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

Raahil Madhok† University of Minnesota

Leikun Yin‡ University of Minnesota

Zhenong Jin§ University of Minnesota

May 17, 2025

Abstract

Cash crops can boost local economies but compete with natural forests for land. We quantify this tension in the context of cashew agroforestry in Benin, West Africa, by exploiting variation in cashew cultivation driven by heterogenous local responses to global cashew price volatility. In doing so, we develop the first high-resolution cashew maps for 2015-2021 using a deep learning model trained on data from field visits and pair it with newly released gridded GDP data. We find that expansion of cashew cultivation degrades forests without generating detectable income gains. These muted income effects reflect (i) concentrated economic benefits obscured in coarse GDP data, and (ii) farmers valuing second order income-smoothing benefits of cashews over first order income gains. Cost-benefit calculations show that each dollar earned from cashews incurs $14-$18 in ecological costs.

Keywords: Agriculture, development, deforestation, Benin

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

*This draft supersedes a previous version distributed as AgEconSearch Discussion Paper No. 347614. We thank Ryan Abman, Marc Bellemare, Emily Pakhtigian, Joakim Weill, and Weizhe Weng for helpful feedback. All errors are our own.

†Department of Applied Economics. Email: rmadhok@umn.edu

‡Department of Bioproducts and Biosystems Engineering. Email: yin00406@umn.edu

§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). While smallholder agriculture can transform rural livelihoods, boost local economies, and improve food security, 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 economic benefits and ecological costs of smallholder agriculture in developing countries. An extensive literature by economists, ecologists, and practitioners studies solutions for balancing this development-environment 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 have quantified the size of the tradeoff in the first place.

We address this gap by estimating the development-environment tradeoff in the context of cashew tree crops in Benin, West Africa. Our objective is to quantify income effects of expanding cashew cultivation as well as forest loss for each unit of expansion. We achieve this objective by developing the first high-resolution cashew maps for Benin using remote sensing, deep learning, and validation data from the field. We pair this novel data with a shift-share instrument for cashew cultivation based on fluctuations in global cashew price volatility. Benchmarking the development-environment tradeoff from agriculture is important for helping developing country governments understand the ecological effects of structural transformation and agricultural modernization.

The cashew sector in Benin offers an ideal study setting since cashews are a high-value crop, provide stable incomes, and are grown widely by smallholders across West Africa. Moreover, Benin is a top-10 global cashew producer, the third largest in West Africa, and has invested heavily in the sector over the past decade (Lin et al., 2021; Yin et al., 2023; Duguma et al., 2021). This smallholder-led growth allows us to examine how agricultural expansion impacts rural incomes and forest loss. While cashew production can increase by forestland conversion (extensive margin) or technology adoption on existing farmland (intensive margin), we focus primarily on the former.

To measure cashew cultivation over time, we built a remotely-sensed data product using image classification, deep learning, and field visits to classify cashew trees at 3m resolution for the years 2015, 2019, 2020, and 2021. Our novel gridded data product is the first of its kind and fills a key gap: most studies of tree crops focus on rubber and oil palm, which are grown on large commercial farms and can be detected by low-resolution satellites (Putra and Wijayanto, 2023), whereas cashew trees grow on small fields, preventing classification by satellites. Our image classification algorithm overcomes this issue, enabling one of the first studies of cashew farming in the applied economics literature[1].

To measure local incomes, we use the new (released February 2025) annual gridded GDP product by Rossi-Hansberg and Zhang (2025) (hereafter, RHZ), which provides modeled GDP estimates at 0.25° resolution globally. These are the most precise subnational GDP estimates available and we are among the first to pilot its use. Yet at 25°, roughly the size of Chicago, income changes at lower levels go undetected. To capture household-level effects, we pair the panel data with the 2018 Demographic and Health Survey (DHS), which surveys 3000 households in our study area.

Lastly, we measure time-varying forest cover using 200m resolution satellite data. To address concerns that cashew trees may be misclassified as forest, we conduct a validation exercise which shows that over 90% of cashew pixels are *not* classified as forest. This means that cashew encroachment into forest will be registered as a decrease in forest cover, without contamination from misclassifying cashew trees as afforestation.

The main econometric challenge in estimating the impact of cashew cultivation on incomes and forests is reverse causality: higher incomes may enable cashew expansion, and degraded forestland may be more likely to be converted to cashew plantations. We

[1] We wish to clarify that the development of our data product is the contribution of our companion paper (Yin et al., 2023), not this one. In this paper, we are simply consumers of our own gridded data product.

address this with a shift-share instrument, motivated by the idea that specialized and non-specialized farmers respond differently to price uncertainty. The "shift" consists of global cashew price volatility and the "share" is baseline cashew land share. This setup isolates local variation in land allocated to cashews generated by heterogeneous farmer responses to global price uncertainty. We support the exclusion restriction by controlling for the corresponding price *level* shift-share—in case price volatility and level are correlated—since high cashew prices can raise incomes of cashew sellers without affecting land use.

Our analysis uncovers an interesting and somewhat puzzling pattern: cashew cultivation in Benin sharply degrades forests without generating meaningful income gains, at least none observable in the satellite GDP data. A 10 percentage point (pp.) increase in land share under cashews causes a 2.7 pp. decline in forest cover. The same expansion of cashew farming increases local GDP per capita by 4.5%, though the estimate is imprecise. Sharp forest loss and weak income gains also appear under a variety of alternative price volatility instruments, with different fixed effects and time trends, and with alternative data sources. Dynamic estimates of the tradeoff also show that forest loss from cashew expansion is lasting, whereas economic benefits are not detected even after several years.

The remainder of the paper investigates mechanisms. Why do farmers grow cashews despite their apparent unprofitability? We find support for two explanations: (i) cashew farming *is* profitable, but income gains are highly concentrated and thus obscured in coarse GDP data; and (ii) farmers pursue *second-order* benefits—reduced income volatility—not first-order income gains. To test the first explanation, we use DHS household data and find that landowning households near cashew plantations are significantly wealthier than those farther away. Though cross-sectional, this points to income gains from cashew cultivation *within households* that may be missed in coarse GDP measures. We then turn to higher-resolution panel GDP data (0.1°) from Chen et al. (2022). Though only for two years and unsuitable for shift-share analysis due to insufficient variation, corresponding TWFE estimates reveal strong income gains. Lastly, to deploy our 2SLS strategy, we use the shift-share on a restricted RHZ sample of areas where cashews are widespread within GDP pixels and, once again, find significant income gains. These exercises suggest that muted effects in the main analysis reflect localized gains blurred in coarse GDP data.

To test the second explanation—that the absence of first order income benefits is because farmers are interested in second order benefits (reduced income volatility)—we use annualized volatility of monthly nightlights an outcome in the shift-share analysis. Unlike annual crops, this view envisions perennial cashew trees as capital assets: requiring initial investments but yielding returns over time that smooth income in uncertain environments. We find that cashew cultivation reduces income volatility across four measures of volatility, though only one measure is statistically significant. While this tempers strong conclusions, our estimates suggest that farmers prioritize income stability over immediate gains, helping explain muted income effects in our main analysis.

The paper concludes by incorporating our estimates into a cost-benefit analysis. We use the social cost of carbon to convert forest loss into dollars. A 10 pp. expansion in land share under cashews costs $USD 2 billion in terms of forest loss and generates between $USD 111-144 million in income gains depending on if we use our qualitatively positive or statistically precise GDP estimates. For every dollar earned from cashew cultivation, the ecological cost is between 14-18 times more. 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; Foster and Rosenzweig, 2003; Jayachandran, 2022). 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 villages and even households.

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 summarizes how we built gridded cashew maps and describes other data. Section 4 distills three stylized facts from the data and Section 5 outlines our instrumental variable strategy. Section 6 presents the main results. Section 7 investigates mechanisms and Section 8 concludes.

2 Background

Benin lies on the West African coast, bordered by Togo to the west and Nigeria to the east. The administrative structure features 12 departments, divided into 77 communes and 546 arrondissements. Arrondissements comprise several villages and form the main administrative unit for local governance. While we present some correlations and mechanisms at the household level (Section 4), most of our analysis is at the arrondissement level.

Agriculture is the backbone of Benin's economy, accounting for 30% of GDP and supporting the livelihoods of 70% of the population. Cashews are among the main cash crops,

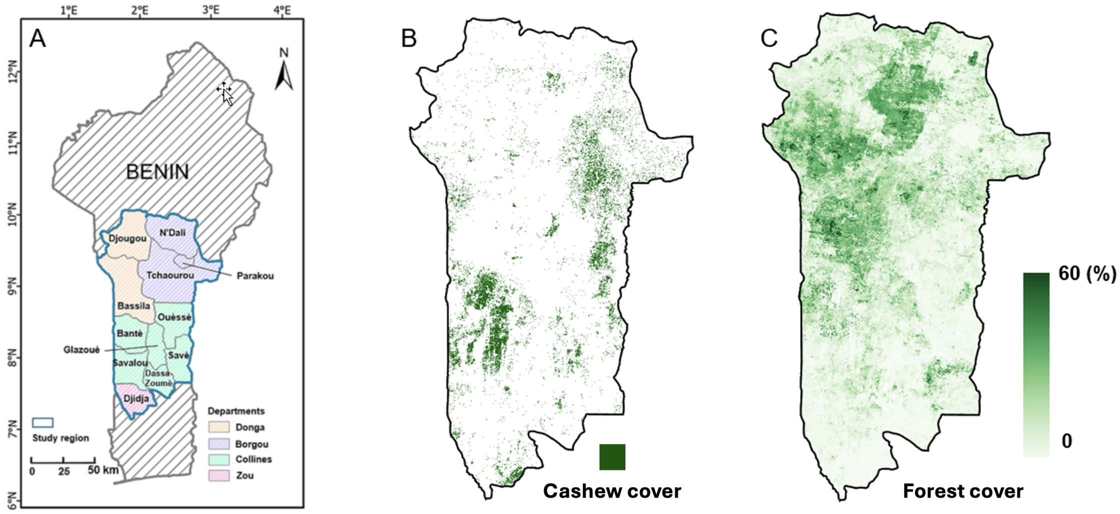

Figure 1: Study Area, Cashew Distribution, and Forest Cover
Note: Panel A shows the study area. Panel B shows the distribution of cashew plantations in the study area at 3m resolution. Panel C plots 2015 forest cover at 200m resolution. Cell values denote % forest cover.

in part due to the tropical savanna climate. Nearly 200,000 smallholders contribute to the industry, which generates about 15% of export revenue (Yin et al., 2023). Farmers first sell raw kernels to cooperatives or directly to intermediaries who, in turn, package, store, and sell to industrial processors. These processors handle shelling, roasting, grading, and preparation of the nuts for export (Tandjiékpon, 2010)[2]. Note that our results are not particular to this market structure, as the same supply chain is followed in other major West African cashew-producing countries (Mighty Earth, 2023).

