoffmann et al., 2019). Drought intensity is measured using the gridded ( $0.5^{\circ} \times 0.5^{\circ}$ ) Standardized Precipitation Evapotranspiration Index (SPEI) from the SPEIbase. SPEI measures the difference between potential evapotranspiration and precipitation. For all gridded covariates, we extract the mean over cells within arrondissements for each year of the study.

## 3.4 Summary Statistics

Table 1 summarizes the main outcome and explanatory variables in the DHS (Panel A) and gridded panel dataset (Panel B). In Panel A, “N” is the number of households surveyed in the sample region where cashew data was collected. Cashew data are from 2015

Table 1: Summary Statistics

	N	Mean	Std. Dev.
<i>Panel A: DHS (2018)</i>			
Dist. to Nearest Plantation (km)	2866	4.26	11.17
# Plantations w/n 10km	2866	934.95	1088.66
Wealth Index (scale 1-5)	2866	3.00	1.28
<i>Panel B: Panel (2015-2021)</i>			
Cashew Plantation Coverage (%)	412	0.39	0.25
Cashew Tree Coverage (%)	412	0.13	0.15
Cashew Density (per $km^2$ )	412	7.04	4.57
GDP Per Capita (USD)	206	1759.69	908.63
Forest Cover (share of land area)	412	0.07	0.03

Note: Panel A summarizes household variables from the 2018 DHS survey. Cashew variables (first and second row) are from 2015. Panel B summarizes the panel data at the arrondissement-year level. “Cashew plantation coverage” is the fraction of 200m grid cells in an arrondissement with cashews. “Cashew tree coverage” is computed in the same way using 3m grid cells.

and DHS data are from 2018. The typical household lives about 4km from the nearest cashew farm. The standard deviation is nearly three times the mean, indicating substantial variation in cashew proximity across space. In terms of density, households are surrounded by 935 cashew farms within 10km, which is about three per sq.  $km^6$ . Lastly, the average household has a wealth score of 3 out of 5.

In Panel B, data are at the arrondissement-year level. GDP is available only for 2015 and 2019. The typical arrondissement cultivates cashew trees on 13% of its land area over the study period. Cultivation is relatively intense: there are 7 plantations per  $km^2$ . In terms of wealth, GDP per capita is about \$USD 1,760, which matches official national statistics. Lastly, the typical arrondissement has about 7% forest cover. This highlights the need for sustainable agriculture given that forest cover is already so low.

## 4 Research Design

We estimate the development-environment tradeoff in agriculture using three complementary research designs. The cross-sectional design compares households close and

<sup>6</sup>The area of a circle with radius 10km is  $\pi(10)^2 \approx 314km^2$ . 935 farms in  $314km^2$  is  $935/314 \approx 3/km^2$ .

far from cashew farms. The TWFE design compares arrondissements across time and space with different levels of cultivation. The IV strategy exploits global commodity price shocks to generate plausibly exogenous variation in cashew cultivation.

## 4.1 Cross-Sectional Design

**Main Equation:** The cross-sectional design compares household wealth, measured in the 2018 DHS survey, between households living near and far from cashew farms in 2015. This elicits medium-term associations between cashew cultivation and wealth. We also compare contemporaneous forest cover around each survey cluster among cashew-exposed and non-exposed households. We estimate the following equation:

$$Y_{iacd} = \beta_1 \text{Cashew}_{iacd} + \Gamma X'_{iacd} + \alpha_a + \epsilon_{iacd} \quad (1)$$

where  $Y_{iacd}$  is the wealth index or forest quality for household  $i$  in arrondissement  $a$  of commune  $c$  and department  $d$ .  $\text{Cashew}_{iacd}$  is the inverse distance between household  $i$ 's DHS cluster and the nearest cashew plantation.  $X'_{iacd}$  is a vector of climate, geography, and household covariates (Section 3.2 for details). The arrondissement fixed effect,  $\alpha_a$ , absorbs time-invariant differences across arrondissements, leaving  $\beta_1$  to be estimated off of comparisons between households *within the same arrondissement* that are close to and far from cashew plantations. We use household sampling weights in Equation 1 to ensure estimates accurately reflect the population<sup>7</sup>.

The coefficient of interest is  $\beta_1$ , the change in wealth or forest cover for a unit change in cashew exposure. While  $\alpha_a$  helps remove endogeneity from cross-unit comparisons,  $\beta_1$  still cannot be interpreted causally. First, conditional on controls and fixed effects, there may still be unobserved differences across households that are correlated with both  $\text{Cashew}$  and  $Y$ . Second,  $\beta_1$  represents a snapshot during the DHS survey year, whereas

---

<sup>7</sup>The survey weight is calculated as the inverse of household  $i$ 's selection probability multiplied by the inverse of the household response rate in the stratum.



cashew cultivation, income, and forest cover are dynamic processes influenced by past decisions and time-varying shocks. These processes are overlooked in a cross-section, limiting the ability to draw causal inferences. Third, when the outcome is wealth, the prospect of high returns may draw households into cashew farming, leading to reverse causality. We therefore interpret  $\beta_1 > 0$  as a positive *association* between cashew cultivation and wealth and address the endogeneity of *Cashew* with TWFE and IV strategies in Section 4.2 and 4.3, respectively.

**Treatment Heterogeneity:** We leverage our extension data to investigate heterogeneity of  $\beta_1$  in Equation 1 by extension service access. This reveals how important extension services are for enhancing cashew income. At the same time, this heterogeneity may exacerbate forest degradation. We estimate the following specification:

$$Y_{iacd} = \beta_1 \text{Cashew}_{iacd} + \beta_2 (\text{Cashew}_{iacd} \times \text{Extension}_{iacd}) + \beta_3 \text{Extension}_{iacd} + \Gamma X'_{iacd} + \alpha_a + \epsilon_{iacd} \quad (2)$$

where  $\text{Extension}_{iacd}$  is either the inverse distance from household  $i$ 's DHS cluster to the nearest seedling nursery, or the number of smallholder training events held in the arrondissement in 2018. All other terms are the same as Equation 1. The interaction term  $\beta_2$  is the coefficient of interest. When  $Y$  indicates wealth,  $\beta_2 > 0$  implies that the cashew-wealth relationship is accentuated when households have better extension access. When  $Y$  indicates forest quality,  $\beta_2 < 0$  implies that extension access exacerbates forest degradation, in line with [Abman and Lundberg \(2024\)](#).

## 4.2 Two Way Panel Fixed Effects Design

Whereas the cross-sectional approach exploited only spatial variation for identification, the TWFE strategy adds time as a second layer of variation since the share of arrondisse-

ment land area planted with cashews changes over time. Our fixed effects strategy compares GDP and forest cover *within* arrondissements at different cultivation levels, controlling for time fixed effects. We estimate the following equation:

$$Y_{acdt} = \beta_1 \text{CashewShare}_{acdt} + \Gamma X'_{acdt} + \alpha_a + \gamma_{dt} + \epsilon_{acdt} \quad (3)$$

where  $Y_{acdt}$  is the outcome of interest, either log nightlights, log GDP per capita or forest cover, in arrondissement  $a$  of commune  $c$  and department  $d$  at time  $t$ . Nightlights and forest cover data span all four years whereas GDP is available only for 2015 and 2019.  $\text{CashewShare}_{acdt}$  is the share of land area in arrondissement  $a$  under cashew cultivation.  $X'_{acdt}$  is a vector of covariates including temperature, rainfall, and drought intensity. Arrondissement fixed effects,  $\alpha_a$  absorb time-invariant differences between arrondissements. Department-by-year fixed effects,  $\gamma_{dt}$ , account for department-specific factors that change over time, such as regional agricultural policy or growth trajectories.

