Protecting Ecosystems and Alleviating Poverty with Parks and Reserves: 'Win-Win' or Tradeoffs?

Paul J. Ferraro[†] and Merlin M. Hanauer[‡]

Abstract: National parks and reserves are globally popular approaches to protecting biodiversity and the supply of ecosystem services. Because these protected areas limit agricultural development and exploitation of natural resources, they are frequently opposed in developing nations where reducing poverty is an important social objective. Conservation proponents argue that protected areas can alleviate poverty by supplying ecosystem services, promoting tourism and improving infrastructure. Thus "win-win" scenarios may be possible in which ecosystems and their services are protected and poverty is alleviated. Previous studies (Andam et al. 2008, 2009) suggest that Costa Rica's protected area system induced both reduced deforestation and alleviated poverty. We demonstrate that these environmental and social impacts were spatially heterogeneous. Importantly, the characteristics associated with the *most* avoided deforestation are the characteristics associated with the *least* poverty alleviation. In other words, the same characteristics that have limited the conservation effectiveness of protected areas may have improved the social welfare impacts of these areas. These results suggest that 'win-win' efforts to protect ecosystems and alleviate poverty may be possible when policymakers are satisfied with low levels of each outcome, but tradeoffs exist when more of either outcome is desired.

Key words: ecosystems, poverty, protected areas, impacts, program evaluation, econometrics, Costa Rica

[†]Department of Economics, Andrew Young School of Policy Studies, Georgia State University, PO Box 3992, Atlanta, GA 30302-3992, USA; pferraro@gsu.edu

[‡]Department of Economics, Andrew Young School of Policy Studies, Georgia State University, PO Box 3992, Atlanta, GA 30302-3992, USA; mhanauer1@gsu.edu

1 Introduction

National parks and reserves are globally popular approaches to protecting biodiversity and the supply of ecosystem services (MEA 2005). These protected areas now cover approximately 12% of the world's terrestrial surface, with few nations lacking a protected area system (WPDA 2009). Despite the ubiquity of protected area systems, the published scientific evidence related to their environmental impacts is sparse and comprises predominantly case study analyses (MEA 2005; Joppa and Pfaff 2010). The evidence base related to their impacts on neighboring human communities is much weaker (Coad et al. 2008). A debate has emerged over whether the environmental goals of protected areas conflict with poverty alleviation goals, particularly in developing nations (Adams et al. 2004; Wilkie et al. 2006; Coad et al. 2008). Opponents highlight the role that protected areas can play in limiting agricultural development and exploitation of natural resources. Proponents highlight the role that protected areas can play in supplying ecosystem services, promoting tourism and improving infrastructure.

Empirical studies have found that protected areas, on average, are effective in reducing deforestation, although not as much as proponents may have expected (e.g., Cropper et al. 2001; Andam et al. 2008; Pfaff et al. 2009). Only a few well designed empirical studies have examined protected area impacts on socioeconomic outcomes in surrounding populations. They have found either no effect (Dufy-Deno 1998; Lewis, Hunt and Plantinga 2002, 2003) or a positive average effect (Andam et al. 2010; Sims 2010). As with most empirical studies in environmental policy, prior research on protected area impacts tends to focus on either environmental or social outcomes, but not both, and estimate only mean treatment effects.

In order to better understand the way in which a protected area system affects environmental and social outcomes, one must examine the two outcomes jointly and elucidate how different subpopulations are impacted. The econometric and program evaluation literature tends to focus primarily on the estimation of mean treatment effects, paying little attention to the impacts of treatment on population subgroups (Manski 2005; Crump et al. 2008). Yet, as noted by Manski (2005), average treatment effects may not provide sufficient information to a social planner whose goal is to maximize a specific social welfare function. For example, a medication may have positive mean health impacts on the treated population as a whole, yet men and women may respond differently. Suppose that the positive treatment effects are driven by males' strong responses whereas the medication has no, or deleterious, impacts on women. A physician would be remiss in prescribing such a medication without conditioning on subgroup characteristics.