Cashew cultivation is mainly practiced in Central Benin. Our study area was chosen to encompass the main cashew-growing communes, including all 6 communes in Collines department and 7 others across the departments of Donga, Borgou, and Zou (Figure 1A). The 13 communes in our study area comprises 103 arrondissements, collectively home to 25% of Benin's population. Figure 1B maps all cashew plantations within our study area (Section 3.1 for data details). Although plantations are spread across Central Benin, production is concentrated near the eastern and western borders.

[2]Over 95% of Benin cashews are exported (Tandjiékpon, 2010). Although our analysis focuses on the first step of the supply chain, much before export, an emerging economics literature studies links between agriculture, international trade, and deforestation (Carreira et al., 2024; Farrokhi et al., 2023; Hsiao, 2021).

In 2015, the first year of our study period, Benin implemented the BeninCaju program, an initiative to stimulate the cashew sector. The program offered subsidized seeds, farmer training, and subsidized credit (USDA and TechnoServe, 2015). Importantly, the program lasted throughout our full study period and was available to all farmers in our study region. Because of its ubiquitous coverage and lack of targeting, we do not expect the program to confound our research design. We discuss such threats to identification in more detail in Section 5.1.2.

Cashew agroforestry threatens local ecology since agriculture and forests compete for land. 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). This tension can also be observed locally: Figure 1C shows that cashew plantations are situated in 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 farm labor and capital investment. 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

We developed the first high-resolution cashew maps for Benin using remote sensing, deep learning, and validation data from the field. We complement this with newly released gridded GDP data. Forest cover is also measured with satellite data. The final panel is unbalanced and spans 2015, 2019, 2020, and 2021.

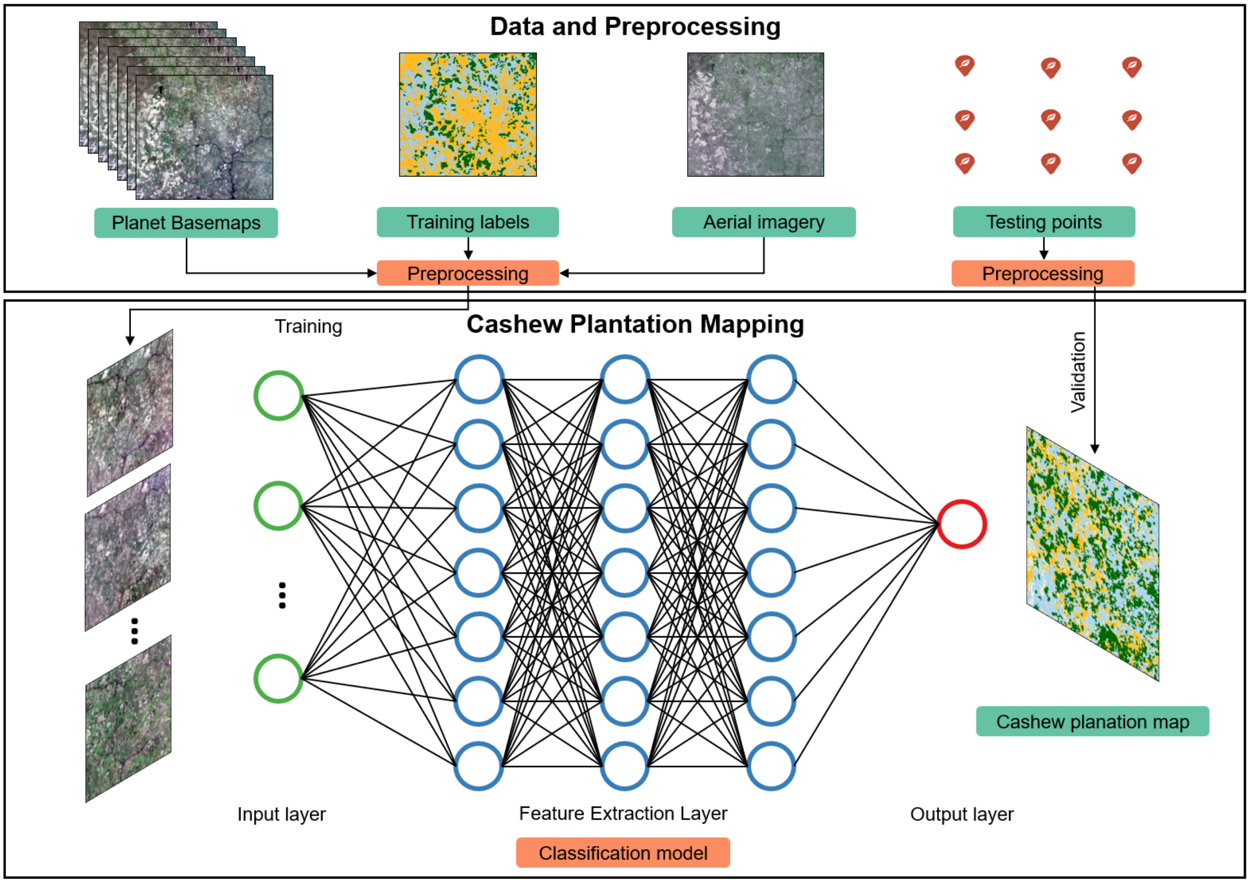

Figure 2: Cashew Image Classification Pipeline

Note: Overview of data processing pipeline along with the methods employed and the maps generated. Detailed technical description provided in our companion paper (Yin et al., 2023).

3.1 Building A Novel Geo-coded Dataset on Cashew Cultivation

Most studies of tree crops have been restricted to commercial crops on large farms (palm oil, rubber) since these can be mapped by low-resolution satellites (Putra and Wijayanto, 2023). In contrast, cashews have small crowns (<5m) and are grown on small fields, leading to a dearth of satellite data on their distribution and, therefore, a knowledge gap about their environmental and economic impacts.

We overcome this data gap by developing the first remotely-sensed cashew maps in Benin for the years 2015, 2019, 2020, and 2021 using cutting-edge image classification techniques paired with field data for ground-truthing. Our algorithmic design details and final data product are the contribution of our companion paper (Yin et al., 2023). In

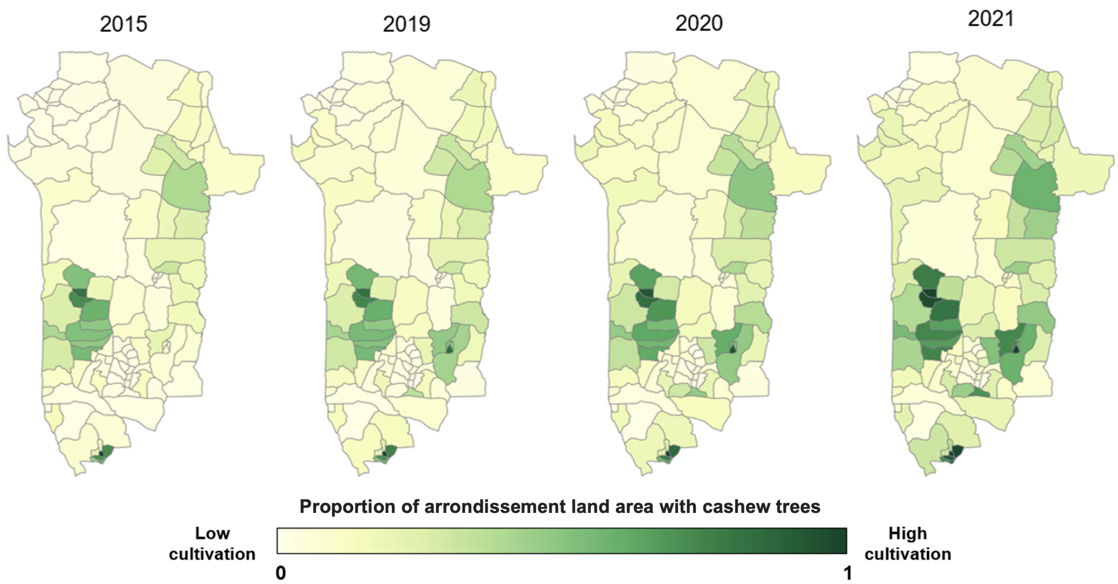

Figure 3: Spatial distribution of cashew coverage (%)

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.

the present paper, we are simply consumers of our own data. To avoid repeating details, we illustrate our data construction procedure in Figure 2 and summarize key steps below.

To build our own gridded cashew data product, we begin by downloading Planet Basemaps for 2019-2021, which provide high-resolution (3 meter) images of Earth. Base data for 2015 are from a Benin government aerial imaging exercise, as Basemaps is unavailable before 2019. Next, we trained a neural network model on validation labels drawn from site visits carried out by our local partner, TechnoServe, to "learn" whether pixels have cashew trees or not. Lastly, we applied the trained model to the Basemaps data for our study area to generate the final gridded cashew data product. To assess accuracy, field teams recorded land cover types in 1,400 validation sites. About 85% were correctly classified by our model, representing top-tier performance for smallholder tree crop classifications. Figure B1 shows classification output for four example cashew plantations. More technical details can be found in our companion paper (Yin et al., 2023).

For the analysis, we aggregate to the fraction of grid cells in an arrondissement growing cashews. This preserves the smallholder nature of cashew farming while aligning with the coarser resolution of GDP pixels (Section 3.3)[3]. Figure 3 plots cashew cultivation at the arrondissement level across the study period: an increase in the density of cashew plots is observed as green areas become darker over time. We also observe arrondissements with no cashews in 2015 come under cultivation in later periods (southeast region).

Our mapping procedure is designed to scale across time and space. It can be easily extended past 2019, when Planet Basemaps are available, and before that if local imaging products are available. The procedure can also be applied in other regions that grow cashews, such as West Africa, South Asia, and Southeast Asia. The main caveat is that training labels must be available from field visits or from another source.