When GDP is the outcome,  $\beta_1 > 0$  indicates that expanding land under cashew cultivation is associated with local economic growth. When forest cover is the outcome,  $\beta_1 < 0$  indicates that this cashew expansion is at the expense of forests. These two relationships are evidence of a development-environment tradeoff in agriculture.

We investigate heterogeneity by extension access by interacting  $\text{CashewShare}_{acdt}$  with measures of seedling nurseries and smallholder training events. We measure seedling access by the number of seedling plants available per nursery in arrondissement  $a$ . This takes advantage of our detailed extension data and accounts for the fact that extension access is low when seeds are unavailable, even if there are many nurseries nearby.

**Threats to Identification:** TWFE addresses several endogeneity concerns in the cross-sectional design. Whereas  $\beta_1$  in Equation 1 is biased by omitted household variables, Equation 3 is at the arrondissement level, and arrondissement fixed effects remove bias from all time-invariant unobservables across arrondissements, such as geography, crop

suitability, and baseline poverty. Other sources of bias in Equation 3 may arise from changing regional politics or agricultural demand. Department-by-year fixed effects account for these factors along with any other time-varying factors that equally affect arrondissements within a department. This leaves arrondissement-level confounding variables that change over time as the main “backdoor” channel, such as drought intensity. The vector  $X'_{acdt}$ , which includes climate covariates, partially mitigates this threat.

Despite these advantages of the TWFE strategy, reverse causality still poses a crucial concern for identifying economic impacts in particular. Cashew cultivation can benefit the local economy, yet better local opportunities can reallocate labor out of agriculture. These opposing forces may attenuate  $\beta_1$  when estimated with OLS. We turn to an instrumental variable solution to this problem next.

### 4.3 Two-Stage Least Squares Design

#### 4.3.1 Instrument Definition

To address reverse causality between cashew cultivation and economic outcomes, we construct an instrument for cashew cultivation based on global wholesale cashew prices. The objective is to isolate variation in local cashew production generated by farmer responses to global cashew price shifts. Since individual farmers are price-takers, their cultivation response to global price changes represents a source of variation that is plausibly uncorrelated with other local determinants of cashew cultivation.

We compute a shift-share style instrument (Goldsmith-Pinkham et al., 2020) by interacting global cashew prices (the shift) with baseline production (the share) as follows:

$$z_{acdt} = \text{CashewPrice}_t \times \text{CashewShare}_a.$$

$\text{CashewPrice}_t$  is the global cashew price in year  $t$  (Section 3.1.2).  $\text{CashewShare}_a$  is the fraction of land area with cashew plantations in 2015, a fixed value that measures arrondisse-

ment exposure to global cashew prices. The combination of shift and share yields an instrument that exploits plausibly exogenous variation in arrondissement exposure to global cashew price shocks. The exclusion restriction rests on the argument that the shift (price) is *global* and cashew-specific. Therefore, conditional on controls and fixed effects, it affects *local* Benin GDP only through influencing the extent of local cashew production<sup>8</sup>. While this is fundamentally untestable, we provide corroborating evidence and discuss threats to identification further below.

### 4.3.2 2SLS Equation

The first stage of the 2SLS model tests whether global cashew price shocks disproportionately impact arrondissements that are more engaged in cashew production. We estimate the following first stage equation:

$$CashewShare_{acdt} = \delta \cdot z_{acdt} + \Gamma X'_{acdt} + \alpha_a + \gamma_{dt} + \epsilon_{acdt} \quad (4)$$

where  $z_{acdt}$  is the instrument described above and remaining terms are the same as Equation 3. The first stage coefficient,  $\delta$ , captures variation in cashew cultivation that is plausibly orthogonal to local agricultural incentives.  $\delta > 0$  indicates that “cashew-exposed” arrondissements supply more output as cultivation becomes more lucrative globally.

In the second stage, we use predicted cultivation,  $\widehat{CashewShare}_{acdt}$ , to explain local economic development and forest cover:

$$Y_{acdt} = \beta_1 \widehat{CashewShare}_{acdt} + \Gamma X'_{acdt} + \alpha_a + \gamma_{dt} + \epsilon_{acdt} \quad (5)$$

where  $Y_{acdt}$  represents economic outcomes or forest cover in arrondissement  $a$  at time  $t$ . Remaining terms and subscripts are the same as Equation 3.

---

<sup>8</sup>Anecdotal conversations with Beninese farmers reveals that farmers are indeed sensitive to global cashew prices, even though cashews are perennial crops and may be less price sensitive than annual crops. Farmers stated that when global prices are low, they reduce inputs rather than reduce production directly.

**Outcome of Interest:** We use log nightlight intensity as the main outcome in Equation 5, in contrast to the cross-sectional and panel design where GDP was the main outcome. We make this choice because GDP is available for only two years, whereas nightlights is available for the full four-year study period. The sample size reduction from using GDP as an outcome was acceptable in the TWFE design since all available variation in  $CashewShare_{acdt}$  was used. However, in the 2SLS design,  $\widehat{CashewShare}_{acdt}$  contains less variation by construction which, after partialling out  $\alpha_a$  and  $\gamma_{dt}$ , leaves little left to identify  $\beta_1$  with only two data points per arrondissement. Using nightlights can identify  $\beta_1$  with more precision while doubling the estimation sample. To see this, Table A1 reports residual variation ( $1 - R^2$ ) after partialling out arrondissement and department-year fixed effects from log nightlights and GDP per capita. There is 11 times more identifying variation remaining when nightlights is the outcome, largely due to their being more years of data for this variable.

### 4.3.3 Instrument Validity

Our shift-share design mitigates at least three threats to identification. One concern is that global cashew price shifts can affect local development through demand-side factors. In our model,  $\gamma_{dt}$  absorbs regional demand shifts and, moreover, the interaction of price with  $CashewShare_a$  ensures identification is based off of *differential* responses of arrondissements to global price shifts.

A second concern is that global cashew prices may co-vary with other agricultural commodity prices. This would violate the exclusion restriction since  $z_{acdt}$  would impact  $Y_{acdt}$  through changes in the cultivation of other crops. Yet we expect weak covariance since  $CashewPrice_t$  is based on customs duties paid on actual cashew trade (INDFC, 2023) whereas global agricultural commodity prices typically reflect futures prices. As a formal check, we test robustness to controlling for the interaction of  $CashewShare_a$  with IMF-reported agricultural commodity prices (which excludes tree nuts) in Equation 5.

The third concern is that  $CashewShare_a$  is potentially endogenous if agricultural areas systematically differ from non-agricultural ones. Although *fixed* differences are absorbed by  $\alpha_a$ , baseline cultivation may predict differential *changes* in outcomes. For example, places with better credit access may experience relatively higher income growth. If credit access is correlated with  $CashewShare_a$ , then  $z_{acdt}$  would pick up differential time paths of income in high- and low-cultivation areas—an alternative channel between the instrument and income—thereby violating the exclusion restriction.

We address the third concern in two ways. First, we include arrondissement-specific linear time trends as a robustness check, which controls for divergent trends across arrondissements that vary at a constant rate. Second, we follow [Goldsmith-Pinkham et al. \(2020\)](#) and provide a formal validity test that correlates  $CashewShare_a$  with a selection of potential confounders from the DHS. The variables are aggregated to the arrondissement level and include the share of households with credit union access, an educated household head, cattle, agricultural jobs, and professional jobs. We use the DHS because gridded data on such detailed demographics are unavailable. Table [A2](#) presents the key result:  $CashewShare_a$  is uncorrelated with the majority of “backdoor channels” that violate the exclusion restriction. An exception is that places growing cashews are seemingly less educated. We therefore include the interaction of education with global cashew prices in a robustness check (Table [4](#)) to account for the instrument potentially picking up differential income trends by household education status regardless of their involvement in cashew cultivation. Overall, lack of correlation between  $CashewShare_a$  (the share) with all other confounders in Table [A2](#), coupled with plausible exogeneity of global cashew prices (the shift), builds confidence in the validity of the shift-share instrument.

