

Materials and Methods:

1 Description of PES program

The PES contract was designed by CSWCT with input from other project partners, including a subset of the authors, using PES programs in other countries as a guide, but also tailoring the program to the context. The contract stipulated that enrollees could not cut down any trees with a trunk diameter at breast height (DBH) of 10 to 50 centimeters (cm) on their forestland. They were allowed to cut selected mature trees (>50 cm DBH), determined by the number of mature trees per species in a given forest patch. This provision was to give PFOs a small amount of leeway in case of emergencies. Participants were allowed to cut small trees (<10 cm DBH) for home use and to gather firewood from fallen trees.

An important consideration in designing the contract was that it be simple enough so that CSWCT could explain it to PFOs and could monitor compliance. As such, the contract is blunter than what would be prescribed by forestry experts to optimally manage the forest for its long-term health. The contract goal was to incentivize improved forest management relative to the status quo rather than to incentivize optimal forest management. This tradeoff between nuanced, multi-faceted requirements, on the one hand, and comprehension of the contract by all parties and feasibility of monitoring compliance, on the other hand, is inherent to all PES programs.

Costa Rica's PAS program, launched in 1997, and Mexico's Pagos de Servicios Ambientales Hidrológicos (PSAH) program, launched in 2003, are two of the major deforestation PES programs worldwide, and a comparison to them is useful (53). Both make annual cash payments like the CSWCT program. The payment level (in nominal USD) is \$65 per hectare per year for PAS (1.2% of Costa Rica's 2006 GDP per capita) and \$27 to \$36 per hectare per year (0.4% to 0.6% of Mexico's 2003 GDP per capita) for PSAH. The \$28 per hectare per year payment in the CSWCT program is 4.7% of Uganda's 2011 GDP per capita. PSA requires a minimum forest area to enroll of 2 hectares, and PSAH requires 50 hectares; the CSWCT program did not have a minimum size. PAS and PSAH offer five-year contracts; the CSWCT contract length was two years. Monitoring for PSA and PSAH is done via annual remote sensing and periodic on-the-ground spotchecks, while CSWCT used frequent on-the-ground monitoring. Among other differences, PSAH deemed forest with tree density less than 80% ineligible, and PSA required fencing the enrolled forest.

To gauge how the \$28 per hectare payment level compares to the opportunity cost of keeping forest intact, in qualitative fieldwork before the program rollout, PFOs reported that a large tree sold for \$20 to \$40. At baseline, 29% of PFOs reported earning revenue from timber products in the past one year, and among these, the mean (median) revenue was \$151 (\$40). The second type of forgone income is from cultivation that would have occurred on cleared land. While most households consume all of the crops that they grow, among households that sell crops for cash, self-reported income is in the range of \$30 to \$100 per hectare of cultivated land; if a PFO clears new land, he or she usually clears one or two 40 square meter (0.16 hectare) plots.

PES enrollees had the option to reforest up to 2 hectares of land. CSWCT provided the seedlings, and the PFO received 70,000 UGX per hectare per year if the seedlings survived. About half of enrollees took up the option. We focus on the averted deforestation component of the program because it was more significant in terms of payments, take-up, and contributions to tree cover. However, when comparing the amount of money paid to the amount of avoided deforestation, we

include the payments for reforestation, as this extra option and payment might have induced some PFOs to enroll. PFOs were not allowed to take up just the reforestation component; all enrollees were required to avoid deforestation on their forest land.

2 Experimental design and data

2.1 Sample of villages and forest owners

To determine the study sample, we first conducted a census of all PFOs in villages with primary forest in Hoima (late 2010) and northern Kibaale (early 2011). Villages with forest were identified using Landsat satellite imagery overlain with administrative boundaries. The field team visited these villages, met with the LC1 chairperson (elected village head), explained the study, and asked him or her to assemble 3 to 4 knowledgeable PFOs. This group then drew a rough map of the village and listed all PFOs in the village. The field team followed up with at least 3 spot checks per village.

Through the census, we identified 189 villages in Hoima and 91 villages in Kibaale with at least 1 PFO. Note that while most PFOs do not have a formal land title, which family de facto owned each plot of land was generally agreed upon within the village; the most common form of land dispute is among family members.

We narrowed this set of 280 villages to our sample of 121 villages by first excluding villages with fewer than 6 or more than 25 PFOs. We excluded villages with more than 25 PFOs because the project had a fixed budget for payments; with village-level randomization, additional PFOs per village add only limited statistical power. We excluded villages with fewer than 6 PFOs due to the fixed cost of working in a village and desire to have statistical power to measure village-level effects. We then excluded parishes (the administrative unit above a village) with only 1 eligible village because we had not yet determined whether a lottery to choose treatment villages would be conducted in each parish or each subcounty (administrative unit above a parish). We excluded two parishes in Kibaale district with very little intact forest; these were the only parishes with forest in their subcounties, so excluding them reduced the geographic spread of the study. Finally, one subcounty (19 villages) was set aside for a pilot. Kyabigambire was chosen as the pilot subcounty because the implementing NGO had close ties to the community and conducted its other activities there. At the NGO's request, four villages where they had very close ties were guaranteed the program. In the other 15 villages, we conducted a pilot of the baseline survey, subcounty lottery, program implementation, and endline survey; five villages were chosen for the treatment group. The lottery and program launch occurred in June 2011, two months before the main program launch.

We then conducted a baseline survey of PFOs in the sample villages in April to May 2011. In these villages, 1449 PFOs had been listed in the census, and we completed baseline surveys of 1174 (81%) of them. The main reasons for non-response were that we could not locate the PFO, or the individual stated that he or she was not a PFO. Some PFOs also refused to participate. An additional criterion to be in our final sample is that we collected valid GPS coordinates for the PFO's home at baseline. GPS coordinates allow us to identify the PFO's home in the satellite imagery and, thus, are necessary for having valid remote sensing data. The main reasons for missing GPS data were malfunctioning of the GPS units or enumerator error.

2.2 Randomization

After the baseline survey, 60 of the 121 villages were randomly selected to be in the treatment group. The randomization was conducted via public lotteries held in each of the 7 subcounties in our sample. In advance of each lottery, we divided the sample villages in the subcounty into two sets, balanced on covariates. Specifically, we first generated permutations that divided the villages into two sets differing in size by no more than one village. In subcounties with few sample villages, we constructed all possible permutations and for subcounties with many villages, for computational reasons, we generated a random sample of 1000 permutations. We then tested for balance on four village-level variables: number of PFOs in the sample, distance to a road, average per capita income, and average size of landholding. We considered the two sets of villages balanced if the standardized difference in the mean of each variable was less than 0.25. For each subcounty, among the balanced permutations, we randomly chose one to use in the public lottery. This procedure of prespecifying the set of balanced randomizations has advantages over re-randomizing until a threshold level of balance is achieved or choosing the most balanced permutation (54). At the public lotteries, which were conducted by the research team, The LC3 chairperson (elected head of the subcounty) or a stand-in drew one of the two lists out of a bin. This selected list constituted the treated villages. The lotteries occurred between August and December 2011.

After the lotteries occurred, we gave CSWCT the list of treatment villages, and they implemented the program. They rolled out the program, subcounty by subcounty, beginning in August 2011 and reaching the last subcounty in January 2012. Their first step when entering a community was to hold a parish-level meeting for eligible PFOs to advertise and explain the program. They then worked with interested PFOs to measure their area of eligible forest. The program did not have further eligibility requirements. For example, recent tree-cutting did not disqualify a PFO. CSWCT helped PFOs fill out and sign the PES contract form. For those who signed up, CSWCT monitored their land via spot checks and made annual payments to those who complied with the contract. Some PFOs who signed up for the program are not in our sample, either because we did not identify them as PFOs in our census or they did not complete our baseline survey. The monitoring occurred through in-person spot checks once every one or two months, during which the CSWCT employee checked for signs of recently cleared forest.