Cashew Prices We obtain country-month level cashew prices from FAOSTAT to build a global price volatility instrument for local cashew production (Section 5). Country-year data on cashew production (tonnes) are also obtained from FAOSTAT and used as weights in the aggregation from country to global level. In a robustness check, we use alternative prices from the International Nut and Dried Food Council (INDFC, 2023)[4]

3.2 Cross-Sectional Household Data

While the majority of our analysis uses satellite data, we establish stylized facts (Section 4) and study mechanisms (Section 7.1) using household data from the 2018 DHS survey. The survey is nationally representative and covers 14,156 households, of which 20% are in our study area. Household wealth is measured on a scale from 1 (poor) to 5 (rich). We measure household "exposure" to cashew cultivation by the euclidean distance from the centroid of their sampling cluster to the nearest plantation[5]. To estimate associations

[3]Moreover, In high-resolution (1-30m) crop mapping, small fields can result in mixed pixels that contain multiple land cover types, leading to classification errors. When data from these pixels are aggregated over larger areas, mixed pixel errors tend to average out, resulting in more accurate crop area estimates (Ozdogan and Woodcock, 2006; Husak et al., 2008).

[4]Yearly prices are computed by dividing global supply value ($USD billions) by production volume (metric tons) reported by INDFC. We convert to 2015 prices to account for inflation.

[5]For computational simplicity, we aggregating the 3m cashew tree rasters to 200m "plantations".

between cashew cultivation and forest cover, we use the Enhanced Vegetation Index (EVI) reported by DHS for each survey cluster. EVI, which is a commonly used measure of forest cover, is reported as an average within a 2km and 10km radius for rural and urban survey clusters, respectively. At these small scales, a negative correlation between cashew cultivation and EVI more likely indicates localized forest clearing for agriculture.

3.3 Time-Varying Satellite Data

Gridded GDP Measuring local agricultural development over time requires a dynamic wealth measure at high resolution. We use new (released February 2025) gridded GDP data developed by RHZ (Rossi-Hansberg and Zhang, 2025), which predicts GDP per capita at 0.25° resolution using a random forest model trained on subnational GDP shares and other predictors like nightlights, population, and CO_2 emissions. Final predictions are rescaled to match national or state GDP totals. For robustness, we also use nightlight intensity, which is a strong proxy for local GDP (Henderson et al., 2012).

Despite enabling measurement of local income in developing countries, gridded GDP suffers at least two shortcomings. The first is conceptual: GDP is an inherently aggregate object, reflecting gross value added from all production or consumption activities in an area. RHZ assume we can measure the contribution of a handful of individuals to total GDP based on how their economic production is picked up by nightlights and other proxies. 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 in rural areas where cash injections from cashew farming may go undetected from outer space. We partially address this with our instrumental variable (IV) design, as orthogonality between the IV and luminosity error allow us to circumvent issues of non-classical error in our 2SLS estimates.

Forest Cover Forest cover is from the Vegetation Continuous Fields (VCF) product (Townshend et al., 2017), which measures percent forest cover at 200m resolution. Our main outcome is forestland share, which is the weighted sum of pixel values in each arrondissement-year, with weights equal to pixel area, divided by arrondissement land area.

The main alternative to VCF is Hansen et al. (2013), which measures forest loss/gain at 30m. We do not use this data because its definition of forest includes "plantations in sub-tropical and tropical ecozones."[6]. This means that replacing natural forest with agroforestry is (mis)recorded as zero forest change, leading to severely attenuated estimates of forest loss from cashew cultivation[7]. VCF avoids this pitfall: as we show in Section 4.1, the percent forest cover values in VCF exclude cashew trees 90-99% of the time.

Covariates We include two sets of covariates in the TWFE and 2SLS analyses: weather and agricultural input use. Weather includes temperature, rain, and drought intensity, which control for climatic factors that co-determine agricultural productivity and economic output. Gridded annual temperature (°C) and rainfall (mm) are from the ERA5 product at 0.125° resolution (Hoffmann et al., 2019). Drought intensity is from the gridded (0.5° resolution) Standardized Precipitation Evapotranspiration Index (SPEI), which measures the difference between potential evapotranspiration and precipitation.

Agricultural input use is measured in kg/ha using the PEST-CHEMGRIDS data product at 0.1° resolution (Maggi et al., 2019). This variable help partial out intensive margin responses and allows us to isolate the impact of cashew production through extensive margin changes in land under cashews. Our measure of input use is the mean application rate of the 20 most common active chemical ingredients applied to fruit crops, which includes the cashew apple[8]. For both weather and agrochemical covariates, we extract the mean over cells within arrondissements for each year.

[6]Quote from `https://data.globalforestwatch.org/documents/gfw::tree-cover-loss/explore`.
[7]This misclassification also arises in Dynamic World (Brown et al., 2022), an alternative forest dataset.
[8]PEST-CHEMGRIDS is for 2015, 2020, and 2025. We linearly interpolate values for 2019 and 2021.

Table 1: Summary Statistics

	Observations	Mean	Std. Dev.
Panel A: DHS (2018)			
Dist. to Nearest Plantation (km)	2866	4.26	11.17
Wealth Index (scale 1-5)	2866	3.00	1.28
Panel B: Panel (2015-2021)			
Cashew Tree Coverage (%)	412	0.13	0.15
Cashew Density (per km^2)	412	7.04	4.57
GDP Per Capita (USD)	412	1172.44	1546.18
Forest (share of land area)	412	0.07	0.03
Agrochemicals (kg/ha)	412	0.41	0.04

Note: Panel A summarizes household variables from the 2018 DHS survey. Panel B summarizes the panel data at the arrondissement-year level. "Cashew plantation coverage" is the fraction of 200m grid cells in an arrondisseemnt with cashews. "Cashew tree coverage" is computed in the same way using 3m grid cells.

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, "Observations" is the number of households surveyed in the study region where cashew trees were mapped. The typical household lives about 4km from the nearest cashew plantation. The standard deviation is nearly three times the mean, indicating substantial variation in cashew exposure across space. The average household has a wealth score of 3 out of 5. In Panel B, data are at the arrondissement-year level. 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 1,170, which matches official national statistics. Lastly, the typical arrondissement has about 7% forest cover.

4 Three Stylized Facts

Having described the data, we next present three stylized facts about agriculture and deforestation in Benin. The first fact verifies that cashew trees are not misclassified as forests. The second and third establish correlations between cashew cultivation, wealth, and forest cover. These insights motivate our empirical strategy in Section 5.

4.1 Fact I: Cashew trees are not classified as forests

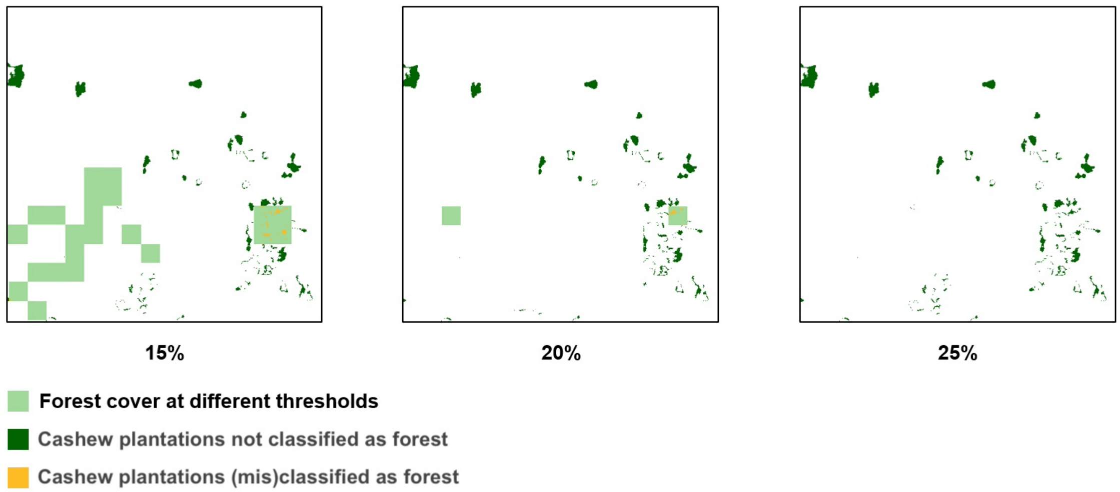

Figure 4: Validation of VCF Forest Cover Data

Note: 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 between 2020 and 2021 that are (correctly) not classified as forest by VCF. Yellow polygons are new plantations that are (incorrectly) classified as forest.

The first fact is that new cashew plantations are not measured as afforestation by our forest cover data. If they were, then the impact of cashew cultivation on forest cover would be attenuated since declines in forest cover from agroforestry encroachment would be offset by misclassifying new plantations as forest gain.

To establish this fact, we first define VCF pixels with forest cover $> 15\%$ as forest. Second, we overlay our cashew maps in year t on year $t-1$ and define non-overlapping polygons as new plantations or plantation expansions in year t. Third, we overlay these on year t forest pixels based on the 15% threshold. If any cashew expansions fall inside forest pixels, then the corresponding percent forest cover value from VCF includes both natural forests and cashew trees, which threatens our research design.

Figure 4 visualizes our procedure. For transparency, we show a part of the study area where some cashew plantations are correctly not classified as forest (dark green),

Table 2: Correlation: Cashew Proximity, Household Wealth, and Forest Cover

	(1) Wealth Index	(2) Log EVI
Near cashew (=1)	0.343**	-0.022***
	(0.145)	(0.002)
Household Controls	Yes	Yes
Geography Controls	Yes	Yes
Outcome Mean	3.004	8.118
Arrondissement FEs	✓	✓
Observations	2866	2866
R^2	0.512	0.985

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are a household cross section. Column 1 is household wealth on a scale of 1-5. Column 2 is log of EVI in a small radius around each household (2km radius for rural and 10km for urban). "Near Cashew" indicates whether households' DHS cluster is less than the median distance to the nearest cashew plantation. Regressions include survey weights and controls for household size, language, rain, temperature, lat/lon, and nightlights. Standard errors are heteroskedasticity-robust.

and where some plantations are incorrectly classified as forest (yellow). The VCF forest cover percentage in light green pixels with yellow polygons will therefore mistakenly include cashew trees. Yet when characterizing misclassification in the full study area, 90% of cashew plantations *are not* classified as forest by VCF (Table A1)[9]. This implies that cashew encroachment into forests will be measured as a decrease in forest cover by VCF data, without contamination from misclassifying cashew trees as afforestation.

4.2 Fact II: Cashews, wealth, and forests are correlated *across space*

The second fact is that households near cashew plantations are wealthier, yet are surrounded by more degraded forests. We establish this using DHS data by estimating:

$$Y_{iad} = \delta \cdot NearCashew_{iad} + \Gamma X'_{iad} + \alpha_a + \epsilon_{iad} \quad (1)$$

[9] 99% of VCF pixels do not include cashew trees using a 25% threshold for what constitutes a forest.