Table 2: DHS Results: Cashew Plantations and Household Wealth Index

	(1)	(2)	(3)
Near cashew (=1)	0.343** (0.145)		
Log(Proximity Cashew)		0.066 (0.045)	0.247** (0.119)
Log(Proximity Cashew) $\times$ Log(Proximity Nursery)			0.164*** (0.033)
Log(Proximity Cashew) $\times$ Trainings			0.015*** (0.004)
Household Controls	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes
Outcome Mean	3.004	3.004	3.004
Arrondissement FEs	✓	✓	✓
Observations	2866	2866	2866
$R^2$	0.512	0.511	0.521

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are a household cross section. The outcome is wealth on a scale of 1-5. “Near Cashew” indicates whether households’ DHS cluster is less than the median distance to the nearest cashew plantation. “Proximity Cashew” and “Proximity Nursery” are inverse distance to the nearest cashew plot and seedling nursery, respectively. “Trainings” is the number of trainings held in the arrondissement in 2018. Regressions include survey weights and controls for household size, language, rain, temperature, lat/lon, and nightlights. Standard errors robust to heterogeneity.

## 5 Main Results

This section presents evidence on the development-environment tradeoff in agriculture in Benin. Across three different research designs, we find evidence of economic benefits from cashew agriculture, but at the cost of forest degradation.

### 5.1 Cross-Sectional Estimates

**Outcome: Wealth Index** Estimates of Equation 1 are shown in the first two columns of Table 2. The outcome is the DHS-provided wealth index (scale 1-5). In column 1, the explanatory variable is a dummy for whether household  $i$  lives near a plantation, measured as below-median proximity to the nearest plantation. Column 2 uses log proximity to the

nearest plantation. Proximity in both columns is measured by inverse distance so that larger values indicate greater exposure to cashew cultivation.

The coefficient in column 1 is positive and statistically significant, suggesting that exposure to cashew cultivation is associated with higher wealth. The point estimate implies that cashew exposure is associated with an 11% ( $=0.343/3.004$ ) increase in wealth. In column 2, the point estimate implies that households twice as proximate to cashew plantations are 2% ( $=0.066/3.004$ ) wealthier, although the coefficient is insignificant.

**The Role of Extension Services:** Column 3 presents estimates of Equation 2, where plantation proximity is interacted with two measures of extension services. The first interaction coefficient indicates whether seedling access accentuates the wealth impact of cashew cultivation. The point estimate is positive and statistically significant, implying that the wealth impact is larger when households have better seedling access. The point estimate is large: wealth impacts of households twice as close to the nearest seedling nursery are 66% ( $=0.164/0.247$ ) larger than households with poor extension access.

The second interaction coefficient is also positive and significant, suggesting that access to smallholder trainings accentuates the wealth impact of cashew cultivation. The point estimate suggests that an additional training in the arrondissement increases the wealth impact of cashew cultivation by 6% ( $=0.015/0.247$ ). These results imply that extension services play an important role in enhancing wealth impacts of cashew farming.

**Sensitivity: Constructed Wealth Score:** The cross-sectional results are robust to an alternative wealth score based on inverse covariance weighted wealth proxies (see Section 3.2 for measurement). Appendix Table A3 presents estimates of Equation 1 and 2 using our constructed wealth score as the outcome (in standard deviations). Across all columns, the main and interaction coefficients are very similar to Table 2. Living near a cashew plantation is associated with a wealth increase of  $0.26\sigma$ . For comparison, the coefficient in Table 2 column 1 is equivalent to  $0.27\sigma$ <sup>9</sup>. As before, the continuous distance measure

---

<sup>9</sup>The SD of the DHS wealth index is 1.27. The coefficient of 0.34 units equals  $0.34/1.27 = 0.27SD$ .



has no direct effect on wealth (column 2), but the interaction with distance to the nearest nursery is again positive and significant (column 3). Lastly, The heterogeneous effect of training events remains positive, but loses precision. These results bolster confidence that the estimates in Table 2 are not artifacts of the way DHS measures wealth.

**Outcome: Forest Quality:** Table A4 reports estimates of Equation 1 with forest quality as the outcome, measured by log of EVI in a close radius around each DHS sampling cluster. The point estimate in column 1 is negative, statistically significant, and suggests that farms near cashew plantations have 2.2% less surrounding forest cover compared to far-away farms. The coefficient remains negative and statistically significant when the explanatory variable is a continuous distance measure (column 2). Column 3 tests heterogeneity by extension access and shows that better access to seed nurseries exacerbates forest degradation. The interaction coefficient implies that seed access worsens forest degradation by 20% ( $=0.002/0.01$ ). In contrast, training workshops neither mute nor exacerbate forest loss among cashew-exposed farmers.

Overall, these results point to a clear positive association between cashew cultivation and household wealth, and a clear negative association with forest quality. This points to a potential tradeoff between agricultural development and environmental degradation. Yet cross-sectional associations cannot be interpreted causally. We move towards causal estimates with panel data next, and with IV estimates afterwards.

## 5.2 Panel Estimates

Next, we show that the development-environment tradeoff in agriculture is visible in a dynamic setting with panel data. In line with the results of the previous section, we find evidence of agricultural development at the expense of ecological degradation.

**Expanding Cultivated Area:** Columns 1, 3 and 5 of Table 3 report estimates of Equation 3. Column 1 presents estimates with log nightlights as the outcome, using the same

Table 3: Panel Estimates: Cashew Coverage, Income, and Forest Cover

	(1)	(2)	(3)	(4)	(5)	(6)
Outcomes are in logs	NTL	NTL	GDP	GDP	Forest	Forest
Cashew Coverage (%)	0.408 (0.340)	0.305 (0.498)	0.134*** (0.045)	0.153*** (0.051)	-0.360* (0.182)	-0.647** (0.272)
Cashew Coverage (%) × Plants/Nursery		-0.032 (0.026)		-0.002 (0.001)		0.025* (0.014)
Cashew Coverage (%) × Trainings		0.023 (0.070)				-0.020 (0.030)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department × Year FEs	✓	✓	✓	✓	✓	✓
Observations	412	309	206	206	412	309
R <sup>2</sup>	0.946	0.967	0.998	0.998	0.876	0.899

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcomes are log nightlight intensity, log GDP per capita, and log forest cover. “Cashew coverage” is the fraction of 200m grid cells in an arrondissement growing cashews. “Plants/Nursery” is the mean number of seedlings (in hundreds) per nursery in the arrondissement at baseline. “Trainings” is the number of training events held in the arrondissement in 2018. All specifications include controls for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

base data as [Chen et al. \(2022\)](#) before they convert to GDP. We observe positive economic effects, although the coefficient is noisy. Column 3 reports estimates with log GDP per capita as the outcome. The point estimate is positive, statistically significant, and implies that a 10 percentage point (p.p.) increase in area under cashew cultivation increases local GDP by 1.34%. In column 5, the outcome is log forest cover, measured as log share of land area under forest. The point estimate implies that a 10 p.p increase in cashew cultivation reduces forest cover by 3.6%. These results suggest that income benefits from cashew cultivation are at the expense of forests degradation.

Columns 2, 4, and 6 estimate heterogeneity by extension access. Seedling access is measured as number of plants per nursery. This accounts for low access when limited seedlings are available, even if there are many nurseries nearby. Since trainings are measured in 2018, we restrict the sample period to 2019-2021 in columns 2 and 6. We cannot estimate heterogeneity by training access in column 4 since GDP is not available after

2019. Unlike the cross-sectional case, we find virtually no evidence of heterogeneity in the panel design. Extension services appear to play a limited role in enhancing the impacts of cashew cultivation.