Understanding subgroup impacts allows for the formulation of what Manski (2005) terms conditional empirical success (CES) rules. CES rules select treatments that maximize average impacts based on observable covariates (Manski 2005 pp.75). In the context of environmental policy, decisionmakers must possess an understanding of the heterogeneous impacts of ecosystem protection conditional on biophysical and demographic characteristics. For example, a planner may generate little avoided deforestation when establishing protected areas on high slope land if this land would likely remain forested in the absence of protection because it is less suitable for agriculture. Similarly, in an attempt to minimize negative socioeconomic impacts from land-use restrictions, a planner may not want to place protected areas in regions that comprise high proportions of agricultural workers if the opportunity costs of conservation in such regions greatly outweigh the local benefits from protected areas.

Costa Rica is an ideal setting for studying CES rules related to protected areas. Costa Rica is a biodiverse developing nation with rich and reliable spatially explicit data on biophysical and demographic characteristics. It was an early adopter of protected areas in the late 1960s and early 1970s and, by 2000, had protected about 25% of the nation. Despite these efforts to protect ecosystems, however, Costa Rica experienced a substantial amount of deforestation over the last 50 years: of the approximately 3 million hectares of forest in 1960, more than 1 million had been deforested by 1997 (Andam et al. 2008). The Costa Rica government has established a goal to be a model of sustainable development in Central America (Rubin and Hyman 2000). Most importantly, the available empirical evidence (Andam et al. 2008, 2010) suggests a 'win-win' scenario in which both avoided deforestation and poverty alleviation were, on average, achieved in and around Costa Rican protected areas. In order to examine this conjecture more deeply, we examine the heterogeneity of the protected area impacts conditional on biophysical and demographic characteristics. We find that the characteristics associated with the most avoided deforestation are the characteristics associated with the least poverty alleviation. While our analysis confirms that Costa Rica's protected areas system did lead to moderate levels of avoided deforestation and poverty alleviation,even among high-poverty areas, it also points to tradeoffs if decisionmakers desire higher levels of either outcome.

1.1 Background

Two studies of the impacts of protected areas on avoided deforestation (Andam et al. 2008) and poverty (Andam et al. 2010) comprise the point of departure for our study. Both studies use quasi-experimental matching techniques to obtain estimates of the average treatment effect on the treated (ATT). Estimating the ATT is akin to asking, "what would deforestation or socioeconomic outcomes have been had these areas not been protected?" Using digital forest cover data, Andam et al. (2008) estimate the amount of avoided deforestation between 1960 and 1997 that can be attributed to the designation of protected areas prior to 1980.[1] Conventional methods of analysis in the conservation literature simply compare deforestation outcomes on protected and unprotected parcels. Using these methods yields estimates that imply protected areas were accountable for a 44% reduction in deforestation. These estimates are inherently biased due to the nonrandom designation of protection. Protected land parcels are observably different from unprotected parcels based on covariates that have been found in other studies to affect deforestation. To control for selection on observable characteristics, the authors create a representative counterfactual group by matching unprotected land parcels to protected parcels based on key observable covariates. The resulting estimate of avoided deforestation is a more modest 11% reduction in deforestation attributable to protection. Their study confirmed that protected areas did indeed prevent deforestation, but because they tend to be placed on land that is undesirable for agriculture, the deforestation they avoid is modest. The placement of protected areas on land poorly suited for agriculture is a global phenomenon (MEA 2005).