Table 3: TWFE: Cashew Coverage, Local Income, and Forest Cover

	(1) Log NTL	(2) Log GDP	(3) Forest Share
Cashew Share	0.376**	0.062	-0.031***
	(0.172)	(0.077)	(0.011)
Controls	Yes	Yes	Yes
Data Source	✓	✓	✓
Arrondissement FEs	✓	✓	✓
Department × Year FEs	412	412	412
Observations	0.977	0.857	0.882

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. Outcomes are transformed by $log(x + 0.01)$ to account for zero values. Column 1 is nightlight radiance, column 2 is GDP per capita from Rossi-Hansberg and Zhang (2025), and column 3 is forest share of land area. "Cashew share" is the fraction of arrondissement land area under cashews. All specifications control for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

where i, a, and d index households, arrondissements, and departments, respectively, and Y_{iad} is wealth or forest quality. $NearCashew_{iad}$ equals one if household i lives less than the median distance to the nearest cashew plantation, and zero otherwise. X'_{iad} is a vector of geography and household covariates[10]. The arrondissement fixed effect, α_a, absorbs time-invariant differences across arrondissements, leaving δ to be estimated off of comparisons across households *within the same arrondissement*.

Household exposure to cashew cultivation is positively correlated with wealth (Table 2). Living near a plantation is associated with 11% (=0.343/3.004) higher wealth (column 1). However, these "cashew-exposed" households also experience more environmental degradation, with 2.2% less surrounding forest cover compared to non-exposed households (column 2). Of course, these patterns are correlational: omitted variables, reverse causality, or unobserved dynamic effects could drive the results. The goal here is simply to motivate the main analysis (Section 5).

4.3 Fact III: Cashews, GDP, and forest cover are correlated *over time*

The third fact is that more cashew agroforestry *over time* also leads to higher local incomes and lower forest cover. We establish this by comparing GDP, nightlights, and forest cover *within* arrondissements at different cultivation levels, controlling for time fixed effects:

$$Y_{adt} = \beta \cdot C_{adt} + \Gamma X'_{adt} + \alpha_a + \gamma_{dt} + \epsilon_{adt} \tag{2}$$

where a, d, and t index arrondissements, departments, and years, respectively. Y_{adt} is either log nightlights, log GDP or forest cover. C_{adt} is the share of land area under cashew cultivation. X'_{adt} is a vector of covariates including temperature, rainfall, and drought intensity. Arrondissement fixed effects, α_a, absorb time-invariant differences between arrondissements. Department-by-year fixed effects, γ_{dt}, account for department-specific factors that change over time, such as regional agricultural policy or growth trajectories.

Yearly economic benefits from cashew cultivation are associated with forest degradation (Table 3). A 10 pp. cashew expansion is associated with a 3.8% rise in nightlight radiance (column 1) and increases GDP by 0.6%, but the effect is imprecise (column 2). These economic gains are accompanied by a 0.3 pp. decline in forest cover (column 3).

While TWFE improve on cross-sectional comparisons (Fact II), the results remain correlational due to potential reverse causality: income may influence selection into cashew cultivation. Although this concern may be limited given widespread seed subsidies and formal credit access throughout our study region (USDA and TechnoServe, 2015), we nonetheless interpret these findings with caution and proceed to a 2SLS strategy next.

[10]Covariates include rain, temperature, lat/lon, language, nightlights, and household size.

5 Research Design

We exploit *global* cashew price volatility to identify *local* development-environment trade-offs from cashew cultivation in Benin. While local cashew cultivation itself may be endogenous to income and forest cover, global cashew price volatility (i) strongly affects land allocated to cashews, especially in places already specialized in cashews, and (ii) as we argue, only affects local income and forest cover through the extensive margin (land), conditional on price levels and intensive margin (input) responses. Our research design leverages these features to construct a volatility-based IV for cashew cultivation.

5.1 Empirical Setup

5.1.1 Measurement

The target parameter for our analysis is given by β in Equation 2, the impact on incomes or forests from a marginal change in land share under cashew cultivation. To identify this parameter, we construct a shift-share instrument, z_{acdt}, for cashew cultivation that combines global cashew price volatility (the shift) with measures of local exposure to price uncertainty (the share). When combined, the shift and share yield an instrument that isolates the plausibly exogeneous component of land allocation decisions generated by heterogeneous local responses to global price uncertainty[11].

The shift-share IV consists of two components (Goldsmith-Pinkham et al., 2020). The global "shift" is measured by the variance, V_{it}, of monthly cashew prices for country i in year t, and then taking a weighted average of these variances using country output shares as weights. The local "share" is measured as baseline arrondissement land share under

[11] We consider short-term responses to price volatility since we only have a four-year unbalanced panel. The literature on dynamic models shows that this type of variation yields a lower bound estimate of the elasticity of land use changes to prices (Scott, 2014; Araujo et al., 2020).

cashews, $C_{ad(t=2015)}$. Letting Q_{it} be output, global price volatility and the shift-share are:

$$\text{[Volatility]} \qquad \sigma_t^{\text{global}} = \sum_i \frac{Q_{it}}{\sum_i Q_{it}} \cdot V_{it}$$

$$\text{[Shift-Share IV]} \qquad z_{adt} = \underbrace{log(\sigma_t^{\text{global}})}_{\text{shift}} \quad \times \quad \underbrace{C_{ad(t=2015)}}_{\text{share}} \tag{3}$$

Importantly, our definition of z_{adt} is grounded in economic theory (see Appendix A.1 for details). Classical economic theory states that price volatility discourages cultivation of risky crops (Sandmo, 1971). A contrasting view from developing countries is that specialized farmers facing limited alternatives and market imperfections may expand cultivation under price risk as a form of self-insurance (Barrett, 1996). This motivates our use of baseline cultivation, which proxies for specialization, as the share component of the shift-share instrument. Whether the self-insurance hypothesis holds is revealed by the first stage. We next discuss instrument validity, and then define estimating equations.

5.1.2 Instrument Validity

Our IV is valid if (i) it strongly predicts the local land share allocated to cashews, and (ii) it meets the exclusion restriction. The first criterion is tested via the first stage (Equation 5) below. The exclusion restriction is that z_{adt} affects local GDP and forest cover only by influencing the extent of local cashew production. While this is fundamentally untestable, we discuss the main threats to identification, and how we address them, here.

The main threat to identification is that price volatility (second moment) and price level (first moment) may be correlated. When price levels are high, local GDP can increase as farmers sell existing cashews at higher prices without converting marginal land, thus violating the exclusion restriction. We address this by controlling for $P_{adt} := log(p_t^{\text{global}}) \times C_{ad(t=2015)}$ in all our regressions, where p_t^{global} is measured similarly to σ_t^{global}, but using average monthly cashew price level in year t of country i instead of price variance. Con-

Table 4: Test for Pre-Trends

Outcomes are for time $t-1$	(1) Log NTL	(2) Log GDP	(3) Forest
Shift-share (time t)	0.043	0.019	-0.005
	(0.040)	(0.024)	(0.003)
Controls	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓
Department × Year FEs	✓	✓	✓
Observations	309	309	309
R^2	0.984	0.887	0.892

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. Outcomes are lagged by one period. "Shift-share" is the interaction of global cashew price variance with baseline cashew coverage. Controls include price level interacted with baseline cashew share, agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

trolling for P_{adt} leaves identification to rely on heterogeneous exposure to shocks in global price *uncertainty*, conditional on first order responses to price *levels*.

A related concern is that land share under cashews, the endogenous variable to be instrumented by z_{adt}, is an extensive margin object. If farmers react to price uncertainty through the intensive margin, i.e., by changing input use, and this affects profits in turn, then local incomes can again be affected without changing land share. Contextually, this concern is minimal because raw cashew nuts are produced "virtually without chemical inputs" (Tandjiékpon, 2010). As a safeguard, we include the application rate (kg/ha) of agrochemicals as a control in all our regressions.

The third concern is that the share, $C_{ad(t=2015)}$, is endogenous if cashew-specialized areas systematically differ from non-specialized areas. Although *fixed* differences are absorbed by arrondissement fixed effects, baseline cultivation may predict differential *changes* in outcomes. This only violates the exclusion restriction if $C_{ad(t=2015)}$ is correlated with the shift, σ_t^{global}, otherwise z_{adt} would not pick up the differential trend. The recent literature formalizes this idea by showing that shift-share IVs with endogenous shares are valid as long as the shock is as-good-as-random (Borusyak et al., 2022). We formally test this by regressing time $t-1$ outcomes, $Y_{ad(t-1)}$, on the time t instrument, z_{adt}, as follows:

$$Y_{ad(t-1)} = \theta \cdot z_{adt} + \Gamma X'_{adt} + \alpha_a + \gamma_{dt} + \epsilon_{adt} \tag{4}$$

where θ is a balance coefficient. If z_{adt} is as-good-as-randomly assigned, then $\theta = 0$ which implies that the shock is unrelated to pre-existing levels of income or forest cover. X'_{adt} is a covariate vector including price level and input use, as described above, as well as rain, temperature, and drought intensity. Indeed, we find $\theta = 0$ for all outcomes (Table 4), suggesting that price volatility does not systematically affect arrondissements with certain pre-existing outcome levels.

5.2 Estimating Equations

Having established validity of the instrument, we specify the first stage equation as:

$$C_{adt} = \pi \cdot z_{adt} + \delta P_{adt} + \zeta A_{adt} + \mathbf{X}'_{adt}\Sigma + \alpha_a + \gamma_{dt} + \epsilon_{adt} \tag{5}$$

As before, a, d, and t index arrondissements, departments, and years. z_{adt} is the shift-share IV described above. $P_{adt} := log(p_t^{\text{global}}) \times C_{ad(t=2015)}$ is the interaction of global average cashew prices with baseline cashew cultivation, which controls for farmers' responses to price changes. A_{adt} is the application rate of agrochemicals, which controls for intensive margin responses to price uncertainty. Including P_{adt} and A_{adt} are key to mitigating the main threats to the exclusion restriction (Section 5.1.2). Remaining terms are the same as Equation 2. The first stage coefficient π captures extensive margin variation in cashew cultivation that is plausibly orthogonal to local agricultural incentives.