**Impact of a Marginal Plantation:** Table A5 presents estimates of Equation 3 with the explanatory variable measuring the number of cashew plantations (200m pixels with cashew trees) per  $km^2$  in an arrondissement. Once again, when log nightlights is the outcome (column 1), the point estimate is positive but noisy. When log GDP is the outcome (column 3), the point estimate is positive, statistically significant, and implies that each marginal plantation is associated with 0.7% higher GDP. The point estimate in column 5 implies that each marginal plantation is associated with 2% forest loss.

Columns 2, 4, and 6 explore heterogeneity in cashew density by extension access. Similar to Table 3, we document virtually no role of extension access in accentuating the development-environment tradeoff.

**Sensitivity: Ultra High-Resolution Data:** Table A6 estimates Equation 3 using cashew *trees* rather than *plantations* to construct explanatory variables. Specifically, we use the ultra-high resolution (3 meter) cashew maps to measure the share of arrondissement land area with cashew trees as well as the number of cashew trees per  $km^2$ .

The income effect of expanding cashew area (Table A6, Column 2) remains positive, statistically significant, and similar to the baseline estimate. The point estimate implies that a 10 p.p. increase in land area with cashew trees is associated with a 2.3% GDP increase, a quantity within the confidence interval of the plantation-based estimate in Table 3. The negative forest cover coefficient (column 3) also remains similar, but precision declines. This is likely because the 3m cashew pixels are substantially smaller than the 200m forest cover pixels. Additional cashew trees may not be captured by forest satellite images until a large enough forest area is diverted, leading to noisier correlations between cashew coverage and forest cover even when aggregated at the arrondissement level.

The coefficient in columns 4 and 5 describes the economic impact of planting a single cashew tree. The point estimate on GDP remains positive, statistically significant, and implies that planting 10,000 more trees is associated with a 1% GDP increase. Column 6 again shows a noisy negative effect on forest cover. The magnitude is near-zero as it describes the impact of planting a single cashew tree on forest cover in the arrondissement.

### 5.3 Instrumental Variable Estimates

Having established a development-environment tradeoff under the cross-section and TWFE designs, we now investigate whether the result holds under an IV strategy that isolates plausibly exogenous variation in cashew production.

#### 5.3.1 First Stage Estimates

Table A7 presents first stage estimates of Equation 4. The outcome is  $CashewShare_{acdt}$ , the arrondissement land area under cashew cultivation. Column 1 is our preferred specification, with  $z_{acdt}$  as the explanatory variable. Column 2 adds the interaction of global agricultural commodity prices with baseline cashew coverage as a control to account for correlation of the instrument with commodity prices of other agricultural goods. Both explanatory variables are standardized for ease of interpretation.

Across both columns, the instrument strongly predicts local cashew cultivation. The point estimate in column 1 implies that a  $1\sigma$  increase in global cashew prices leads farmers to increase production by 12% relative to the mean. The opposing coefficients in column 2 indicate crop substitution: when the price of other agricultural commodities increases (row 2) relative to cashews, cashew production declines. The cashew price effect (row 1) nearly triples after accounting for this downward pressure. The F-statistic is above rule-of-thumb levels in both columns, supporting the case for instrument relevance.

Table 4: Second Stage: Cashew Plantations and Economic Activity

Outcome in logs	(1)	(2)	(3)	(4)	(5)	(6)
Cashew Coverage (%)	0.263** (0.131)		0.267* (0.152)	0.252 (0.195)	0.229 (0.244)	0.219** (0.105)
Cashew Density (per $km^2$ )		0.015** (0.007)				
Ag. Price $\times$ Crop Coverage (%)	No	No	Yes	No	No	No
Cashew Price $\times$ Education	No	No	No	Yes	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Linear Trend	No	No	No	No	Yes	No
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department $\times$ Year FEs	✓	✓	✓	✓	✓	✓
Resolution	200m	200m	200m	200m	200m	3m
Observations	412	412	412	264	412	412
F-Stat	16.74	16.02	14.19	19.67	6.13	28.33

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcome is log nightlight intensity. “Cashew Coverage” is the fraction of 200m cells growing cashews, instrumented with the interaction of global cashew prices with baseline cashew coverage. “Cashew Density” is the number of cashew plots per  $km^2$ , instrumented in the same way. “Ag. Price” is the global food price index. All regressions control for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

### 5.3.2 Second Stage Estimates: Economic Activity

Estimates of Equation 5 are presented in Table 4. The outcome is log nightlights, which is available for the full study period and contains substantially more identifying variation than the limited GDP data<sup>10</sup>. We find  $\beta_1 > 0$  across a variety of specifications: cashew cultivation has economic benefits. Column 1 is the main specification, where cashew coverage is instrumented with the interaction of global cashew prices with baseline cashew land share. The point estimate is positive, statistically significant, and implies that a 10 p.p. increase in the cashew land share increases economic activity by 2.6%. Column 2 investigates the impact of a marginal plantation by instrumenting the number of cashew

<sup>10</sup>There is 11 times less identifying variation remaining after partialling out fixed effects from log GDP per capita (Table A1). Intuitively, we are underpowered when log GDP is the outcome since it is available for only half of the study period. Moreover, 2SLS exploits only a small fraction of available variation in  $CashewShare_{acdt}$  for identification. Low sample size combined with less variation means there is insufficient variation to detect a signal.

plots per  $km^2$ . The point estimate implies that converting an additional plot for cashew cultivation increase local economic benefits by 1.5%. Given our supporting evidence for instrument validity in Section 4.3.3, we interpret these two estimates as causal evidence of positive economic benefits from cashew cultivation in Benin.

**Robustness Checks:** The remaining columns explore sensitivity of the estimates. Column 3 controls for the interaction of global agricultural commodity prices and baseline crop coverage. Commodity prices exclude tree nuts and baseline crop coverage measures agricultural land fraction *excluding* cashews. This specification thus controls for income effects from farmers reacting to agricultural prices potentially correlated with global cashew prices. The estimate remains remarkably similar, highlighting the credibility of our instrument since it seemingly does not pick up adjustment margins through other crops. As discussed in Section 4.3, this may be because cashew prices are based on physical trade whereas commodity prices reflect futures markets.

Column 4 controls for the interaction of global cashew prices with household education. This accounts for the correlation of baseline cashew coverage with education (Section 4.3.3, Table A2) and, therefore, the possibility that educated households experience stronger income growth even in the absence of cashew cultivation. One limitation is that education is only available in the arrondissements covered by the DHS, which more-than-halves the sample size and reduces statistical power. The point estimate remains virtually unchanged, suggesting that differential income growth by education level does not bias our main estimates. Precision declines due to the sample size reduction.

Column 5 adds an arrondissement-specific linear time trend i.e., the interaction of years with a full set of arrondissement fixed effects. While column 4 accounted for differential trends by an observed variable, the set of controls in column 5 flexibly accounts for any “drift” in expected income driven by *unobserved* factors that vary across arrondissements at a constant rate over time. For example, incomes may grow at a faster rate in arrondissements with more fertile land. The point estimate remains very similar. Preci-

Table 5: Second Stage: Cashew Plantations and Forest Cover

Outcome in logs	(1)	(2)	(3)	(4)	(5)	(6)
Cashew Coverage (%)	-0.257*** (0.076)		-0.249*** (0.084)	-0.181** (0.069)	-0.328** (0.156)	-0.286*** (0.067)
Cashew Density (per $km^2$ )		-0.015*** (0.004)				
Ag. Price $\times$ Crop Coverage (%)	No	No	Yes	No	No	No
Cashew Price $\times$ Education	No	No	No	Yes	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Linear Trend	No	No	No	No	Yes	No
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department $\times$ Year FEs	✓	✓	✓	✓	✓	✓
Resolution	200m	200m	200m	200m	200m	3m
Observations	412	412	412	264	412	412
F-Stat	16.74	16.02	14.19	19.67	6.13	28.33

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcome is share of arrondissement land area under forest cover. “Cashew Coverage” is the fraction of 200m cells growing cashews, instrumented with the interaction of global cashew prices with baseline cashew coverage. “Cashew Density” is the number of cashew plots per  $km^2$ , instrumented in the same way. “Ag. Price” is the global food price index. All regressions control for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

sion declines since adding 103 controls (one for each arrondissement) makes for a very demanding specification.