2.2 Satellite data

The QuickBird satellite is operated by DigitalGlobe. We contracted with an image reseller, Apollo Mapping, to task our images. QuickBird's swath width is such that it requires multiple flyovers to image our entire study region. Thus our baseline and endline image each consist of multiple strips taken during north-south flyovers on different days. It sometimes required multiple passes over a strip to obtain an image that, in aggregate for the study region, had less than 15% cloud cover, which is the quality standard that the vendor guarantees.

Most of the baseline images were collected in May to June 2011. Due to an extended rainy season, the last strips of the image were collected in December 2011 and January 2012, a few months after the program was rolled out. This timing would likely attenuate our estimated program effects; relative to control villages, treated villages should (and do) look more forested at baseline if the program has been in place and has been having impacts for a few months, which attenuates the treatment effect size that is estimated. In the analysis, we show that the

results are robust to restricting the sample to observations where the lottery occurred after the baseline image was taken (and indeed these results are larger).

We obtained our endline images while the program was still in effect. At scale-up, this type of PES program would likely be in place indefinitely or at least for a long duration, but this trial PES program lasted for two years. Once the financial payments end, PFOs have no incentive not to deforest their land, and thus a zero impact after the program is over does not mean that the program did not delay carbon emissions. The estimate of interest for a temporary program is to measure the postponement of deforestation it caused.

To avoid delays in the endline images such as occurred at baseline, we tasked the satellite to begin taking images in December 2012. The first strip was taken in December 2012, and most of the area was imaged between January and March 2013. The program was in place through at least August 2013 in all treatment villages. We estimate effects when the program had been in place for on average 1.5 years.

The remote sensing analysis was carried out by the Stanford Spatial Analysis Center. The initial step was to pre-process the images, for example to combine the image strips that comprise the overall baseline or endline image of the study region and to adjust for solar zenith angle and Earth-Sun distance at acquisition time. Calculations for top-of-atmosphere reflectance, which accounts for factors such as solar zenith angle and Earth-Sun distance at acquisition time, were done following DigitalGlobe's "Radiance Conversion of QuickBird Data" technical documentation. This step and subsequent ones were applied to the image area within the study village administrative boundaries, which were obtained from the Ugandan Bureau of Statistics.

Then the images were analyzed to classify each pixel as having tree cover or not using object-based image analysis (OBIA), a technique that takes advantage of the rich information contained in the relationship between adjacent pixels (55). In OBIA, the image is first segmented into polygons, or objects, with spectral and spatial homogeneity. The goal is for the polygons to correspond to real-life objects or features of the landscape such as a tree crown. OBIA is used for analysis of very high resolution images, i.e., <5 meter resolution, and more precisely, when the objects of interest are larger than the pixel size. For example, the shape, size, and context of a cluster of pixels help us identify it as a tree.

After segmentation, the polygon or object is the unit of analysis for classification. In our case, the goal is to classify each object as tree, non-tree, or cloud cover. (Pixels for which the landscape is obscured by clouds become missing data in our analysis.) The classification step uses a knowledge-based expert system in which the researcher defines a mutually exclusive and collectively exhaustive set of classes and a rule set for assigning objects to classes. The rules use several attributes of the objects such as color, shape, and size. In practice, classification is not done after the segmentation is finalized; the entire process is iterative, and the researcher visually inspects the classification results to refine the rules and parameters. Specifically, we used a change detection algorithm in which we first segmented and classified the baseline image, and then segmented the combined multidate images, using the initial classification as an object attribute in the multidate processing (56). We classified objects as persistent tree cover, persistent non-tree, tree gain, tree loss, or cloud-covered. To define the rules, we applied a classification and regression tree (CART) algorithm which determined the best rules for using the object attributes to predict forest classification, and "trained" the CART system with some of the ground-based data and validated it using the rest of the ground-based data.

The ground-based forest measurements were collected by a Ugandan forestry NGO, Nature Harness Initiatives (NAHI). After an initial rough classification of the study area by the Stanford lab, we sampled 440 geographic locations in the study area for the measurement. The ground measurements took place in spring and summer 2012 (so while the experiment was underway). NAHI followed the standard protocol for conducting a forestry inventory of a 20 meter by 25 meter plot. They identified all trees with a DBH of 10 centimeters or more, recorded the DBH, and estimated crown height and crown diameter. They also counted the number of trees of each species.

To improve statistical precision and adjust for any pre-trends in deforestation, we also use Landsat satellite images from the pre-intervention period, specifically from 1990 and 2010. Given the coarser (30 meter) resolution, we use a pixel-by-pixel technique and calculate the probabilistic fraction of each pixel that is photosynthetic vegetation. The classification does not make a distinction between trees and other photosynthetic vegetation such as grass. However, the technique we employ, the Carnegie Landsat Analysis System, was designed to detect deforestation and forest degradation using Landsat images (57).

2.4 Unit of observation and missing data

The main unit of observation in our remote sensing regression analysis is the village. The remote sensing analysis produces a classification of each pixel in the study area, and to convert this to village-level data, we overlay administrative boundaries of the village; the shapefiles were obtained from the Ugandan Bureau of Statistics.

We also use the PFO as the unit of observation. While we do not have the actual boundaries of their land, we have the geolocation of their home and their self-reported landholding. As a proxy for their land, we use a circle, centered on their home that is twice as large as the area of land they own. Note that homes are typically spread apart, and an individual's land is usually a contiguous plot including his or her home.

We use a circle larger than the actual land owned with the goal of including all or most of the land they own and nearby land to which they might shift their tree-cutting. If the land circles excluded much of their land, an estimated reduction in deforestation could simply reflect shifting of tree-cutting from near their home to the periphery. On average, PFOs own 10.8 hectares of land, so the average size of the circles we analyze is 21.6 hectares. We also show the results using circles that are 1 and 3 times the size of the landholding, and because PFOs' estimates of the area of their land might be inaccurate, we also use circles that are equal-sized for each PFO, based on the median land size of 5.3 hectares. Note that 1 PFO has missing data for the amount of land owned and 5 PFOs reported owning no land; for these observations, we have data for a circle around their home using the sample-median land size but not self-reported size. These PFOs reported owning forest in the screening questions at the start of the survey, and 3 of them reported owning land at endline.

Tree classification data are missing for some pixels due to cloud cover. For the village polygons, for on average 17% of pixels, the landscape is obscured by cloud cover at either baseline or endline. In the change-detection remote-sensing algorithm we use, pixels with cloud cover at either baseline or endline are treated as missing. The regression analysis weights observations by the proportion of the polygon with available tree cover data. For the PFO-level data, a second

source of missing data is that a part of the land circle might fall outside the study villages and therefore we do not have forest classification for that region. For 98 PFOs, or 9% of the sample, the entire land circle is missing when using PFO land circles that are twice the self-reported land area in size. These are missing observations. Table S1 breaks down the sample by availability of different outcome data. Note that, mechanically, smaller PFO circles are more likely to have completely missing data. Table S2 shows that attrition due to missing satellite data is uncorrelated with treatment status or program enrollment, but is correlated with land size (column 4). Conditional on land size, the observations with missing satellite data are similar to the main sample (column 5).

Shortly after the launch of the intervention, in early 2012, we mapped the boundaries of the forest plots for 707 of the 1099 PFOs. In the treatment group, enrollment is highly correlated with whether the PFO agreed to the land mapping boundary. (Program non-enrollees refusing to participate in the land mapping is consistent with fears of land grabs depressing program enrollment.) Because of the highly selected sample, we do not use these boundaries in the analysis of treatment effects. However, we can use these data to assess whether forest plots lie within village boundaries and PFO land circles. On average, 91% of the forest plot area lies within the village. In most cases where there is forest outside the village, a plot lies both within and outside the village. For 2.5% of PFOs, their forest is entirely outside their village. For the PFO circles, 23% of the forest plot lies within the circle that is two times the land area. Despite this small overlap, we do not enlarge the circles, as doing so leads to significant overlap between PFO circles.