The corresponding second stage equation can be written as:

$$Y_{adt} = \beta \cdot \widehat{C}_{adt} + \delta P_{adt} + \zeta A_{adt} + \mathbf{X}'_{adt}\Sigma + \alpha_a + \gamma_{dt} + \epsilon_{adt} \tag{6}$$

where Y_{adt} measures economic outcomes or forest cover. We use log GDP per capita from RHZ as the key economic outcome. β is the coefficient of interest: when GDP is the outcome, $\beta > 0$ indicates that expanding cashew cultivation increases local incomes. When forest cover is the outcome, $\beta < 0$ indicates that cashew expansion is at the expense of forests. Since cashew specialization partly depends on geography, $\mathbf{X'}_{adt}$ includes temperature, rain, and drought intensity to absorb relationships between climatic suitability and outcomes. Arrondissement fixed effects, α_a, absorb time-invariant differences between arrondissements. Department-year fixed effects, γ_{dt}, absorb regional trends.

β is thus identified by changes in land allocated to cashews generated by farmers' heterogeneous responses to cashew price volatility over time within an arrondissement, conditional on their response to changes in price levels, and holding fixed any average differences between arrondissements and any regional climate dynamics. The main identification assumption is that the propagation of global cashew price volatility across cashew-specialized and non-specialized arrondissements is independent of potential outcomes. We provided support for this assumption in the previous section (Section 5.1.2).

6 Results

This section presents formal evidence on the development-environment tradeoff from agroforestry in Benin. Our IV estimates show that cashew cultivation degrades natural forests without generating meaningful gains in local incomes. We explore the reasons for muted income effects in Section 7.

6.1 Main Findings

First Stage Estimates First stage estimates of Equation 5 are reported in Table A2. The outcome is C_{adt}, the share of land area under cashew cultivation. In columns 1-3, z_{adt} is constructed using output-weighted price variance (Section 5.1.1). Column 2 controls

Table 5: The Development-Environment Tradeoff from Cashew Cultivation

	(1) Log NTL	(2) Log GDP	(3) Forest
CashewShare	0.258	0.448	-0.274***
	(0.965)	(0.389)	(0.070)
Log Price × CashewShare	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓
Department × Year FEs	✓	✓	✓
KP (2006) F-Stat	19.84	19.84	19.84
Observations	412	412	412

Note: * $p < .1$, ** $p < .05$, *** $p <.01$. Data are at the arrondissement-year level. Outcomes are transformed by $log(x + 0.01)$ to account for zero values. Column 1 is nightlights, column 2 is GDP per capita, and column 3 is forest share of land area. "CashewShare" is the fraction of grid cells with cashew plantations, instrumented with the interaction of global cashew price variance with baseline cashew coverage. All regressions control for price level interacted with baseline cashew coverage, agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

for price level and column 3, our preferred specification, also controls for input use. In columns 4-6, z_{adt} is constructed using alternative volatility measures. The IV strongly predicts cashew cultivation in all specifications. The point estimate in column 3 implies that a 10% rise in price volatility leads farmers in cashew-exposed arrondissements to expand cultivation by 0.6 pp.[12]. The F-statistic is well above rule-of-thumb levels.

Second Stage Estimates Table 5 reports IV estimates of Equation 6 for economic activity and forest cover. The corresponding OLS estimates are reported in Section 4.3. The outcome variables in columns 1-3 are log nightlights, log GDP per capita, and forest share of land area, respectively. Standard errors are clustered by arrondissement.

Estimates of β imply that cashew cultivation sharply degrades local forests (column 3) without generating discernible income effects (columns 1-2). Although the point estimates in columns 1 and 2 are positive, we cannot interpret these economic gains as

[12] Appendix A.1 discusses the literature on positive associations between price volatility and production.

meaningful since statistical precision is low. In contrast, $\beta < 0$ and is statistically significant for forest cover in column 3. The point estimate implies that a 10 pp. increase in land share under cashew cultivation reduces forest cover by 2.7 pp. Given our supporting evidence for instrument validity in Section 5.1.2, we interpret these estimates as causal evidence that expanding cashew cultivation comes at the cost of forest loss, with little to no compensating gains in local livelihoods.

One reason that income gains are not statistically precise is that GDP data is relatively coarse (25km×25km), obscuring highly localized changes, unlike the much finer 200m forest cover data. Another reason is that farmers may be interested in second order benefits, i.e., income smoothing opportunities offered by cashew trees, rather than first order income gains. We provide modest evidence for both explanations in Section 7.1.

Sensitivity: Alternative Instruments To demonstrate robustness to alternative instruments, Table A3 estimates Equation 6 with other measures of global cashew price volatility. Instead of output-weighted price *variance*, we construct z_{adt} using standard deviations (columns 1, 4, and 7), the coefficient of variation (columns 2, 5, and 8), and 6-month rolling price variance (columns 3, 6 and 9). We continue to find statistically insignificant impacts of cashew cultivation on local economic activity (columns 1-6). The impact on forest share (columns 7-9) remains strongly negative, similar to the main estimate, and is statistically significant across all price volatility measures.

6.2 Dynamic Estimates

Our baseline estimates characterize the development-environment tradeoff from cashew expansion within the year. This overlooks important dynamic processes, since cashew trees take at least two years to bear fruit. We estimate these dynamics using cumulative lag models with up to two lags, which can mean several years since the first two periods of our sample are four years apart. To operationalize the method, cashew coverage is predicted from the first stage (Equation 5), after which lags are taken and used as independent variables in the second stage. Standard errors are corrected by the delta method.

Figure B2 presents dynamic estimates for nightlights (Panel A), GDP (Panel B), and forest share (Panel C). White diamonds are the sum of baseline and lagged coefficients, which measure net impacts of cashew cultivation several years later. Corresponding tabular estimates showing are provided in Table A9. Although nightlights significantly increases one period after cashew expansion, the effect fades by the second period. The insignificant impact on GDP remains persistent up to two years later. In contrast, forest loss from cashew expansion is statistically significant and remains persistent across periods. We do not interpret the larger magnitude on the first lag of forest cover as evidence of worsening forest degradation, as confidence intervals overlap the contemporaneous estimate. We also interpret the second lag with caution due to data loss. Overall, these dynamic estimates imply that forest loss from cashew expansion is lasting, whereas corresponding economic benefits do not materialize even after several years.

6.3 Additional Robustness Checks

Having established robustness to alternative instruments (Table A3), we now probe estimate sensitivity further with a variety of additional tests. Table A4 presents robustness tests for GDP. While the main analysis finds a positive yet statistically insignificant impact of land *area* under cashews on local GDP, the results are similar when instrumenting the *number* of cashew plantations in an arrondissement instead (column 1). Estimates also remain stable when controlling for the interaction of price volatility and non-cashew cropland share (column 2), which accounts for substitution effects as a potential channel through which cashew price risk affects local income. We continue to find qualitatively positive impacts on GDP when controlling for price effects with alternative cashew price data from the International Nut and Dried Fruit Council (column 3). Lastly, estimates are robust to controlling for arrondissement-specific linear time trends, which flexibly account for "drift" in expected income driven by unobserved factors that vary across arrondissements at a constant rate over time. We drop arrondissement fixed effects in this specification because the time trends add over 100 controls (one for each arrondissement), which leaves little variation left for identification. Table A5 shows the same robustness tests with log nightlights as the outcome. We continue to find qualitatively positive but statistically insignificant impacts of cashew cultivation on local economic activity.

Table A6 presents robustness tests for forestland share. We find remarkably stable and robust negative impacts on forest cover across the same five robustness tests. Taken together, these results reinforce our finding that cashew cultivation in Benin degrades local forests without generating meaningful economic benefits. To further probe this finding and understand why farmers continue to cultivate cashews, we study mechanisms next.

7 Discussion

This section discusses why we find statistically imprecise income gains from cashew cultivation. We then use our econometric estimates to calculate a back-of-the-envelope estimate of aggregate environmental costs from cashew cultivation in Benin.

7.1 Mechanisms: Why do farmers cultivate cashews?

Our analysis yields a somewhat puzzling pattern: cashew cultivation in Benin causes substantial forest loss without generating measurable economic gains—at least none observable in satellite GDP data. Why, then, do farmers continue growing cashews? We consider two explanations: (i) cashews *are* profitable, but benefits are highly localized and obscured in coarsely measured satellite data, and (ii) farmers grow cashews for their second order income-smoothing benefits rather than for first order income opportunities. We find modest evidence for each explanation.

Reason I: Localized income effects The first reason we may not see strong economic benefits is that, despite its major advance, gridded GDP data from RHZ may be too coarse to capture concentrated income gains. By this line of reasoning, we should then be able to detect income gains in household data, higher-resolution GDP data, and in sub-samples of RHZ data where nearly all farmers grow cashews. We explore each of these in turn.

First, we already estimated household-level correlations in Section 4.2, which show that households closer to cashew plantations tend to be wealthier. But since DHS does not report if households cultivate cashews, it is unclear whether this correlation reflects farmers profiting from their own cashew cultivation. To investigate this, we estimate Equation 1 by land ownership (Table A8), and find that landowning households within 900m (median distance) of a cashew plot are significantly wealthier than landowners farther away (column 1). This heterogeneity is unlikely to reflect other spatial differences since we control for household latitude, longitude, and urbanization. Thus, it is hard to imagine why landowners within 900m of cashew plots experience a wealth premium unless they profit from growing cashews themselves. This heterogeneity is also apparent in other wealth measures such as having a flush toilet (column 2), electricity (column 3), fridge (column 4), and an educated household head (column 5). Of course, all of these estimates should be interpreted as correlations, as the survey data are cross-sectional.

Next, we test for income gains at finer scales using higher-resolution GDP data from Chen et al. (2022), which predicts local GDP at 1km resolution using nightlights as weights to distribute national GDP[13]. We did not use this data to begin with because it is only available for 2015 and 2019, halving our sample to just 206 observations. With two sets of shift-share interactions and multiple levels of fixed effects, there is insufficient variation to estimate our 2SLS model with this data. However, TWFE estimates of Equation 2 using this data yield a positive and significant point estimate (Table A7): a 10 pp. cashew expansion is associated with an 1.3% rise in GDP per capita. While not causal, these esti-

[13] Chen et al. (2022) predict total GDP in a cell. To obtain per capita estimates, we sum pixel values in each arrondissement, do the same for gridded population from WorldPop,. and then divide the two.