Column 6 tests robustness to the resolution of the instrument. Global cashew prices are interacted with baseline cashew coverage calculated from the high-resolution (3 meter) maps. The coefficient remains stable and precision improves. Overall, these 2SLS estimates corroborate the cross-sectional and TWFE estimates (Table 3), suggesting that we are capturing a robust relationship between agricultural and economic development.

### 5.3.3 Second Stage: Forest Cover

Table 5 presents the same 2SLS specifications but with log forest cover as the outcome, where forest cover is the share of land area under forest cover.  $\beta_1 < 0$  in all specifications: cashew cultivation degrades forests. The point estimate in column 1 implies that

a 10 p.p. increase in the land share under cashew cultivation reduces forest cover 2.6%. Column 2 documents the impact of a marginal plantation. The point estimate implies that converting an additional plot for cashew cultivation reduces local forest cover by 1.5%

The remaining columns document sensitivity to the same robustness tests. Estimates of forest cover loss are remarkably stable when controlling for correlated agricultural commodity price changes (column 3), the interaction of global cashew prices with household education (column 4) and a linear time trend (column 5). Results are also stable when using higher-resolution (2.4m instead of 200m) cashew data (column 6).

Overall, these results provide causal evidence that cashew cultivation is at the expense of forest cover. The question remains *how* cashew cultivation translates into economic benefits and environmental degradation. We turn to an exploration of mechanisms next.

## 5.4 Mechanisms

Are the economic benefits documented in previous sections reflective of farmers reaping the rewards of their own cultivation? Does forest loss reflect farmers clearing forests for agriculture? We turn back to the DHS to elucidate mechanisms. One limitation is that the survey is cross-sectional. Another is that we do not know whether households grow cashews. We thus view the this investigation as suggestive evidence of mechanisms.

Recall that Table 2 showed that living near cashew plantations is associated with 11% higher household wealth. To investigate whether these benefits accrue to cashew farmers themselves, we estimate heterogeneity of Equation 1 by household land ownership. A positive interaction term indicates that households living near cashew plots, who are also farmers themselves, derive more income than landowners living further away.

Columns 1 and 2 of Table A8 present the results. Among households living near cashew plots, landowners are nearly twice as wealthy as non-landowners. This is suggestive evidence that the wealth effect in Table 2 is driven by farmers directly benefiting from growing cashews. To understand why, recall that “Near Cashew” equals one



for households within median distance (900 meters) from the nearest cashew plantation. While landowners may be wealthier than non-landowners for many reasons, it is hard to imagine why the wealth gap is exacerbated among households 0-900 meters from cashew plots unless the land-owning households in this group are the ones growing cashews themselves. Column 2 shows sensitivity to using our constructed wealth score as the outcome. Coefficient sign and magnitude is roughly similar, but precision declines.

Next, we explore whether cashew cultivation reduces forest cover by agricultural encroachment into forestland. Column 3 of Table A8 presents estimates of Equation 1 with log farm size as the outcome. Land owners living near cashew plantations cultivate 3.4% ( $=0.114/3.360$ ) larger farms than land owners living further away in the same arrondissement. Despite our household and geographic controls to account for other differences across these households, the fact that cashew-proximate landowners continue to cultivate larger farms suggests that they may be converting adjacent natural land for cashew agroforestry. This aligns with our estimates of lower forest cover in a small radius around these same cashew-proximate households (Column 3, Table A4)

## 5.5 Cost-Benefit Estimates

As a final exercise, we calculate aggregate costs and benefits of doubling cashew cultivation. We study this scenario because the Benin government announced plans to double cashew production during our study period ([Ministere de l'Agriculture de l'Elevage et de la Peche, 2017](#)). The challenge is that our benefit measure (income) and cost measure (forest cover) are in different units, preventing direct comparison. We therefore convert forest loss into dollars using the social cost of carbon.

To calculate costs, we use the IV coefficient to estimate deforestation from doubling cashew land share. The average arrondissement is 36,000 ha., with 13% (Table 1) under cashews (4,680 ha.) and 7% forests (2,520 ha.). Since the outcome in Equation 5 is log forest share, the new forest share after increasing cashew coverage by another 13p.p. can

be written as  $e^{\log(0.07)+(-0.257 \times 0.13)} \approx 0.068$ , which is  $0.068 \times 36,000 = 2,448$  ha. Therefore, doubling cashew land share in an arrondissement reduces forest cover by 72 ha. We next convert this to dollars using the social cost of carbon. While we acknowledge that forests provide a variety of ecosystem services, we focus on carbon for simplicity. The carbon stock of forests in Benin is about 389 tons of  $CO_2$  per hectare (Houssoukpèvi et al., 2022). Using the most recent social cost of carbon estimate of \$51 per/ton of  $CO_2$ , and the fact that there are 103 arrondissements in the study area, the aggregate cost of doubling cashew cultivation is about \$USD 147,126,024.

To calculate benefits, we use the TWFE coefficient since the outcome is in dollars. Given mean GDP per capita of \$USD 1,760 in an arrondissement, the new amount after doubling cashew land share is  $e^{\log(1760)+(0.134 \times 0.13)} \approx$  \$USD 1,791 per capita. Doubling cashew land share therefore generates a gain of \$31 per person during the study period. Given mean arrondissement population of 20,644 and 103 arrondissements in the study area, this amounts to an aggregate gain of \$USD 65,916,292.

Putting together costs and benefits, our back-of-the-envelope calculation implies a cost-benefit ratio of 2.23. For each dollar earned from cashew cultivation, the ecological cost is about twice as large. This represents a lower bound since the cost estimate is entirely based on carbon and excludes other foregone ecosystem services.

## 6 Conclusion

This paper investigates the development-environment tradeoff from cultivating cash crops. The empirical setting is Benin, West Africa, during 2015-2021, a period when the cashew sector experienced a dramatic expansion. We use novel, remotely-sensed data on cashew plantations to study impacts of cultivation on local incomes and forest cover. While many studies have documented the development impacts of agriculture, and other have documented environmental impacts, few have studied them together in the same context.

We uncover a clear development-environment tradeoff: cashew cultivation raises local incomes, but at the expense of local forest cover. We document this tension across several research designs, ranging from cross-sectional comparisons on household survey data, to a more robust instrumental variable design with panel data using global wholesale cashew prices as an instrument. Our estimates imply that doubling land under cashew cultivation increases local economic activity by 15-30% and reduces forest cover by a similar amount, depending on the specification.

Our results have several broader implications. First, they provide a benchmark for the costs and benefits of cash crops, which is especially relevant as other West African countries look to reinvigorate the agricultural sector. Our cost-benefit analysis suggests that the ecological cost per dollar generated from cash crops is about \$2.23. Second, our findings underscore the amount of savings that can be realized through sustainable development policy. For example, through investments in agricultural productivity on existing land, there is less need to encroach into neighboring forestland, leading to gains from avoided deforestation ([Abman and Lundberg, 2024](#); [Abman et al., 2024](#)).

Our findings should be interpreted with certain caveats. In terms of external validity, the estimates pertain to a small country with a unique ethnic composition, land tenure system, and climate characteristics suitable for cashews. It is unclear whether our findings generalize to settings outside of West Africa. In terms of internal validity, our panel estimates rely on gridded GDP data, which has several conceptual and measurement limitations. For this reason, we complemented this data with alternative income proxies from surveys and nightlights. Lastly, since we lack survey data on cashew farmers, we are unable to verify whether local benefits accrue to cashew farmers specifically, or whether forest loss is driven by farm encroachment into adjacent forestland. Nevertheless, we provide indirect and suggestive evidence that this is the case.