2.5 Other data

Survey data

The baseline survey collected data on characteristics of the PFO's land, past tree-cutting behavior, attitudes toward the environment, consumption, and other topics. We also recorded the location of the PFO's home using hand-held GPS devices.

We conducted a similar survey at endline to collect data on secondary outcomes, and we also asked questions about the program such as why treatment PFOs did or did not take up the program. We successfully re-surveyed (93%) of the 1099 baseline PFOs. The survey completion rate was higher in control (95%) than treatment villages (91%), and in particular, PFOs in treatment villages who did not take up the program were less likely to participate, as shown in Table S2. Some of them had moved or died, but most just did not want to participate in the survey; anecdotally, PFOs who were wary of outsiders were less likely to participate in both the program and the endline survey.

In 53 cases, the endline respondent was a different member of the household than the baseline respondent. In 18 of these cases, the baseline respondent had died, in 3 cases he or she was too ill to participate, and in 22 cases he or she was temporarily or permanently away from the village. The remaining 10 cases are miscellaneous reasons.

When we estimate treatment effects on outcomes from the PFO survey, we control for the baseline outcome if it was collected, include an indicator variable for whether the baseline value is missing for that observation, and impute the baseline value to be the sample mean in these cases.

Administrative data on program enrollment and payments

From CSWCT's administrative records, we have data for each sample PFO on whether the PFO enrolled in the program, how much forest area the PFO owned and therefore enrolled, and how much the PFO was paid each of the two years of the program. We also have data on how much land PFOs set aside for reforestation, how many trees they planted, how many survived, and the payment received for this activity.

CSWCT deemed some enrollees as violating the contract and received no payment, some as fully compliant who received 100% payment, and some as partially compliant who received partial payment, at the discretion of CSWCT. Their records indicate whether non-compliance is due to forest conservation, reforestation, or both. Most of the non-compliance was related to the reforestation component of the project. CSWCT has missing data on monitoring and payments for 4 enrollees.

Variable construction

IHS transformation: The inverse hyperbolic sine transformation uses the function $\sinh^{-1}(x) = \ln(x + (x^2 + 1)^{1/2})$. Except for values of x close to 0, it approximates $\ln(x) + \ln(2)$. Before transforming variables in hectares with the IHS function, we rescale the variables in levels such that the 10th percentile of the baseline value for the PFO-level tree cover is 1. For baseline per capita income, the 10th percentile is 2500 UGX, and we use the same scale factor for all other monetary values before applying the IHS transformation.

Food and non-food expenditures: The categories for food were as follows: tea; soda; milk; sugar; meat; fish; rice; beans; salt; snacks and meals consumed outside the home; other foods. The categories for non-food were as follows: petrol and diesel; paraffin; body soap; clothing soap; other cosmetics, combs, razors; transport (excluding petrol and diesel); air time and use of public phones; domestic assistant or farm help; funeral expenses; shoes and sandals; clothing and bedding (not including school uniforms); livestock care (medicine, food, enclosures); school fees (not including uniforms and supplies); school supplies; bride price expenses; religious tithes.

Number of treatment villages within 5 kilometers: When calculating the number of villages near village A, a village is considered within 5 km if any portion of its polygon is within a 5 km distance of village A's centroid.

3 Additional analyses

3.1 Robustness checks

Table S5 shows unweighted results, which do not take into account the extent of missing tree-classification data. Table S6 addresses the skewness of the land area distribution by dropping the top 1% of PFOs in terms of baseline forest cover. (The results are similar when dropping outliers in terms of endline forest cover or baseline land ownership). Table S6 also shows results using equally-sized land circles for each PFO that are based on the median land area in the sample as an alternative approach to ensure that large PFOs are not driving the result. The median circles are on average smaller because of the skewed distribution. The effect size in levels is (almost

mechanically) smaller when using smaller circles, as is the average tree loss in the control group. The proportional effect size in these alternative specifications is similar to the main results.

Table S7 estimates the effects using alternatively-sized land circles. The first three columns use circles whose area is the PFO's land area; unless the PFO's landholding is exactly a circle centered on the PFO's home, these circles will omit much of his or her land. We continue to find statistically significant treatment effects, though again not surprisingly, the amount of tree gain in hectares is smaller because the circles are half the size. The next three columns show the results using circles that are 3 times as large as the PFO's land area. We find positive, statistically significant, and larger impacts on tree cover. The proportional effect size remains similar to the main results.

We placed an order for the baseline QuickBird images in May 2011, requesting that the study region be imaged as soon as possible. The majority of the region was imaged in May and June, but some of the area was not imaged until December 2011. The subcounty lotteries occurred between August and December 2011, which means the baseline image was taken after randomization in some cases. Hence, an important robustness check is to restrict the sample to cases where the subcounty lottery occurred after the date of the baseline satellite image. When we do so, we continue to find positive, statistically significant effects on tree cover, as shown in Table S8. The time span between baseline and endline images is longer in this subsample compared to the full sample (because we have dropped observations with late baseline data), so with a constant rate of deforestation, more deforestation should occur in the control group in this subsample, and, likewise, more deforestation should be averted in the treatment group. Indeed, we find a larger treatment effect than in the full sample, as well as a higher amount of deforestation in the control group.

3.2 Heterogeneous impacts

We examine heterogeneity in the effects by baseline PFO characteristics. We do so by estimating PFO-level regressions based on Table 3, column 5. We first analyze how the effects vary with the initial amount of forest. On the one hand, those with more forest have more potential forest to keep intact, which could lead to larger program impacts. On the other hand, the fact that more of their forest is intact suggests that these PFOs might rarely deforest, leading to smaller program effects. We find that the first of these possibilities is more relevant: Tree gain is larger for PFOs with more tree cover at baseline (Table S9, column 1). Column 2 shows this same pattern holds using the proportion of land that is tree-covered.

The next four columns examine heterogeneity by whether and why the PFO reported clearing trees recently. Those who had cut trees recently exhibit larger treatment effects, and this appears to be especially true for those who cut trees for timber products as opposed to cultivation. It has been hypothesized that PES programs are less effective for those who cut trees for large emergency expenses (58). Column 6 shows that the intent-to-treat program effect is not lower for this group, despite their lower likelihood of enrolling in the program. Column 7 shows that the higher the revenue from timber products at baseline, the larger the treatment effect.

The pattern seen across the first seven columns is that if a characteristic is predictive of more deforestation in the control group between baseline and endline (negative main effect of *Characteristic*), it is also associated with more averted deforestation in the treatment group

(positive interaction effect). Column 8 tests for this pattern more comprehensively. We use the control group and baseline data to predict the change in tree cover (just as we did in our analysis testing if enrollment is correlated with predicted tree loss; see Table S3). We then can calculate the predicted change in tree cover for both control and treatment PFOs and test for heterogeneity by this variable. The predicted value used as the regressor for control group observations is estimated excluding that observation itself to avoid bias (59). The negative interaction effect in column 8 indicates that the PES program caused larger gains in tree cover for treated PFOs whose tree loss, absent the program, would have been larger. This pattern reaffirms what earlier results suggested: To first approximation, enrollment in the program was unrelated to predicted counterfactual deforestation, enrollees complied and refrained from deforesting, and as a result, the largest program impacts are seen for those who would have deforested the most had the program not been offered to them.

3.3 Hypothetical scenarios of enrollment

To assess how important the composition of enrollment is to the magnitude of the program effect, we can calculate how the effect size would vary if enrollment had been (a) representative of all PFOs (b) concentrated among those who would have deforested the most absent the program (c) or concentrated among those who would have deforested the least. We use the distribution of tree loss (or gain) in the control group (shown in Fig. S5) as the measure of counterfactual deforestation absent the program and assume 32% enrollment to match the actual rate. For simplicity, suppose that all enrollees would fully comply with the program rules and avert their deforestation. Further, we assume that for enrollees who would have had net tree gain absent the program, the program has no effect, and we assume that the program has no effect on non-enrollees. Tree gain is due to natural regeneration of the forest. There is some, but quite limited, tree planting in the study area. Some of the observed tree gain is also likely due to measurement error.