Table 6: Mechanisms: Income Volatility

Outcome: Log(Volatility + 1)	(1) Variance	(2) Std. Dev.	(3) Coeff Var.	(4) Rolling
CashewShare	-0.729	-0.659	-1.594**	-0.715
	(0.748)	(0.518)	(0.773)	(0.733)
Log Price × CashewShare	Yes	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓	✓
Department × Year FEs	✓	✓	✓	✓
KP (2006) F-Stat	9.45	9.45	9.45	9.45
Observations	412	412	412	412

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. Outcomes are log of annualized monthly nightlight volatility measured by variance (column 1), standard deviation (column 2), coefficient of variation (column 3), and 6-month rolling variance (column 4). "CashewShare" is the fraction of cells with cashew plantations, instrumented with the interaction of global cashew price variance with baseline cashew coverage. All regressions control for price level interacted with baseline cashew coverage, agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

mates illustrate that resolution matters: the development benefits of cashew agroforestry are sharply visible in fine-grained but not in coarsely measured GDP data. This logic explains why we find strong negative effects in fine-grained 200m forest cover data.

Lastly, to probe the income-concentration hypothesis with a more credible design, we re-estimate our 2SLS model with RHZ data on a restricted sample of places where cashew cultivation is widespread. If highly local (within-pixel) income gains are averaged out in coarse GDP data, then they should be detectable in coarse data if cashew farming is widespread across the entire GDP pixel. Conveniently, the RHZ resolution is 0.25°, roughly the area of an arrondissement. Figure B3 plots arrondissement level second-stage estimates (β from Equation 6) across quintiles of baseline cultivation ($C_{ad(t=2015)}$). Positive and significant (p=0.08) income gains emerge in the fifth quintile, where most land is planted with cashews. This exercise provides modest evidence that the noisy effects in our main analysis reflect localized income gains blurred in coarse satellite data.

Reason II: Second order income effects A second reason why we may not be seeing strong income gains in the main analysis is that farmers may in fact be seeking second order benefits—reducing income volatility—rather than first order income gains. Unlike annual crops, this logic envisions perennial cashew trees as capital assets: requiring up-front investment but yielding returns over time, helping to smooth income in uncertain environments. Income stabilization may outweigh the immediate profit motive, making cashew cultivation attractive even in the absence of observable gains in aggregate data. Since our GDP data is only available annually, we instead test this idea by calculating the volatility of monthly nightlights at the annual level for each arrondissement and using it as an outcome in Equation 6. The shift-share instrument is the same as before (Section 5.1).

Table 6 presents results for four outcome measures of income volatility: variance, standard deviation, coefficient of variation, and 6-month rolling variance. In line with the income-stabilizing motive, cashew cultivation reduces income (nightlights) volatility across the board, though statistical precision is only achieved when income volatility is measured by the coefficient of variation. The point estimate implies that a 1 pp. increase in land share under cashew cultivation reduces income volatility by 1.6%. While low precision across other volatility measures tempers strong conclusions, we interpret this table as suggestive evidence that cashew cultivation provides income-stabilizing benefits.

7.2 Aggregate Cost-Benefit Estimates

As a final exercise, we use our econometric estimates to calculate aggregate costs and benefits of expanding cashew cultivation. 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 (SCC).

To calculate costs, we use the IV coefficient in Table 5 to estimate deforestation from increasing land share for cashews by 10 pp. Since $\beta = -0.274$ (column 3) and the average arrondissement is 36,000 ha., then a 10 pp. increase in cashew land share causes forest loss of $0.10 \times -0.274 \times 36,000 = -986.4$ ha. per arrondissement. We convert this to dollars using the SCC. We focus on carbon for simplicity, acknowledging that forests provide a variety of ecosystem services. The carbon stock of Benin's forests is about 389 tCO_2/ha (Houssoukpèvi et al., 2022). Using the most recent SCC estimate of \$USD $51/tCO_2$ from the Interagency Working Group convened by the Obama Administration (IWG on Social Cost of Carbon, 2013), and the fact that there are 103 arrondissements in the study area, the aggregate ecological cost of cashew expansion is about \$USD 2 billion. The corresponding cost is \$USD 7.3 billion using the SCC from Rennert et al. (2022).

To calculate aggregate benefits, we conduct a similar exercise using the GDP coefficient in column 2 of Table 5 ($\beta = 0.448$). This calculation is meant only for qualitative comparison since β is statistically insignificant. Coefficient magnitude implies that increasing cashew land allocation by 10 pp. raises GDP per capita by 4.48%. Relative to mean GDP per capita of \$USD 1,172 (Table 1), this amounts to $1,172 \times 0.0448 =$ \$52.5 per person. Given mean arrondissement population of 20,644 and 103 arrondissements in the study area, the aggregate gain is \$USD 111.6 million. If we use the GDP coefficient from cashew-dominated areas (quintile 5 from Figure B3), which is statistically significant, the aggregate benefit is \$USD 144.5 million.

Thus, the implied cost-benefit ratio using the most conservation SCC is 18 using the main GDP estimate and 14 using the binned estimate. In other words, for each dollar earned from cashew cultivation, the ecological cost is \$14-\$18. This represents a lower bound since the cost is entirely based on carbon and excludes other ecosystem services.

8 Conclusion

This paper investigates the development-environment tradeoff from cultivating cash crops in Benin during 2015-2021. 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 nuanced development-environment tradeoff: cashew cultivation degrades local forests without generating measurable economic gains, at least none observable in coarsely measured GDP data. We document strong forest loss and weak income gains in both a TWFE design as well as a robust IV design using global cashew price volatility as part of a shift-share instrument.

Probing our results further, we find that they do not imply that cashew cultivation is unprofitable. Rather, we find evidence that income gains are concentrated in localized pockets, which are averaged out in coarsely measured GDP data. We also find modest evidence for a second explanation: farmers value cashew trees for their second order benefits—reducing income volatility—rather than first order income gains.

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 agroforestry sector. Our cost-benefit analysis suggests that the ecological cost per dollar generated from cash crops is about $USD 14-18. 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 land tenure system and climate suitable for cashews. It is unclear whether our findings generalize to settings outside of West Africa. In terms of internal validity, our IV estimates rely on relatively coarse gridded GDP data, which has several limitations. Although we complemented this data with data from cross-sectional surveys and nightlights, more research is needed to measure income effects of cashew agroforestry at the household level.

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.

Araujo, Rafael, Francisco Costa, and Marcelo Sant'Anna, "Efficient forestation in the brazilian amazon," *SocArXiv. December*, 2020, *10.*

Assunção, Juliano, Molly Lipscomb, Ahmed Mushfiq Mobarak, and Dimitri Szerman, "Agricultural productivity and deforestation in Brazil," *Climate Policy*, 2017.

Barrett, Christopher B, "On price risk and the inverse farm size-productivity relationship," *Journal of Development Economics*, 1996, *51* (2), 193–215.

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.

_ , Yu Na Lee, and David R Just, "Producer attitudes toward output price risk: Experimental evidence from the lab and from the field," *American Journal of Agricultural Economics*, 2020, *102* (3), 806–825.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," *The Review of Economic Studies*, 2022, *89* (1), 181–213.

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.

Brown, Christopher F, Steven P Brumby, Brookie Guzder-Williams, Tanya Birch, Samantha Brooks Hyde, Joseph Mazzariello, Wanda Czerwinski, Valerie J Pasquarella, Robert Haertel, Simon Ilyushchenko et al., "Dynamic World, Near real-time global 10 m land use land cover mapping," *Scientific Data*, 2022, *9* (1), 251.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural productivity and structural transformation: Evidence from Brazil," *American Economic Review*, 2016, *106* (6), 1320–1365.

Carreira, Igor, Francisco Costa, and Joao Paulo Pessoa, "The deforestation effects of trade and agricultural productivity in Brazil," *Journal of development economics*, 2024, *167*, 103217.

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$\times$ 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.

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.

Farrokhi, Farid, Elliot Kang, Heitor S Pellegrina, and Sebastian Sotelo, "Deforestation: A global and dynamic perspective," *Working Paper*, 2023.

Foster, Andrew D and Mark R Rosenzweig, "Economic growth and the rise of forests," *The Quarterly Journal of Economics*, 2003, *118* (2), 601–637.

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.

Haile, Mekbib G., Matthias Kalkuhl, and Joachim von Braun, "Worldwide Acreage and Yield Response to International Price Change and Volatility: A Dynamic Panel Data Analysis for Wheat, Rice, Corn, and Soybeans," *American Journal of Agricultural Economics*, 2016, *98* (1), 172–190.

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.

Hsiao, Allan, "Coordination and commitment in international climate action: evidence from palm oil," *Working Paper, MIT*, 2021.

Husak, Gregory J, MT Marshall, Joel Michaelsen, Diego Pedreros, Chris Funk, and Gideon Galu, "Crop area estimation using high and medium resolution satellite imagery in areas with complex topography," *Journal of Geophysical Research: Atmospheres*, 2008, *113* (D14).

INDFC, "Nuts and Dried Fruit Statistical Yearbook 2022/23," *Annual Report*, 2023.

IWG on Social Cost of Carbon, "Technical support document: Social cost of carbon for regulatory impact analysis under Ex-ecutive Order 12866.," *Technical Report*, 2013.

Jayachandran, Seema, "How economic development influences the environment," *Annual Review of Economics*, 2022, *14* (1), 229–252.

—, **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.

Krah, Kwabena, "Maize price variability, land use change, and forest loss: evidence from Ghana," *Land Use Policy*, 2023, *125*, 106472.

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.

Lundberg, Clark and Ryan Abman, "Maize price volatility and deforestation," *American Journal of Agricultural Economics*, 2022, *104* (2), 693–716.

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

Maggi, Federico, Fiona HM Tang, Daniele la Cecilia, and Alexander McBratney, "PEST-CHEMGRIDS, global gridded maps of the top 20 crop-specific pesticide application rates from 2015 to 2025," *Scientific data*, 2019, *6* (1), 170.

Mighty Earth, "The Cashew Conundrum," *Working Paper*, 2023.

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

Ozdogan, Mutlu and Curtis E Woodcock, "Resolution dependent errors in remote sensing of cultivated areas," *Remote Sensing of Environment*, 2006, *103* (2), 203–217.

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.

Rennert, Kevin, Frank Errickson, Brian C Prest, Lisa Rennels, Richard G Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum et al., "Comprehensive evidence implies a higher social cost of CO2," *Nature*, 2022, *610* (7933), 687–692.

Rossi-Hansberg, Esteban and Jialing Zhang, "Local GDP Estimates Around the World," Technical Report, National Bureau of Economic Research 2025.

Sandmo, Agnar, "On the theory of the competitive firm under price uncertainty," *The American Economic Review*, 1971, *61* (1), 65–73.

Scott, Paul, "Dynamic discrete choice estimation of agricultural land use," *TSE Working Paper no. 14-526*, 2014.