## References

- Abman, Ryan and Clark Lundberg**, “Contracting, market access and deforestation,” *Journal of Development Economics*, 2024, p. 103269.
- , **Teevrat Garg, Yao Pan, and Saurabh Singhal**, “Agricultural Productivity and Deforestation,” *Working Paper*, 2024.
- Alix-Garcia, Jennifer M, Katharine RE Sims, and Patricia Yañez-Pagans**, “Only one tree from each seed? Environmental effectiveness and poverty alleviation in Mexico’s payments for ecosystem services program,” *American Economic Journal: Economic Policy*, 2015, 7 (4), 1–40.
- Anderson, Michael L**, “Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American statistical Association*, 2008, 103 (484), 1481–1495.
- Assunção, Juliano, Molly Lipscomb, Ahmed Mushfiq Mobarak, and Dimitri Szerman**, “Agricultural productivity and deforestation in Brazil,” *Climate Policy*, 2017.
- Bellemare, Marc F and Christopher B Barrett**, “An ordered Tobit model of market participation: Evidence from Kenya and Ethiopia,” *American journal of agricultural economics*, 2006, 88 (2), 324–337.
- Brewer, Julia, Ashley Larsen, and Frederik Noack**, “The land use consequences of rural to urban migration,” *American Journal of Agricultural Economics*, 2024, 106 (1), 177–205.
- Burgert, Clara, Josh Colston, Thea Roy, and Blake Zachary**, “Geographic Displacement Procedure and Georeferenced Data Release Policy for the Demographic and Health Surveys,” *DHS Spatial Analysis Reports No. 7*, 2013.

- Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli**, “Agricultural productivity and structural transformation: Evidence from Brazil,” *American Economic Review*, 2016, 106 (6), 1320–1365.
- Caunedo, Julieta and Namrata Kala**, “Mechanizing agriculture,” Technical Report, National Bureau of Economic Research 2021.
- Chen, Jiandong, Ming Gao, Shulei Cheng, Wenxuan Hou, Malin Song, Xin Liu, and Yu Liu**, “Global 1 km × 1 km gridded revised real gross domestic product and electricity consumption during 1992–2019 based on calibrated nighttime light data,” *Scientific Data*, 2022, 9 (1), 202.
- Conning, Jonathan and Christopher Udry**, “Rural financial markets in developing countries,” *Handbook of agricultural economics*, 2007, 3, 2857–2908.
- Dasgupta, Susmita, Benoit Laplante, Hua Wang, and David Wheeler**, “Confronting the environmental Kuznets curve,” *Journal of economic perspectives*, 2002, 16 (1), 147–168.
- Demographic and Health Surveys Program**, “Guide to DHS Statistics: DHS-8,” *DHS Manual*, 2023.
- Duguma, LA, Alagie Bah, PA Minang, Tesfaye Woldeyohanes, Malanding Jaiteh, E Sanneh, K Muthee, and M Badjie**, “Cashew: An emerging tree commodity in African drylands for livelihoods improvement and ecosystem restoration,” *Chapter 7: Tree Commodities and Resilient Green Economies in Africa*, 2021.
- Emerick, Kyle**, “Agricultural productivity and the sectoral reallocation of labor in rural India,” *Journal of Development Economics*, 2018, 135, 488–503.
- FAO**, “FRA 2020 Remote Sensing Survey,” *FAO Forestry Paper No. 186*, 2020.
- Gibson, John, Susan Olivia, and Geua Boe-Gibson**, “Night lights in economics: Sources and uses 1,” *Journal of Economic Surveys*, 2020, 34 (5), 955–980.

- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik instruments: What, when, why, and how,” *American Economic Review*, 2020, 110 (8), 2586–2624.
- Green, Rhys E, Stephen J Cornell, Jorn PW Scharlemann, and Andrew Balmford**, “Farming and the fate of wild nature,” *science*, 2005, 307 (5709), 550–555.
- Grossman, Gene M and Alan B Krueger**, “Economic growth and the environment,” *The quarterly journal of economics*, 1995, 110 (2), 353–377.
- Hansen, Matthew C, Peter V Potapov, Rebecca Moore, Matt Hancher, Svetlana A Turubanova, Alexandra Tyukavina, David Thau, Stephen V Stehman, Scott J Goetz, Thomas R Loveland et al.**, “High-resolution global maps of 21st-century forest cover change,” *science*, 2013, 342 (6160), 850–853.
- Henderson, J Vernon, Adam Storeygard, and David N Weil**, “Measuring economic growth from outer space,” *American economic review*, 2012, 102 (2), 994–1028.
- Hoffmann, Lars, Gebhard Günther, Dan Li, Olaf Stein, Xue Wu, Sabine Griessbach, Yi Heng, Paul Konopka, Rolf Müller, Bärbel Vogel et al.**, “From ERA-Interim to ERA5: the considerable impact of ECMWF’s next-generation reanalysis on Lagrangian transport simulations,” *Atmospheric Chemistry and Physics*, 2019, 19 (5), 3097–3124.
- Houssoukpèvi, Issiakou Alladé, Hervé Nonwègnon Sayimi Aholoukpè, Demayi Jorès Mauryo Fassinou, Murielle Nancy Rakotondrazafy, Guillaume Lucien Amadji, Lydie Chapuis-Lardy, and Tiphaine Chevallier**, “Biomass and soil carbon stocks of the main land use of the Allada Plateau (Southern Benin),” *Carbon Management*, 2022, 13 (1), 249–265.
- INDFC**, “Nuts and Dried Fruit Statistical Yearbook 2022/23,” *Annual Report*, 2023.

**Jayachandran, Seema, Joost De Laat, Eric F Lambin, Charlotte Y Stanton, Robin Audy, and Nancy E Thomas**, “Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation,” *Science*, 2017, 357 (6348), 267–273.

**Lin, Chenxi, Zhenong Jin, David Mulla, Rahul Ghosh, Kaiyu Guan, Vipin Kumar, and Yaping Cai**, “Toward large-scale mapping of tree crops with high-resolution satellite imagery and deep learning algorithms: A case study of olive orchards in Morocco,” *Remote Sensing*, 2021, 13 (9), 1740.

**Madhok, Raahil, Frederik Noack, Ahmed Mushfiq Mobarak, and Olivier Deschenes**, “Internal Migration and the Reorganization of Agriculture,” *Mimeo*, 2024.

**Ministere de l’Agriculture de l’Elevage et de la Peche**, “Plan Stratégique de Développement du Secteur Agricole (PSDSA) 2025 et Plan National d’Investissements Agricoles et de Sécurité Alimentaire et Nutritionnelle PNIASAN 2017 - 2021,” *Government Report*, 2017.

**Moscona, Jacob**, “Agricultural Development and Structural Change, Within and Across Countries,” *Mimeo*, 2019.

**Putra, Yoga Cahya and Arie Wahyu Wijayanto**, “Automatic detection and counting of oil palm trees using remote sensing and object-based deep learning,” *Remote Sensing Applications: Society and Environment*, 2023, 29, 100914.

**Ronneberger, Olaf, Philipp Fischer, and Thomas Brox**, “U-net: Convolutional networks for biomedical image segmentation,” in “Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18” Springer 2015, pp. 234–241.

**Stern, David I, Michael S Common, and Edward B Barbier**, “Economic growth and environmental degradation: the environmental Kuznets curve and sustainable development,” *World development*, 1996, 24 (7), 1151–1160.

**The World Bank**, “World Development Report 2008: Agriculture for Development,” *World Bank Report No. 41455*, 2007.

**Townshend, John, Matthew Hansen, Mark Carroll, Charlene DiMiceli, Robert Sohlberg, and Chengquan Huang**, “Annual global automated MODIS vegetation continuous fields (MOD44B) at 250 m spatial resolution for data years beginning day 65, 2000-2014, collection 5 percent tree cover, version 6.,” Technical Report, University of Maryland, College Park 2017.