In scenario (a), if enrollment were drawn evenly from across the distribution of PFOs, the intent-to-treat effect size would be 0.13 hectares per PFO. This value is the average in the control group of $-\min(\Delta\text{tree cover}, 0)$ multiplied by the enrollment rate. The *min* function operationalizes the assumption that the program has no effect on those with a positive counterfactual change in tree cover. In scenario (b), where the enrollees are those with the most deforestation, the effect size would be 0.36 hectares, considerably larger than in scenario (a). Finally, if those who enrolled were the PFOs with the least counterfactual deforestation, the effect size would be 0 hectares. (A third of control-group PFOs have a non-negative change in tree cover.) In this last scenario, despite the program enrolling and paying 32% of PFOs, all of the payments would have been inframarginal to conservation.

4 Cost-benefit calculation

Estimating the delayed CO₂

Global Forest Watch uses Landsat data to estimate the biomass in forests globally at a resolution of 30 meters. The most recent data at this resolution are from 2000. We calculate the average carbon per hectare in forests with at least 67% forest cover within our study village boundaries. Note that the effect size we estimate is in hectares of tree cover, so we might be underestimating

the carbon stored per hectare of tree cover by using the carbon density from areas with 67% forest cover. The value is 307 MT of above-ground biomass per hectare, and we apply the standard factor of 0.5 to obtain MT of carbon per hectare.

Estimating program costs

The value we use for payments to PFOs are inclusive of payments for reforestation and is based on Table 2, column 3 (\$36.09).

Monitoring costs were \$88 per program enrollee, or \$28 per eligible PFO. CSWCT hired forest monitors who each covered 30 enrollees. Monitors were paid \$90 per month, and for transportation, were given on average two bicycles each. Adding in repair costs, the bicycles cost \$480 per monitor over the course of the program. The source for this cost data is CSWCT. We assume a \$30 marketing and program management cost per eligible PFO and a 10% transaction fee for PES payments. These administrative costs are equivalent to \$0.34 per averted MT of CO₂. The monitoring costs could be lower at scale-up because, for the sake of the randomized trial, the program was rolled out in a geographically dispersed set of villages. It seems reasonable that with greater geographic density, monitors could conduct 2 rather than 1 spot check per day, in which case administrative costs would be \$0.26 per averted MT of CO₂. During the RCT, a forest monitor was supposed to conduct 1 spot check every one or two months per PFO; if a monitor spot-checked each enrollee once every 6 weeks, then the forest monitor conducted slightly less than 1 spot check per day, assuming 25 work days per month.

Valuing the delayed CO₂ emissions

We use an SCC of \$39 in 2012 US dollars (USD), which is based on interpolating the EPA estimates for 2010 and 2015. The EPA expresses the SCC in 2007 USD, and its values for the SCC in 2010 and 2015 are \$33 and \$38. Linear interpolation gives an SCC for 2012 of \$35 in 2007 USD, or \$39 converted to 2012 USD.

To quantify the value of the delayed CO₂ emissions, we assume that status quo deforestation is uniform over time. Thus, the averted deforestation over the two year program occurs on average at $t = 1$ year (because of the short time period, using this linear approximation for exponential discounting does not change the results). Our base case scenario assumes that after the program ends, treated PFOs deforest at a 50% higher rate than usual, and so it takes them 4 years to “undo” the averted 2 years of deforestation. That is, over the next 4 years, they carry out $4 \times 150\% = 6$ years’ worth of deforestation. Again using the uniformity assumption, the backlog of deforestation occurs $4 \div 2 = 2$ years after the program ends or at $t = 4$ years. Thus, the deforestation is delayed by 3 years (from $t = 1$ year to $t = 4$ years).

We can then calculate the value of delaying a MT of CO₂ by 3 years, or more generally by any delay length. We denote the delay length in years as D ; in the base case, $D = 3$. Next, we define the effective discount rate, combining the time discount rate and the rate at which the SCC rises: $r = (1.03)/(1.019) - 1 \approx 0.0108$. The effective discount rate is about 1.08%. While costs of CO₂ damage that are postponed are less costly because of the time discount rate, they are more costly because of the rising SCC (which is due to the fact that the marginal damages from CO₂ damages are higher when climate change has advanced further). Finally, the costs are born not when trees are cut but when CO₂ is emitted, which we assume is 10 years after tree-cutting in our base case.

To generalize, we denote this duration from deforestation to carbon emissions as S for storage. The value of delaying a MT of CO₂ emissions can be calculated as the cost of CO₂ emissions that occur in the absence of the program minus the cost of CO₂ emissions with the program:

$$Value = SCC * \frac{1}{1+r^S} \left(1 - \frac{1}{(1+r)^D} \right)$$

In scenario 2, $D = 1$ year because the uniform deforestation over the 2 year period is displaced to the end of the two-year period, so there is an average delay of 1 year.

In scenario 3, PFOs resume deforestation at their usual rate, and we need to make further assumptions about how PFOs deplete their forest for all future time. We assume they clear their forest at the same rate we observe between baseline and endline (9.1% of their stock of trees every two years) until their forest is depleted. This future stream of deforestation, discounted back to the present, is equivalent to 3.0 hectares per PFO, or 9.1 times the amount observed during the intervention. That is, the program delays by 2 years 3.0 hectares per PFO in present value terms. The present value calculation proceeds as follows. If we define the two-year deforestation rate in the control group as θ , then in the two years of the program, the flow of deforestation is θX , where X denotes the initial stock. In the next two-year period, the initial stock is $(1 - \theta)X$ so the flow deforestation is $(1 - \theta)\theta X$, and in the next two year period, the deforestation is $(1 - \theta)^2\theta X$. When calculating the present value of this infinite stream, we use the two year discount rate, which is $(1+r)^2 - 1$. The value of this infinite geometric sum is $\theta X / (1 - \lambda) = 9.1 * \theta X$, where $\lambda \equiv (1 - \theta) / (1+r)^2$. The present value of the postponed deforestation from the treatment effect is, thus, $9.1 * (\text{treatment effect on tree cover})$, or $9.1 * 0.326 = 3.0$ hectares. The program delays the 3.0 hectares of deforestation by 2 years in scenario 3.

When extrapolating scenario 3 to a permanent program, we assume the payments needed to deter deforestation rise over time at the same rate as the SCC. The restrictions on tree-cutting in the PES contract were tailored to the two-year duration; one would likely want to modify the contract conditions for a longer-duration program.

Scenario 4 uses the assumptions of Scenario 1 but the treatment effect size on tree cover is based on the PFO-level results (Table 3, column 5). Scenario 5 scales up the treatment effect to account for the fact that we measure the effects after 1.5 years rather than the full 2 years of the program. In the base case, we assume there are no further treatment effects in the final half year of the program, while in Scenario 5, we assume that the flow of averted deforestation is the same in the final half year as in the period we observe. Scenarios 6 and 7 change S to 5 years and 20 years, respectively. Scenario 7 halves the number of monitoring checks per day that the implementing agency's staff can conduct, increasing program administration costs.

Figure S1: **Study region.** The left panel shows in black the region in Hoima and Kibaale districts for which QuickBird images were obtained. The right panel shows the study villages within the region.