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.

Suri, Tavneet and Christopher Udry, "Agricultural technology in Africa," *Journal of Economic Perspectives*, 2022, *36* (1), 33–56.

Tandjiékpon, AM, "Analysis of the Benin cashew sector value chain," *African cashew initiative (ACi)*, 2010.

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.

USDA and TechnoServe, "TechnoServe (TNS) Benin: USDA Cashew Integration & Acceleration Program (CIAP) Monitoring & Evaluation Plan," *M & E Plan*, 2015.

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

A.1 Price Volatility and Land Cultivation

Identification via our shift-share instrument relies on the interaction of global price volatility with baseline cashew land share to generate a source of exogenous variation land allocated to cashew production (Section 5). This appendix elaborates the economic logic and empirical evidence linking price volatility to land cultivation decisions.

Land allocation decisions under crop price uncertainty depend on farmers' risk preferences and available risk coping strategies. Classical economic theory predicts that price volatility discourages cultivation of risky crops. Sandmo (1971) provides the canonical theory that risk-averse producers facing higher price variance reduce output, particularly at the extensive margin. This theory is supported by substantial empirical evidence from Sub-Saharan Africa and elsewhere (Krah, 2023; Lundberg and Abman, 2022; Haile et al., 2016). Under this view, we would expect cashew farmers in Benin to reduce cultivated area in response to greater cashew price volatility.

A contrasting view suggests that price volatility may in fact increase cultivation—particularly among farmers already specialized in the crop. To see this, note that cashew trees are capital assets that involve high upfront costs. For farmers with established orchards, agronomic know-how, and market relationships, the marginal costs of expansion are low. Even under volatile prices, these farmers may respond positively to price signals as long as average prices remain favorable. This channel is especially salient when market frictions limit risk-mitigation strategies (Suri and Udry, 2022). This logic echoes Barrett (1996), who shows that under price uncertainty, resource-constrained households may respond by over-allocating land to buffer against income volatility. While few empirical studies test this alternative mechanism, the closest is Bellemare et al. (2020), which finds no empirical support for the Sandmo (1971) prediction.

A.2 Appendix Tables

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

Table A1: Share of new cashew plantations not classified as forests.

Note: Cell values are the share of new cashew plantations (planted between the corresponding row year and the prior year) not classified as forests by VCF. VCF pixel values (percent forest) are classified as forest if they are above the threshold (columns).

Table A2: First Stage Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
Log Volatility $\times$ $C_{ad(t=2015)}$	0.057***	0.062***	0.062***	0.134***	0.172**	0.232***
	(0.013)	(0.014)	(0.014)	(0.030)	(0.074)	(0.059)
Log Price $\times$ $C_{ad(t=2015)}$	No	Yes	Yes	Yes	Yes	Yes
Agrochemicals	No	No	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Volatility Measure	Var	Var	Var	SD	CV	Roll
Arrondissement FEs	✓	✓	✓	✓	✓	✓
Department $\times$ Year FEs	✓	✓	✓	✓	✓	✓
KP (2006) F-Stat	18.65	20.17	19.84	19.43	5.40	15.62
Observations	412	412	412	412	412	412

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. The outcome is land area under cashews. Volatility is measured as variance (columns 1-3), standard deviation (column 4), coefficient of variation (column 5), and 6-month rolling variance (column 6). $C_{ad(t=2015)}$ is cashew land share in 2015. "Price" is global average monthly prices in year t. All specifications include controls for agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A3: Robustness to Alternative Volatility Measures

	Outcome: Log NTL			Outcome: Log GDP			Outcome: Forest Share		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
CashewShare	0.124	-1.817	0.007	0.516	1.493	0.575	-0.272***	-0.247**	-0.270***
	(0.965)	(1.898)	(0.973)	(0.406)	(0.953)	(0.422)	(0.070)	(0.124)	(0.070)
Log Price × CashewShare	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IV Volatility Measure	SD	CV	Roll	SD	CV	Roll	SD	CV	Roll
Arrondissement FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓
Department × Year FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓
KP (2006) F-Stat	19.43	5.40	18.75	19.43	5.40	18.75	19.43	5.40	18.75
Observations	412	412	412	412	412	412	412	412	412

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. "Cashew Share" is land share under cashew cultivation, instrumented with global cashew price volatility interacted with baseline cashew share. Price volatility measures include standard deviation (columns 1, 4, 7), coefficient of variation (columns 2, 5, 8), and rolling volatility (column 3, 6, 9). "Price" is global average monthly prices in year t. All specifications include controls for agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A4: Robustness Tests: Log GDP

	(1)	(2)	(3)	(4)
NumberPlantations	0.025			
	(0.022)			
CashewShare		0.571	1.028	0.436
		(0.477)	(0.660)	(0.447)
Log Volatility × CropShare	No	Yes	No	No
Log Price × CashewShare	Yes	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Linear Trend	No	No	No	Yes
Arrondissement FEs	✓	✓	✓	
Department × Year FEs	✓	✓	✓	✓
Price Data	FAO	FAO	INDFC	FAO
Observations	412	412	412	412
F-Statistic	19.03	16.94	9.45	14.91

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. The outcome is GDP per capita and transformed by $log(x + 0.01)$ to account for zero values. "NumberPlantations" is number of cashew plantations per km2 in an arrondissement. "CashewShare" is land share under cashews. Both are instrumented with global cashew price variance interacted with baseline cashew share. "Volatility" is measured by price variance. "Price" is global average monthly prices in year t. "CropShare" is land area under non-cashew crops. All specifications include controls for agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A5: Robustness Tests: Log Nightlights

	(1)	(2)	(3)	(4)
NumberPlantations	0.015			
	(0.055)			
CashewShare		-0.298	1.090	0.134
		(1.105)	(1.048)	(1.130)
Log Volatility × CropShare	No	Yes	No	No
Log Price × CashewShare	Yes	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Linear Trend	No	No	No	Yes
Arrondissement FEs	✓	✓	✓	
Department × Year FEs	✓	✓	✓	✓
Price Data	FAO	FAO	INDFC	FAO
Observations	412	412	412	412
F-Statistic	19.03	16.94	9.45	14.91

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. The outcome is nightlight radiance and transformed by $log(x + 0.01)$ to account for zero values."NumberPlantations" is number of cashew plantations per km2 in an arrondissement. "CashewShare" is land share under cashews. Both are instrumented with global cashew price variance interacted with baseline cashew share. "Volatility" is measured by price variance. "Price" is global average monthly prices in year t. "CropShare" is land area under non-cashew crops. All specifications include controls for agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A6: Robustness Tests: Forest Cover

	(1)	(2)	(3)	(4)
NumberPlantations	-0.016***			
	(0.004)			
CashewShare		-0.304***	-0.290**	-0.268***
		(0.088)	(0.129)	(0.080)
Log Volatility × CropShare	No	Yes	No	No
Log Price × CashewShare	Yes	Yes	Yes	Yes
Agrochemicals	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Linear Trend	No	No	No	Yes
Arrondissement FEs	✓	✓	✓	
Department × Year FEs	✓	✓	✓	✓
Price Data	FAO	FAO	INDFC	FAO
Observations	412	412	412	412
F-Statistic	19.03	16.94	9.45	14.91

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level. The outcome is forestland share. "NumberPlantations" is number of cashew plantations per km2 in an arrondissement. "Cashew-Share" is land share under cashews. Both are instrumented with global cashew price variance interacted with baseline cashew share. "Volatility" is measured by price variance. "Price" is global average monthly prices in year t. "CropShare" is land area under non-cashew crops. All specifications include controls for agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A7: Mechanisms: High-Resolution GDP Data

	(1) Log GDP
CashewShare	0.134***
	(0.045)
Controls	Yes
Data Source	✓
Arrondissement FEs	✓
Department × Year FEs	206
Observations	0.998

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are at the arrondissement-year level for 2015 and 2019. The outcome is GDP per capita from Chen et al. (2022), transformed by $log(x + 0.01)$ to account for zero values. "Cashew share" is the fraction of arrondissement land area under cashews. All specifications control for rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

Table A8: Mechanisms: Wealth and Farm Size

	(1) Wealth	(2) Toilet	(3) Elec.	(4) Fridge	(5) Educated
Land owner	-0.575***	-0.030***	-0.107***	-0.027*	-0.128***
	(0.066)	(0.011)	(0.031)	(0.014)	(0.031)
Near Cashew × Land owner	0.206**	0.029**	0.086**	0.030*	0.076*
	(0.095)	(0.012)	(0.041)	(0.016)	(0.042)
Near cashew (=1)	0.247	-0.029	0.066	-0.069**	-0.013
	(0.155)	(0.021)	(0.067)	(0.030)	(0.067)
Household Controls	Yes	Yes	Yes	Yes	Yes
Geography Controls	Yes	Yes	Yes	Yes	Yes
Arrondissement FEs	✓	✓	✓	✓	✓
Observations	2866	2866	2866	2866	2866
R^2	0.533	0.207	0.366	0.101	0.193

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Data are a household cross section. "Landowner" indicates if the household owns farmland. "Near Cashew" indicates whether the household's DHS cluster is below median distance to the nearest cashew plot. The outcome in column 1 is a wealth index between 1-5. The outcome in remaining columns are indicators for having a flush toilet, electricity, fridge, and an educated (> secondary school) household head. All regressions include survey weights and controls for household size, language, rain, temperature, lat/lon, and nightlights. Standard errors are heteroskedasticity-robust.