**UNCTAD**, “Commodities at a glance: special issue on cashew nuts,” *UN Working Paper No. 14*, 2021.

**World Bank**, “Benin: Country Forest Note,” *Report No: AUS0001346*, 2020.

**Yin, Leikun, Rahul Ghosh, Chenxi Lin, David Hale, Christoph Weigl, James Obarowski, Junxiong Zhou, Jessica Till, Xiaowei Jia, Nanshan You et al.**, “Mapping smallholder cashew plantations to inform sustainable tree crop expansion in Benin,” *Remote Sensing of Environment*, 2023, 295, 113695.



## A Appendix Tables

Table A1: Identifying Variation

	$1 - R^2$	$\sigma_\epsilon$
	(1)	(2)
NTL: Arrondissement + Department-Year FE	0.011	0.054
GDP: Arrondissement + Department-Year FE	0.001	0.030

This table summarizes regressions of log GDP (first row) and log nightlights (second row) on arrondissement and department-year fixed effects. Column 1 reports  $1 - R^2$  i.e., the fraction of variation not explained by the fixed effect effects. Column 2 is the standard deviation of the residuals.

Table A2: Instrument Validity: Correlation of Cashew Coverage with Confounders

	(1)	(2)
Credit Union	0.005 (0.010)	0.018 (0.014)
Educated	-0.023*** (0.008)	-0.021** (0.011)
Owens Cattle	0.001 (0.001)	-0.003 (0.002)
% Agricultural Jobs	0.005 (0.011)	-0.010 (0.013)
% Professional Jobs	-0.006 (0.009)	0.008 (0.011)
Commune FEs	✓	
Department FEs		✓
Observations	1972	1972
$R^2$	0.565	0.276

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are a cross-section. The outcome is percent cashew coverage measured as % of grid cells with cashew plots in 2015. Explanatory variables are from the 2018 DHS survey, aggregated to the arrondissement. "Credit union" is the share of DHS households in an arrondissement with credit access. "Educated" is the share of households with an educated household head (secondary or higher). "% Agricultural jobs" and "% Professional jobs" is the share of household members with agricultural and professional jobs, respectively". Standard errors are robust to heterogeneity.

Table A3: DHS Results: Cashew Plantations and Household Wealth Score

	(1)	(2)	(3)
Near cashew (=1)	0.260** (0.123)		
Log(Proximity Cashew)		0.052 (0.040)	0.236** (0.094)
Log(Proximity Cashew) × Log(Proximity Nursery)			0.087** (0.036)
Log(Proximity Cashew) × Trainings			0.005 (0.003)
Household Controls	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓
Observations	2866	2866	2866
$R^2$	0.229	0.228	0.246

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are a household cross section. The outcome is a wealth score derived by inverse-covariance weighting ownership of various assets (see Section 4.1. “Near Cashew” indicates whether the household’s DHS cluster is below the median distance to the nearest cashew plot. “Proximity Cashew” and “Proximity Nursery” are inverse distance to the nearest cashew plot and seedling nursery, respectively. “Trainings” are the number of training events held in the arrondissement in 2018. All regressions include survey weights and controls for: household size, language, rain, temperature, lat/lon, and nightlights. Standard errors robust to heterogeneity.

Table A4: DHS Results: Cashew Plantations and Forest Quality

	(1)	(2)	(3)
Near cashew (=1)	-0.022*** (0.002)		
Log(Proximity Cashew)		-0.001* (0.001)	-0.010*** (0.001)
Log(Proximity Cashew) × Log(Proximity Nursery)			-0.002*** (0.001)
Log(Proximity Cashew) × Trainings			0.000*** (0.000)
Household Controls	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓
Observations	2866	2866	2866
$R^2$	0.985	0.984	0.985

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are a household cross section. The outcome is the log EVI in a small radius around each household (2km radius for rural and 10km for urban households). “Near Cashew” indicates whether the household’s DHS cluster is below the median distance to the nearest cashew plot. “Proximity Cashew” and “Proximity Nursery” are inverse distance to the nearest cashew plot and seedling nursery, respectively. “Trainings” are the number of training events held in the arrondissement in 2018. All regressions include survey weights and controls for: household size, language, rain, temperature, lat/lon, and nightlights. Standard errors robust to heterogeneity.

Table A5: Panel Estimates: Impacts of Marginal Plantation on GDP and Forests

Outcomes in logs	(1) NTL	(2) NTL	(3) GDP	(4) GDP	(5) Forest	(6) Forest
Cashew Density (per $km^2$ )	0.022 (0.019)	0.016 (0.028)	0.007*** (0.002)	0.008*** (0.003)	-0.020* (0.010)	-0.035** (0.015)
Cashew Density (per $km^2$ ) $\times$ Plants/Nursery		-0.002 (0.001)		-0.000 (0.000)		0.001* (0.001)
Cashew Density (per $km^2$ ) $\times$ Trainings		0.001 (0.004)				-0.001 (0.002)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department $\times$ Year FEs	✓	✓	✓	✓	✓	✓
Observations	412	309	206	206	412	309
$R^2$	0.946	0.967	0.998	0.998	0.876	0.899

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcomes are log nightlight intensity, log GDP per capita, and log share of land area under forest cover. "Cashew Density" is the number of 200m cashew plots per  $km^2$  in an arrondissement. "Plants/Nursery" is the mean number of seedlings (in hundreds) per nursery in the arrondissement at baseline. "Trainings" are the number of training events held in the arrondissement in 2018. All specifications include controls for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A6: Panel Estimates: High-Resolution Cashew Plantations, GDP, and Forests

Outcome in logs	(1) NTL	(2) GDP	(3) Forest	(4) NTL	(5) GDP	(6) Forest
Cashew Coverage (%)	0.404704 (0.637039)	0.228329** (0.101918)	-0.014050 (0.319084)			
Cashew Density (per $km^2$ )				0.000002 (0.000004)	0.000001** (0.000001)	-0.000000 (0.000002)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department $\times$ Year FEs	✓	✓	✓	✓	✓	✓
Observations	412	206	412	412	206	412
$R^2$	0.946	0.998	0.874	0.946	0.998	0.874

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcome in columns 1 and 4 is log nightlight intensity. The outcome in columns 2 and 5 is log GDP per capita. The outcome in columns 3 and 6 is log share of land area under forest cover. "Cashew coverage" is the fraction of 3m grid cells in an arrondissement with cashew trees. "Cashew Density" is the number of 3m resolution cashew trees per  $km^2$  in an arrondissement. All specifications include controls for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A7: First Stage Estimates

	(1)	(2)
CashewPrice $\times$ CashewShare	0.046*** (0.011)	0.149*** (0.043)
AgPrice $\times$ CashewShare		-0.124*** (0.042)
Controls	Yes	Yes
Outcome Mean	0.391	0.391
Arrondissement FEs	✓	✓
Department $\times$ Year FEs	✓	✓
KP (2006) F-Stat	16.74	12.16
Observations	412	412

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are at the arrondissement-year level. The outcome is cashew coverage, the fraction of 200m cells in an arrondissement growing cashews. "Cashew Price" is the global wholesale cashew price in year  $t$ . "CashewShare" is cashew coverage at baseline. "Ag. Price" is the global food price index in time  $t$ . All specifications include controls for rain, temperature, and drought intensity.