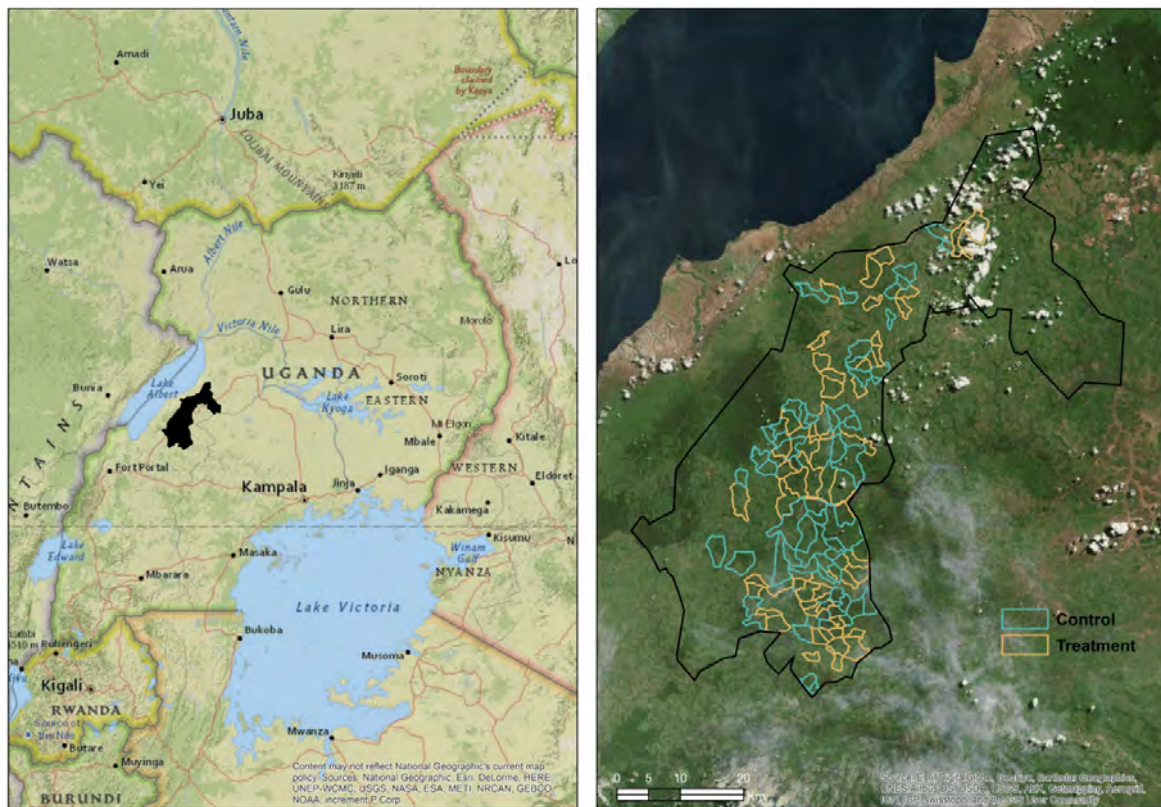


Figure S2: **Examples of tree cover change.** The analysis used a change detection technique that compared baseline and endline images and detected the change in tree cover. Four illustrative examples are shown.

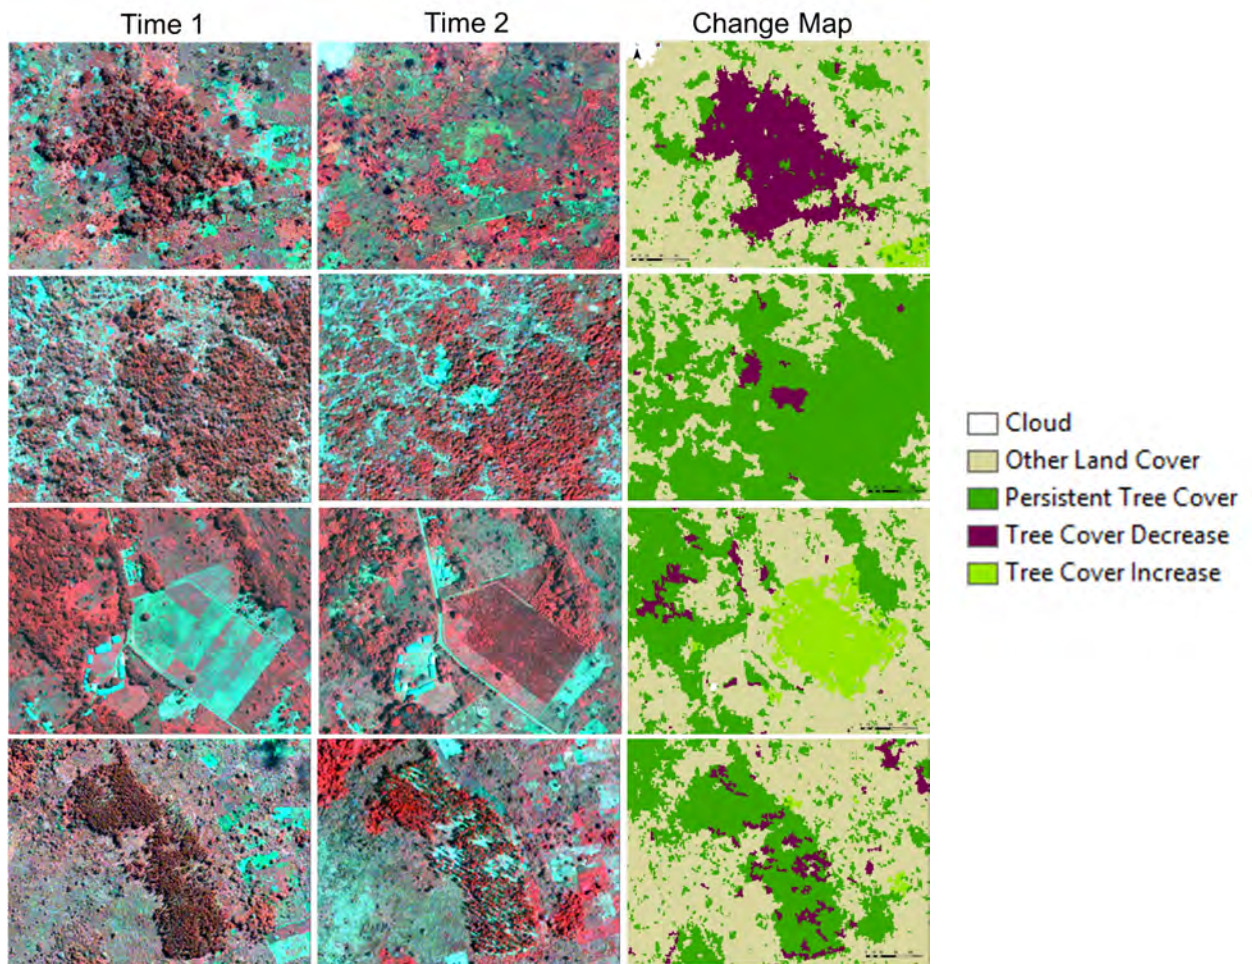


Figure S3: **Proportion of village and PFO polygons with available tree-classification data.** The figure shows histograms of the proportion of sample villages' (black outline) and PFOs' (solid gray) land that has available tree-classification data.