Table A9: Dynamic Estimates

	Log NTL		Log GDP		Forest Share	
	(1)	(2)	(3)	(4)	(5)	(6)
CashewShare (No Lag)	-0.013	3.857	0.248	0.657	-0.240***	-0.709***
	(0.886)	(2.785)	(0.239)	(0.445)	(0.061)	(0.180)
Cashew Share (Lag 1)	4.378*	3.261	-1.802	0.571	-0.215	-1.158***
	(2.216)	(3.370)	(1.240)	(1.163)	(0.134)	(0.317)
Cashew Share (Lag 2)		-2.154		-0.185		-0.736***
		(2.727)		(1.163)		(0.268)
Cumulative Lag	4.365	4.964	-1.554	1.043	-0.455	-2.603
S.E. (Delta Method)	2.233	6.132	1.257	2.249	0.124	0.598
p-value	0.051	0.418	0.216	0.643	0.000	0.000

Note: * $p < .1$, ** $p < .05$, *** $p < .01$. Lagged values are predictions of $CashewShare_{adt}$ from the first stage (Equation 5). "Cumulative Lag" reports the sum of the contemporaneous effect (No Lag) and corresponding lagged coefficients in each column. Standard errors in the footer refer to the cumulative estimate and are computed by the delta method. All specifications include controls for price levels, agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

B Appendix Figures

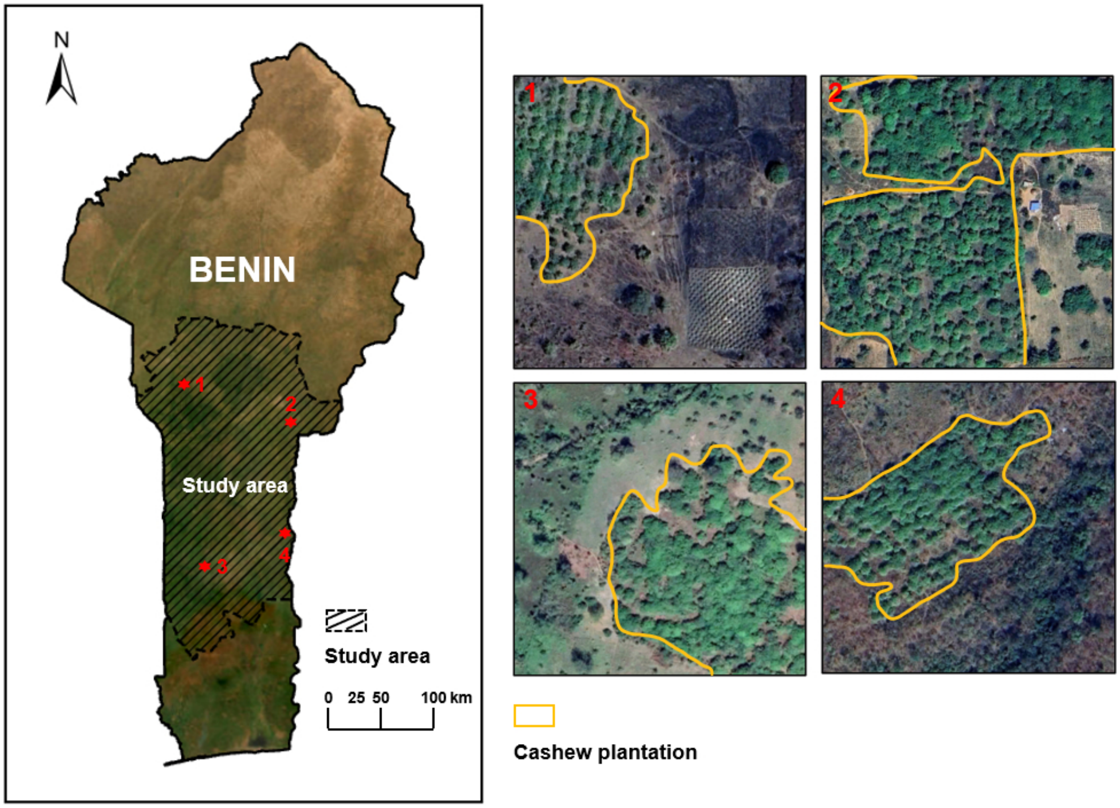

Figure B1: Cashew Plot Delineation

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

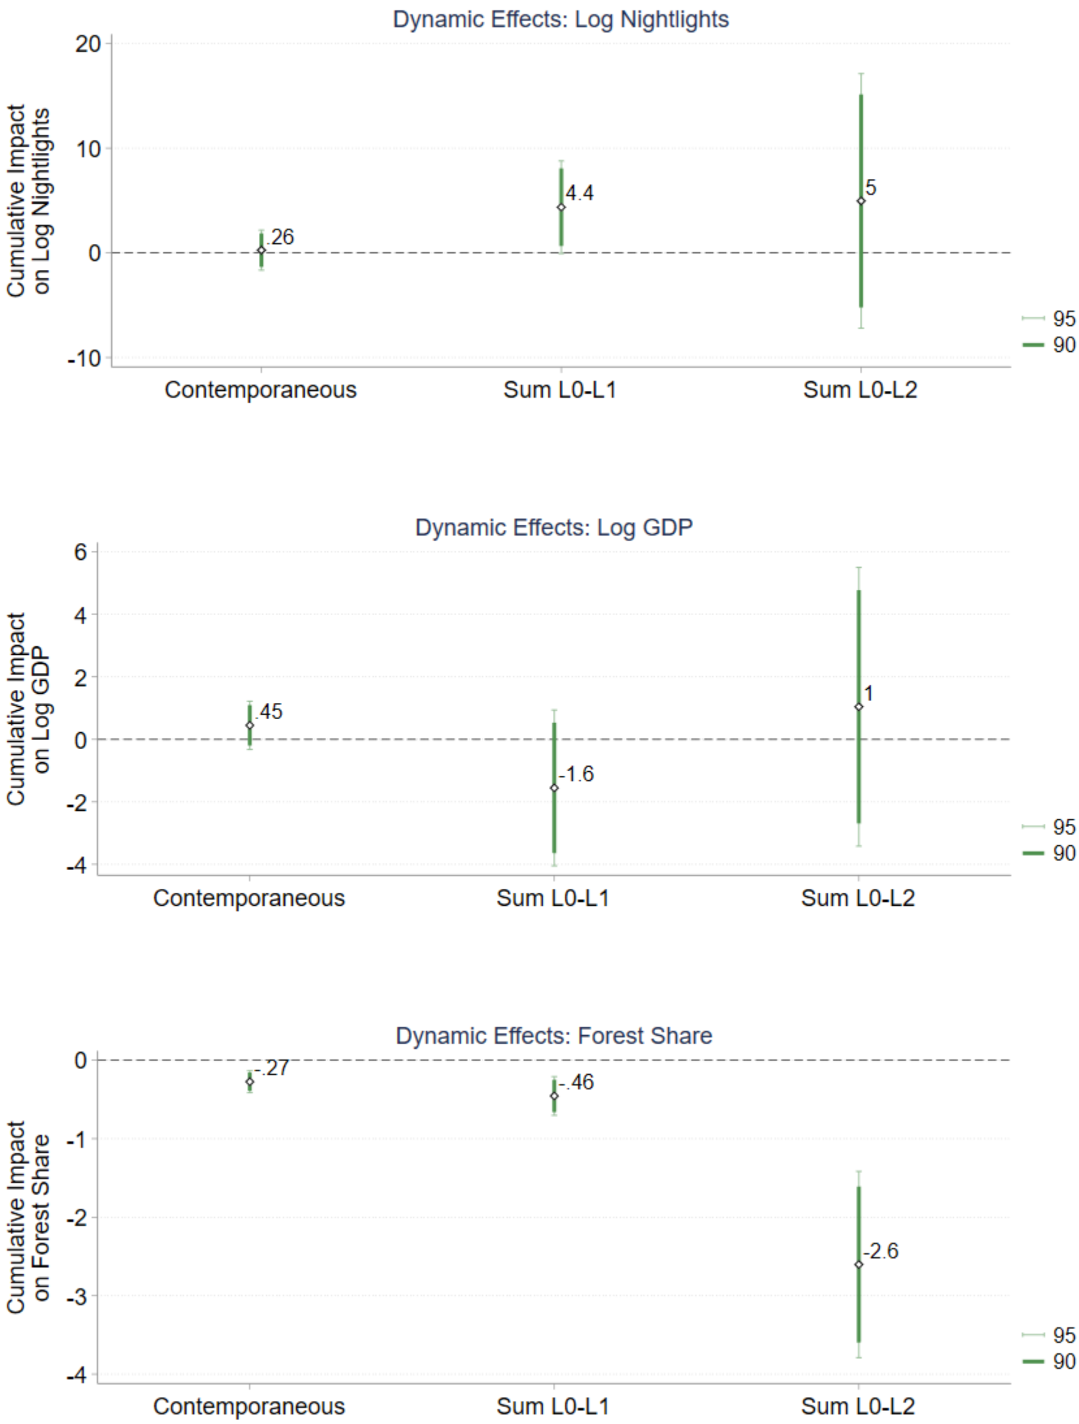

Figure B2: Dynamic Effects

Note: "Contemporaneous" repeats the main 2SLS result. Lagged coefficients (white diamonds) use predicted values of $CashewShare_{adt}$ from the first stage (Equation 5) using the shift-share IV. "Sum L0-L1" adds the first lag to the main specification and reports the sum of coefficients on the first lag and baseline effect, and so on. Shaded bars are confidence intervals. All specifications include controls for price levels, agrochemical use, rain, temperature, and drought intensity. Standard errors clustered by arrondissement.

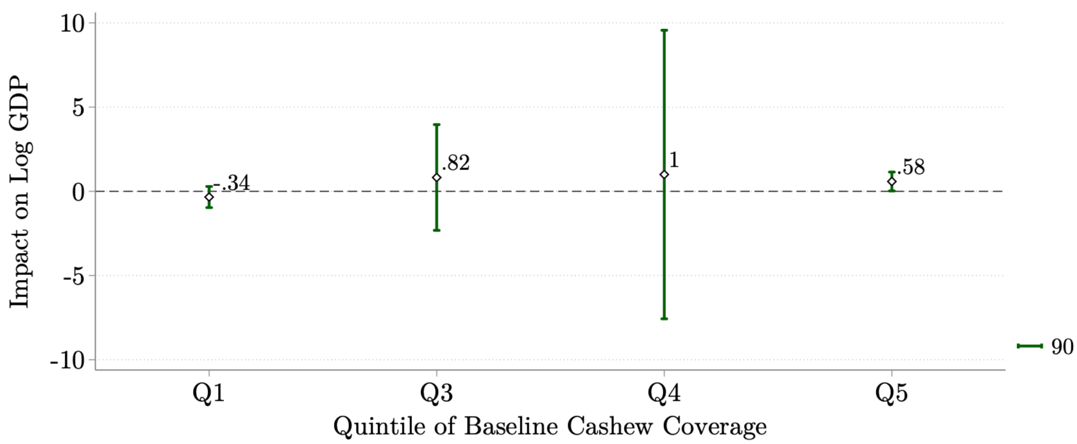

Figure B3: IV Estimates by Quintiles of Baseline Cashew Coverage

Note: White diamonds are coefficients from the second stage equation (Equation 6) across quintiles of baseline cashew cover. Bars are 90 percent confidence intervals. All specifications control for price levels, agrochemical use, rain, temperature, and drought intensity. Q2 is omitted from presentation due to overly large standard errors which warp the graph scale. Standard errors clustered by arrondissement.