Table A8: Mechanisms: Wealth and Farm Size

	(1) Wealth Index	(2) Wealth Score	(3) Log Farm Size
Near cashew (=1)	0.247 (0.155)	0.180 (0.146)	0.006 (0.101)
Near Cashew (=1) $\times$ Land owner (=1)	0.206** (0.095)	0.142 (0.101)	0.114* (0.062)
Land owner (=1)	-0.575*** (0.066)	-0.218** (0.089)	3.360*** (0.052)
Household Controls	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes
Outcome Mean	3.004	-	2.574
Arrondissement FEs	✓	✓	✓
Observations	2866	2866	2806
$R^2$	0.533	0.233	0.839

Note: \*  $p < .1$ , \*\*  $p < .05$ , \*\*\*  $p < .01$ . Data are a household cross section. "Near Cashew" indicates whether the household's DHS cluster is below the median distance to the nearest cashew plot. "Landowner" indicates whether the household owns farmland. All regressions include survey weights and controls for: household size, language, rain, temperature, lat/lon, and nightlights. Standard errors robust to heterogeneity.

## B Appendix Figures

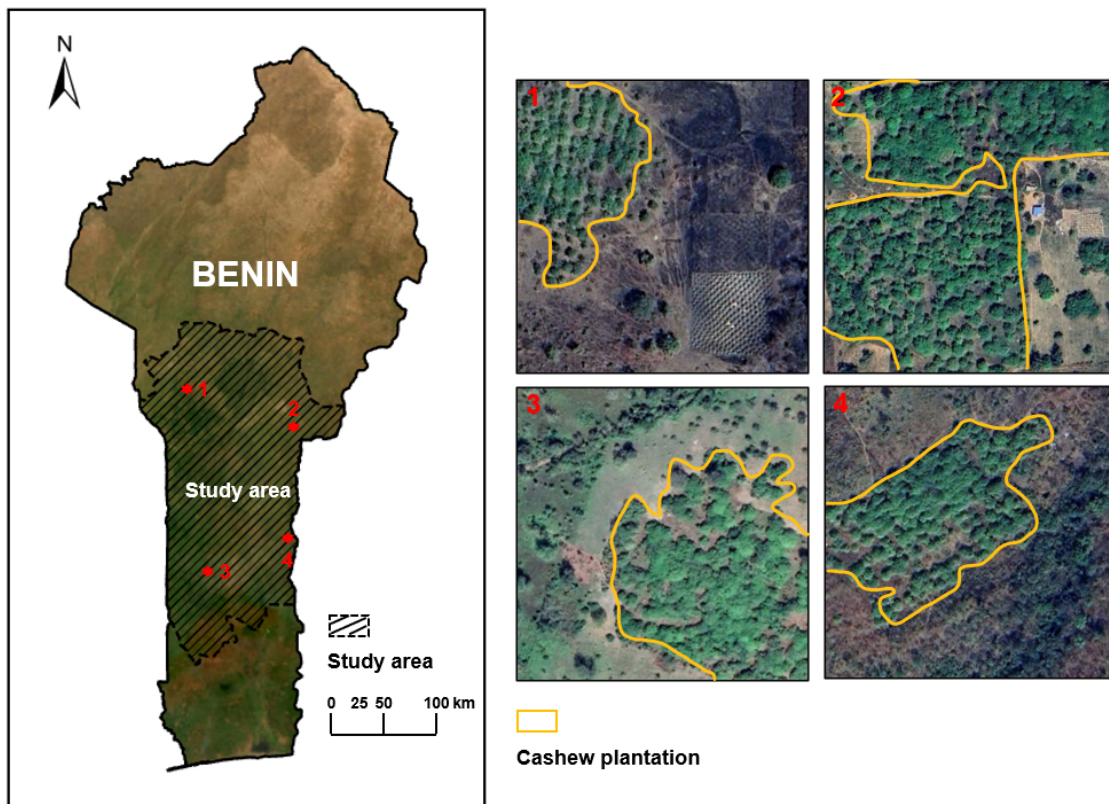


Figure B1: Cashew Plot Delineation

Note: Left panel shows the study area along with four example cashew plantations (red points). Right panel shows high resolution classification model output.

## C Forest Classification Exercise

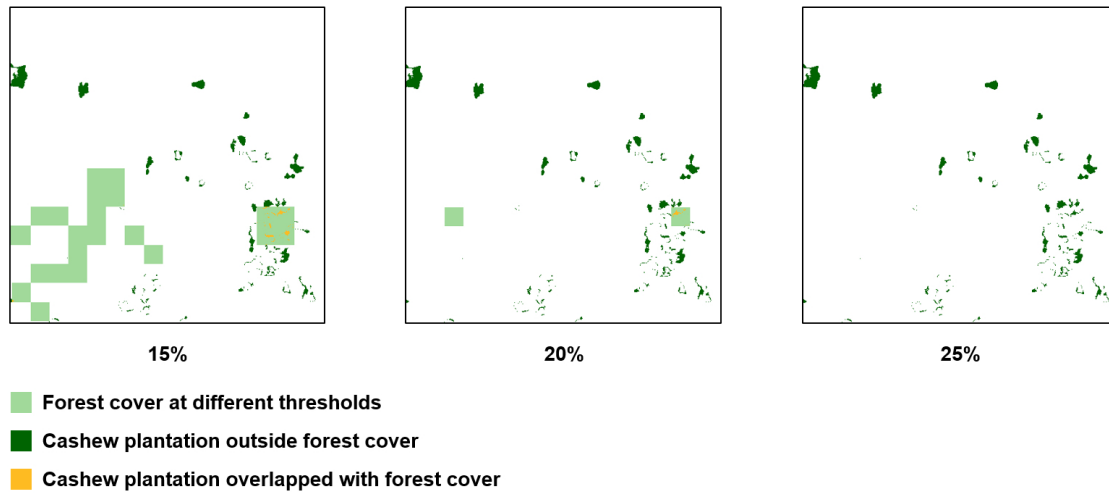


Figure C1: Forest cover at thresholds 15%, 20%, and 25%, and newly developed cashew plantation in 2021

Note: Figure shows new cashew plantations and forest locations in 2021. Panels A, B, and C define forest (light green) as pixels with forest cover above 15%, 20% and 25%, respectively. Dark green polygons are new cashew plantations (planted between 2020 and 2021) outside of forests. Yellow polygons are new plantations inside forests.

In this Appendix, we develop a procedure to determine whether new cashew plantations were classified as forests in the VCF satellite product for the corresponding year. We show that the vast majority of new cashew plantations are not classified as forest. This is important because, if the VCF satellite did classify cashew trees as forest, then our estimates of the forest cover impact of cashew agroforestry would be attenuated.

We conduct our validation exercise in three steps. First, we set a 15% forest cover threshold to classify a VCF pixel as forest. We do this because pixel values denote forest cover, and we need to define what constitutes forest. Second, we overlay our gridded cashew maps in year  $t$  on year  $t - 1$  and define the non-overlapping polygons as cashew expansions in year  $t$ . Third, we overlay these new cashew plantation polygons on the year  $t$  forest cover data product to determine whether they lie in a forest (based on the 15% threshold) or not. To ensure robustness of our exercise, we repeat this exercise for several stricter thresholds for what constitutes forest.

<b>Share of New Cashew Plantations Outside of Forests</b>			
<b>Year</b>	<b>15% Threshold</b>	<b>20% Threshold</b>	<b>25% Threshold</b>
2019	0.908	0.971	0.996
2020	0.950	0.997	0.999
2021	0.908	0.972	0.995

Table C9: Share of new cashew plantations outside of forests.

Note: Cell values are the share of new cashew plantations (planted between the corresponding row year and the prior year) outside of forests. VCF pixel values (percent forest) are classified as forest if they are above the threshold (columns).

Overall, over 90% of new cashew plantations are located outside of forests. For the purpose of visualization, Figure C1 maps cashew plantations inside and outside forests in a zoomed-in region across different forest thresholds in 2021. While this example shows some new cashew plantations inside forests, Table C9 shows the percent of new plantations in each year located outside of forests across the full study area. The results imply that the vast majority of VCF forest cover pixel values *exclude* cashew trees.