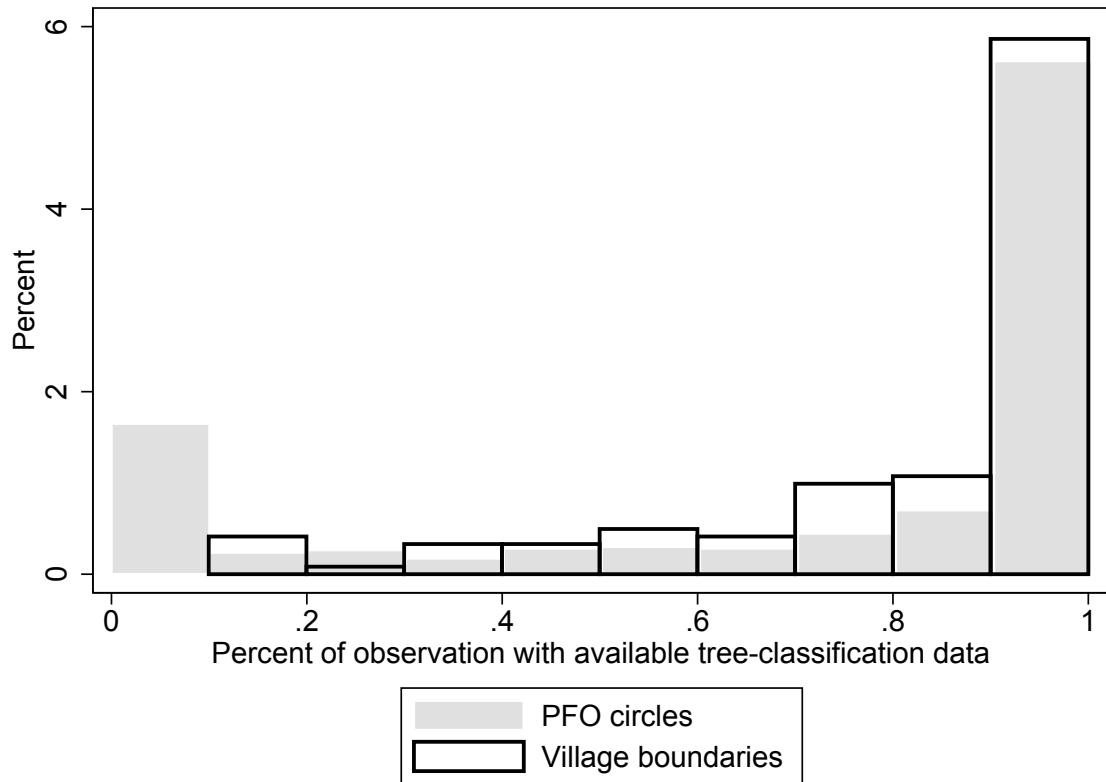


Figure S4: **Reasons for not enrolling in the program.** The figure is based on responses to questions asked on the endline survey to non-enrollee PFOs in treatment villages. PFOs were asked if they were aware of the program. Those who were aware but said they did not enroll were asked their reason for not enrolling.

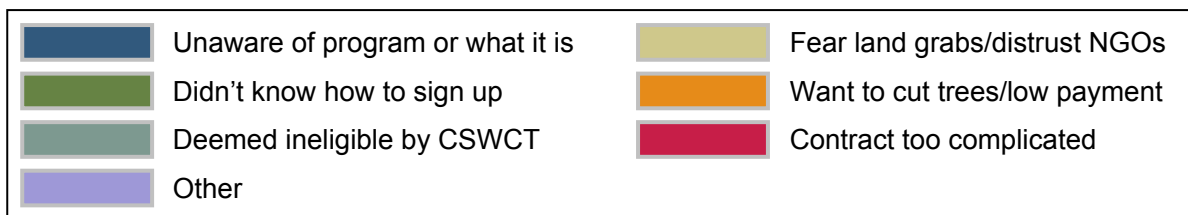
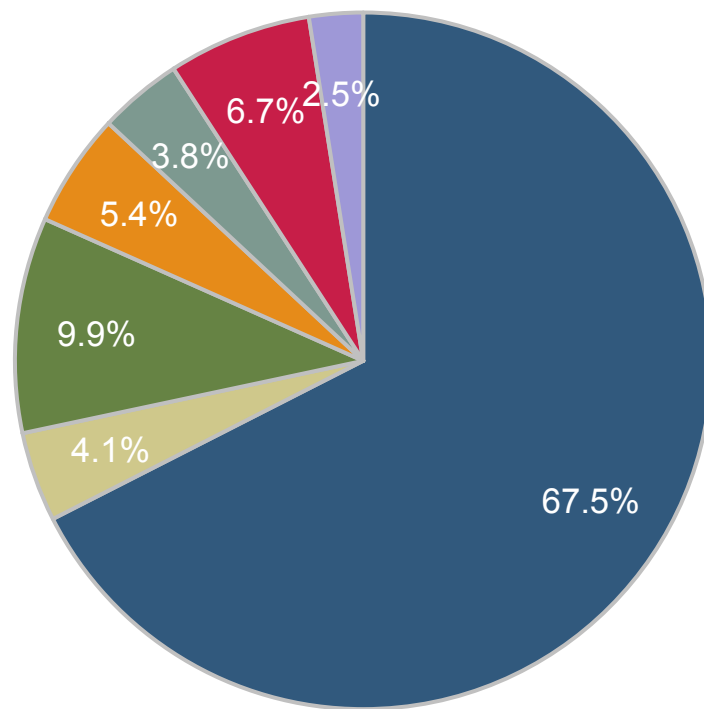


Figure S5: **Distribution of control group's baseline to endline change in tree cover.** The figure plots the density of the change in tree cover (in hectares) between the baseline and endline satellite images for the control group. The thick vertical gray dashed lines mark the 32nd and 68th percentiles of the distribution, which are used in the hypothetical scenarios considered in supplementary materials 3.3.

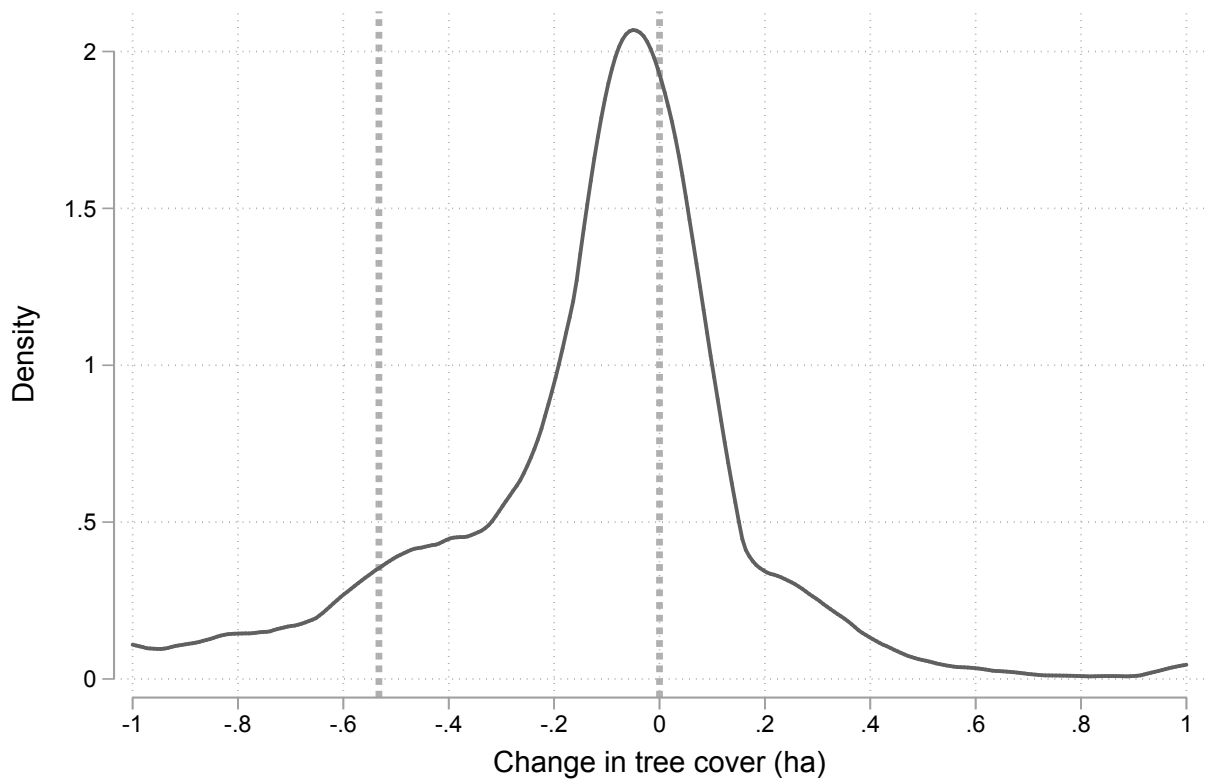


Table S1: **Sample attrition.** The top row lists the number of observations in the full sample, and the following rows provide the number of PFOs with available endline and satellite data.