Andam et al. (2010) use Costa Rica census tracts (segmentos) as the units of analysis to estimate the impact of protected areas established prior to 1980 on poverty between 1973 and 2000. Similar to Andam et al. (2008), the authors use matching techniques to form

[1] They also estimate the impact of protected areas established after 1980, but the focus of our analysis is on the areas established before 1980.

a counterfactual sample that is similar to the treated census tracts based on observable covariates that are believed to affect both designation of protected areas and socioeconomic outcomes. Their results indicate that the mean poverty was 1.3 points lower in census tracts with more than 10% of their area protected compared to similar matched census tracts with less than 1% protected land. This reduction is equivalent to an effect size of 0.2 (impact divided by standard deviation of the matched control group). Selection bias was substantial because protected areas tend to be placed in high poverty areas with low potential for economic growth. A simple comparison of census tracts with and without protected areas would lead to biased estimates that imply protected areas exacerbated poverty.

2 Data

2.1 Baseline Data Sets

We use data from Andam et al. (2008) and Andam et al. (2010) to estimate the heterogeneous impacts of protection, conditional on biophysical and demographic characteristics. The deforestation analyses use digital forest cover boundaries from 1960 and 1997, and georeferenced land characteristics that are believed to influence both the designation of protected areas and deforestation (see Table 1 and Andam et al. 2008 for details). To ensure comparability, the sample land parcels from Andam et al. (2008) are used. Forest cover outcomes are calculated using geographic information systems (GIS) and digital forest cover maps from 1960 and 1997. Twenty thousand three-hectare land parcels (minimum mappable unit) were selected at random from the 1960 forest cover layer. This layer predates protected areas and thus serves as the baseline forest cover,which can be compared across time to the 1997 forest cover. Forest cover is represented by a binary indicator: a land parcel is considered forested if it has greater than 80% canopy cover. The outcome for each land parcel is denoted by a 0 if it had not been deforested by 1997 and a 1 if it had been deforested. To determine if a land parcel is considered protected for the analyses, a layer containing all protected areas established prior to 1980 is overlaid with the land parcels. Costa Rica's protected areas system includes International Union for Conservation of Nature (IUCN) management categories Ia, I, II, IV and VI, which represent the level of land-use restrictions: Ia being the most strict. The proportions of these IUCN categories in our sample are: Ia&I = 0.038; II = 0.43; IV = 0.038; VI = 0.496. Land parcels within the boundaries of a protected area receive and indicator of treatment.[2] Similar overlays are performed with other data layers to create a set of covariates associated with each observation.

In the socioeconomic analyses, the unit of observation is the census tract. The 1973 census is used as the baseline year (see appendix) and demographic data are geocoded to their respective census tracts to form a set of covariates for each observation. In 1973 Costa Rica contained 4,694 census tracts with an average size of 8.82km^2 (range: 0.00466-836 km^2). To determine if a census tract is considered protected for the analyses, a layer containing all protected areas established prior to 1980 is overlaid with the census tracts. As in Andam et al. 2010, a census tract is considered protected if at least 10% of its area is occupied

[2]Of the 20,000 land parcels in the random sample, 3,380 were protected prior to 1980. To avoid potential bias in estimates we follow Andam et al. (2008) and drop any land plot that was protected between 1980 and 1997 from the pool of potential counterfactual observation. 4,717 land parcels are excluded prior to the analysis for various reasons, justification for which can be found here: http://www.pnas.org/content/suppl/2008/10/14/ 0800437105.DCSupplemental/0800437105SI.pdf

by protected land (results are robust to changes in this threshold definition).[3] Conversely, any census tract that contains less that 1% protected land is considered unprotected and a binary treatment indicator is assigned accordingly.[4] A poverty index is derived for each tract from census data following Cavatassi, Davis and Lipper (2004). Higher levels of poverty are associated with greater poverty index values (negative poverty index values indicate low levels of poverty). The censuses from which the poverty index is derived were conducted in 1973 and 2000. In the analyses, the poverty index calculation for 2000 is the outcome of interest. To match tracts on baseline characteristics, we use the matching covariates used in Andam et al. (2010), which include the 1973 poverty index and other baseline covariates that affect both protected area location and economic growth (see Table 1 and Appendix for more details). As noted in Andam et al. (2010) there were some protected areas established prior to our baseline year (1973). However, a majority of the protected areas in our sample (approximately 85%) were established between 1973 and 1979. Further, when we drop the protected areas that were established prior to 1973 from the analysis, the qualitative results remain the same.