	Number of PFOs		
	Treatment group	Control group	Total
Baseline survey (with GPS location of PFO home)	564	535	1,099
Baseline survey and satellite data for PFO land circle	508	487	995
HH reports owning no land	2	3	5
Didn't report land area	0	1	1
Missing tree-classification data for entire PFO land circle	54	44	98
Baseline survey and endline survey	512	508	1,020
Baseline survey, satellite data for PFO land circle, and endline survey	463	464	927

Table S2: **Correlates of sample attrition.** The table reports subsample means with standard deviations in brackets. Column 4 reports the regression-adjusted difference in mean between the full sample and the observations with missing tree-classification data divided by the pooled standard deviation. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. The standardized difference and p-value are based on a regression with subcounty fixed effects, with clustering at the village level. Column 5 is the regression-adjusted standardized difference when the inverse hyperbolic sine (IHS) of land area is included as a control variable. Column 6 reports the regression-adjusted standardized difference between the full sample and the observations with missing endline data. Summary statistics for *Enrolled* are based on the treatment group only.

	All PFOs	PFOs with missing tree classification data	PFOs with missing endline data	Std. diff. (1-2)	Adj. std. diff. (1-2)	Std. diff. (1-3)
	(1)	(2)	(3)	(4)	(5)	(6)
Household head's age	47.543 [14.122]	47.500 [13.555]	44.671 [14.772]	0.035	-0.070	0.207*
Household head's years of education	7.820 [4.093]	8.330 [4.223]	7.385 [3.845]	-0.147	-0.270**	0.119
IHS of self-reported land area (ha)	4.034 [0.996]	3.543 [1.348]	3.804 [1.165]	0.540***		0.253*
Self-reported forest area (ha)	1.893 [8.978]	1.213 [2.500]	1.197 [1.450]	0.072*	-0.088	0.078**
Cut any trees in the last 3 years	0.851 [0.356]	0.760 [0.429]	0.772 [0.422]	0.273**	0.165	0.271**
Cut trees to clear land for cultivation	0.238 [0.426]	0.231 [0.423]	0.152 [0.361]	0.022	0.018	0.199*
Cut trees for timber products	0.712 [0.453]	0.596 [0.493]	0.696 [0.463]	0.254**	0.165	0.088
Cut trees for emergency/lumpy expenses	0.270 [0.444]	0.212 [0.410]	0.316 [0.468]	0.136	0.082	-0.047
IHS of total revenue from cut trees	1.315 [2.183]	1.030 [2.176]	1.086 [1.992]	0.094	-0.010	0.135
Rented any part of land	0.180 [0.384]	0.096 [0.296]	0.139 [0.348]	0.257***	0.132	0.082
Dispute with neighbor about land	0.212 [0.409]	0.173 [0.380]	0.266 [0.445]	0.118	0.081	-0.148
Involved in any environmental program	0.106 [0.307]	0.071 [0.259]	0.013 [0.114]	0.148*	0.103	0.335***
Agree: Deforestation affects the community	0.544 [0.498]	0.531 [0.502]	0.434 [0.499]	0.045	0.068	0.246**
Agree: Need to damage environ. to improve life	0.054 [0.226]	0.071 [0.258]	0.065 [0.248]	-0.071	-0.077	-0.088
Treated	0.513 [0.500]	0.538 [0.501]	0.658 [0.477]	-0.072	-0.098	-0.324***
Enrolled	0.319 [0.467]	0.250 [0.437]	0.154 [0.364]	0.171	0.124	0.306*
Tree cover in PFO land circle (ha)	4.105 [10.978]		3.070 [3.901]			0.121*
% of PFO land circle with tree cover	0.204 [0.159]		0.211 [0.150]			-0.004
% change in vegetation in PFO land circle, 1990–2010	0.036 [0.062]		0.027 [0.060]			0.153
Observations	1,099	104	79			

Table S3: **Correlates of program enrollment in treatment group.** Standard errors are clustered by village. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. All columns include subcounty fixed effects, and the first three columns include the four village-level baseline variables used to balance the randomization. Missing independent variables have been imputed with the sample mean, and indicator variables for having a missing value are included in the regression. IHS denotes the inverse hyperbolic sine function.

	Enrolled (1)	Enrolled (2)	Δ Tree cover (ha) (3)	Enrolled (4)
Household head's age	0.002 [0.001]		0.003 [0.003]	
Household head's years of education	0.004 [0.005]		0.001 [0.014]	
IHS of self-reported land area (ha)	0.055* [0.028]	0.059** [0.024]	-0.318** [0.135]	
Self-reported forest area (ha)	-0.004 [0.006]		-0.073 [0.054]	
Cut any trees in the last 3 years	0.036 [0.094]		0.038 [0.163]	
Cut trees to clear land for cultivation	0.028 [0.057]		0.046 [0.140]	
Cut trees for timber products	0.090 [0.076]		0.049 [0.178]	
Cut trees for emergency/lumpy expenses	-0.129*** [0.041]	-0.099** [0.040]	-0.340* [0.195]	
IHS of total revenue from cut trees	-0.010 [0.010]		-0.049 [0.029]	
Rented any part of land	-0.046 [0.067]		0.004 [0.175]	
Dispute with neighbor about land	0.051 [0.045]		-0.063 [0.109]	
Involved in any environmental program	-0.014 [0.079]		0.216* [0.125]	
Agree: Deforestation affects the community	0.032 [0.039]		-0.037 [0.086]	
Agree: Need to damage environ. to improve life	-0.239*** [0.075]	-0.200*** [0.068]	-0.403 [0.346]	
Tree cover in PFO land circle (ha)	-0.003** [0.001]	-0.003** [0.001]		
% change in vegetation in PFO land circle, 1990–2010	0.285 [0.348]		1.787** [0.886]	
Predicted change in tree cover				-0.024 [0.034]
Sample	Treatment group	Treatment group	Control group	Treatment group
R ²	0.140	0.115	0.280	0.067
Observations	564	564	486	564

Table S4: **Effect of the PES program on tree-planting.** Standard errors are clustered by village. Asterisks denote significance: *** $p < .01$. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. The outcome variables used in columns 1 to 4 are from CSWCT administrative data on program participation. The outcome in column 5 is from the endline survey.

	Took up reforesta- tion option	Reforestation area (ha)	Total trees planted	Total trees survived	Have planted trees in the past 12 mths
	(1)	(2)	(3)	(4)	(5)
Treated	0.149*** [0.018]	0.101*** [0.016]	31.007*** [3.556]	9.813*** [1.555]	0.168*** [0.040]
Lee bound (lower)					0.157*** [0.041]
Lee bound (upper)					0.198*** [0.040]
Control group mean	0.002	0.001	1.710	0.933	0.282
Control group SD	[0.043]	[0.022]	[25.339]	[16.534]	[0.450]
Observations	1,099	1,099	1,099	1,099	1,019
Observations (Lee bounds)					998

Table S5: **Effect of the PES program on tree cover with unweighted regressions.** Standard errors are heteroskedasticity-robust in columns 1 to 3 and clustered by village in columns 4 to 6. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. The first 3 columns use village boundaries and the next 3 use a land circle centered on the PFO's home that is twice as large as his or her self-reported land area. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. Columns 2, 3, 5, and 6 control for dummy variables for the date of the baseline satellite image. Columns 2 and 3 control for 1990 and 2010 area covered by photosynthetic vegetation within the village polygon and in aggregate in PFO land circles for the village; columns 5 and 6 control for 1990 and 2010 area covered by photosynthetic vegetation within the village polygon and in the PFO's land circle.

	Village boundaries			PFO-level land circles		
	Δ Tree cover (ha) (1)	Δ Tree cover (ha) (2)	Δ Log of tree cover (3)	Δ Tree cover (ha) (4)	Δ Tree cover (ha) (5)	Δ IHS of tree cover (6)
Treated	5.395* [3.134]	5.439* [2.834]	0.0362 [0.026]	0.184* [0.099]	0.212** [0.096]	0.0301 [0.025]
Control group mean	-12.273	-12.273	-0.073	-0.319	-0.319	-0.062
Control variables	No	Yes	Yes	No	Yes	Yes
Observations	121	121	121	995	995	995

Table S6: **Effect of the PES program on tree cover, removing outliers.** Standard errors are clustered by village. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. All regressions and means are weighted by the proportion of available tree-classification data. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. Columns 2, 3, 5, and 6 control for dummy variables for the date of the baseline satellite image and 1990 and 2010 area covered by photosynthetic vegetation in the PFO's land circle and in the village polygon. The first 3 columns omit the top 1% of observations based on baseline tree cover. The next 3 columns use land circles that are equal-sized for all PFOs; the size is twice the sample median self-reported land size.

	PFO-level land circles (dropping top 1%)			PFO-level land circles (median-sized)		
	Δ Tree cover (ha) (1)	Δ Tree cover (ha) (2)	Δ IHS of tree cover (3)	Δ Tree cover (ha) (4)	Δ Tree cover (ha) (5)	Δ IHS of tree cover (6)
Treated	0.215** [0.101]	0.259*** [0.099]	0.0436* [0.023]	0.135** [0.057]	0.161*** [0.057]	0.0550** [0.025]
Control group mean	-0.336	-0.336	-0.073	-0.209	-0.209	-0.078
Control variables	No	Yes	Yes	No	Yes	Yes
Observations	986	986	986	1,002	1,002	1,002

Table S7: **Effect of the PES program on tree cover using different-sized land circles.** Standard errors are clustered by village. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. All regressions and means are weighted by the proportion of available tree-classification data. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. Columns 2, 3, 5, and 6 control for dummy variables for the date of the baseline satellite image and 1990 and 2010 area covered by photosynthetic vegetation in the PFO’s land circle and in the village polygon. The first 3 columns use land circles with an area equal to the PFO’s self-reported land area. The next 3 columns use land circles with an area 3 times the PFO’s self-reported land area.

	PFO-level land circles (x1)			PFO-level land circles (x3)		
	Δ Tree cover (ha) (1)	Δ Tree cover (ha) (2)	Δ IHS of tree cover (3)	Δ Tree cover (ha) (4)	Δ Tree cover (ha) (5)	Δ IHS of tree cover (6)
Treated	0.143** [0.068]	0.143** [0.063]	0.0470** [0.024]	0.332** [0.142]	0.371*** [0.137]	0.0385* [0.021]
Control group mean	-0.173	-0.173	-0.073	-0.539	-0.539	-0.073
Control variables	No	Yes	Yes	No	Yes	Yes
Observations	973	973	973	1,008	1,008	1,008

Table S8: **Effect of the PES program on tree cover in subsample with baseline satellite data prior to randomization.** Standard errors are heteroskedasticity-robust in columns 1 to 3 and clustered by village in columns 4 to 6. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. All regressions and means are weighted by the proportion of available tree-classification data. $\Delta Tree\ cover\ (ha)$ is measured in hectares. The first 3 columns use village boundaries and the next 3 use a land circle centered on the PFO's home that is twice as large as his or her self-reported land area. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. Columns 2, 3, 5, and 6 control for dummy variables for the date of the baseline satellite image. Columns 2 and 3 control for 1990 and 2010 area covered by photosynthetic vegetation within the village polygon and in aggregate in PFO land circles for the village; columns 5 and 6 control for 1990 and 2010 area covered by photosynthetic vegetation within the village polygon and in the PFO's land circle. The sample is restricted to observations for which the baseline satellite image was collected prior to the subcounty lottery.

	Village boundaries			PFO-level land circles		
	$\Delta Tree$ cover (ha) (1)	$\Delta Tree$ cover (ha) (2)	$\Delta \text{Log of}$ tree cover (3)	$\Delta Tree$ cover (ha) (4)	$\Delta Tree$ cover (ha) (5)	$\Delta \text{IHS of}$ tree cover (6)
Treated	8.287*** [2.977]	8.252*** [2.877]	0.0851*** [0.027]	0.322** [0.159]	0.417*** [0.155]	0.0866*** [0.030]
Control group mean	-16.576	-16.576	-0.122	-0.459	-0.459	-0.106
Control variables	No	Yes	Yes	No	Yes	Yes
Observations	78	78	78	580	580	580

Table S9: **Heterogeneous effects of the PES program on tree cover.** Standard errors are clustered by village. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. The outcome variable in all columns is $\Delta Tree\ cover\ (ha)$, which is measured in hectares. All regressions are weighted by the proportion of available tree-classification data. All columns include subcounty fixed effects, the four village-level baseline variables used to balance the randomization, dummy variables for the baseline satellite date, and 1990 and 2010 photosynthetic vegetation in the PFO's land circle and in the village boundary. In column 8, predicted tree loss is the predicted value from the regression reported in Table S3, column 3.

<i>Heterogeneous treatment effects on $\Delta Tree\ cover\ (ha)$ by:</i>								
	Above- median tree cover in land circle (1)	% of land circle with tree cover (2)	Cut any trees in the last 3 years (3)	Cut trees to clear land for cultivation (4)	Cut trees for timber products (5)	Cut trees for emer- gency/lumpy expenses (6)	IHS of total revenue from cut trees (7)	Predicted change in tree cover (8)
Treat \times Characteristic	0.472** [0.204]	2.054** [0.929]	0.424** [0.163]	0.025 [0.162]	0.342** [0.172]	0.410* [0.227]	0.122*** [0.044]	-0.686** [0.292]
Treated	0.020 [0.072]	-0.174 [0.145]	-0.094 [0.124]	0.265** [0.118]	0.016 [0.128]	0.140 [0.095]	-0.014 [0.079]	-0.008 [0.088]
Characteristic	-0.577*** [0.183]	-2.659*** [0.804]	-0.329** [0.134]	0.076 [0.120]	-0.335** [0.140]	-0.413** [0.198]	-0.105*** [0.039]	0.532 [0.396]
Observations	995	995	993	995	995	995	993	994

Table S10: **Testing for spillover effects and anticipation effects.** Standard errors are clustered by village in columns 1 and 3 to 7, and heteroskedasticity-robust in column 2. Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. All columns include subcounty fixed effects and the four village-level baseline variables used to balance the randomization. In columns 1, 2, 3, 6 and 7, regressions are weighted by the proportion of available tree-classification data, and the regressions include 1990 and 2010 photosynthetic vegetation, and dummy variables for the baseline satellite date. The outcomes in columns 4 and 5 are from the endline survey. Column 4 controls for the baseline outcome; the variable in column 5 was not collected at baseline.

	Δ Tree cover (ha) (1)	Δ Tree cover (ha) (2)	Δ Tree cover (ha) (3)	Get visits from timber dealers (4)	Increase in timber dealer visits last 2 years (5)	Δ Tree cover (ha) (6)	Δ Tree cover (ha) (7)
Treated	0.186 [0.265]	5.612** [2.664]		0.020 [0.034]	-0.016 [0.025]		
Treat \times Distance to forest reserve	0.011 [0.040]						
Treat \times Contiguous to forest reserve		2.391 [9.165]					
# of treatment villages within 5km			0.069 [0.066]				
Believes program likely to come to village						0.138 [0.084]	
Believes program ends in 2015 or later							-0.189 [0.142]
Sample	PFOs	Villages	Control PFOs	PFOs	PFOs	Control PFOs	Treatment PFOs
Observations	995	121	487	1,020	1,020	487	508