2.2 Subgroup Variables

Agriculture has played a central role in the history of deforestation and economic growth in Costa Rica (de Camino et al. 2000). For protected areas to stem deforestation, they must

[3] We use the 10% threshold in accordance with Andam et al. 2010. A 10% threshold was chosen because protecting 10% percent of the worlds' ecosystems was the goal of the 4th World Congress on National Parks and Protected Areas (Andam et al. 2010). Andam et al. (2010) show that their results are robust to changes in this threshold value (alternatively defined as 20% and 50%).

[4] Of the 4,691 census tracts, 249 are considered protected (treated) prior to 1980 and 4164 are considered potential counterfactual observations. To avoid bias in the analysis, 278 tracts with protection between one and ten percent are dropped from the analysis.

be placed in areas in which the forest was at risk of conversion to other uses and they must be enforced. Thus we wish to estimate treatment effects within subgroup covariates that capture the returns to agriculture, the dependence of an area on agricultural activity, and the ease of enforcement. All threshold values used to define subgroups are baseline, pre-protection values, and we test the sensitivity of our results to the choice of these thresholds.

Land use capacity is a measure of land's suitability for cultivation that takes into account such factors as soil, precipitation, climate and slope (see Table 1). Land parcels designated as land use capacities 1, 2, 3 or 4 are denoted as land with high returns to agriculture. In a related study, Pfaff et al. (2009) estimate how avoided deforestation between 1986 and 1997 on protected Costa Rican land parcels varies according to geographic characteristics that categorize the parcels as either "high" or "low" pressure. They use slope as a subgroup variable under the assumption that high-sloping land is less productive and more costly to cultivate (it is also more costly to log). To permit comparisons between our study and their study, as well as to provide another proxy for returns to agriculture in an area, we designate land with a slope of more than 23% as high-slope areas (the median value of the deforestation analysis sample).

The returns to agriculture are higher on land that is closer to cities with markets. Yet cities also tend to be the seats of government enforcement of deforestation laws and thus their proximity to a plot may have a countervailing effect on ecosystem conversion. In other words, parcels far from cities may have low returns to agriculture, but less enforcement of land-use laws. Cities also provide a tourism gateway and thus may further mediate the economic impacts of protected areas. As a measure of access to markets we use the distance to one of Costa Rica's three major cities. Land parcels more than 57 kilometers of San

Jose, Puntarenas or Limon are considered to be high-distance parcels (the median value of the deforestation analysis sample).[5] We also ran analyses with distance to road, which is a covariate that captures the same economic relationships as distance to cities, but we omit it from the final analyses because it provides qualitatively similar results to distance to major city as a measure of access to markets. Among treated parcels, distance to major city and distance to road have a (Pearson's) correlation coefficient of 0.704.

The aforementioned covariates are measures of the characteristics of the land parcel. To characterize the economic conditions in the surrounding area, we use the percentage of adults employed in the agricultural sector in the census tract. Robalino (2007) presents a theoretical model that predicts negative economic impacts from protected area will be stronger in areas with greater proportions of agricultural workers. We define areas with high-baseline agricultural workers as census tracts with more than 13% of the workers employed in agriculture (the median value of the poverty analysis sample).

As a final variable to form subgroups for analysis, we chose a variable based on policy-relevance rather than theory. As noted in the Introduction, the relationship between protected areas and poverty is important in international environmental policy debates (Adams et al. 2004; Wilkie et al. 2006; Coad 2008). Thus differences in outcomes for low-poverty and high-poverty regions are of interest to decisionmakers. We define an area as high-poverty if it has a baseline poverty index of greater than 18 (the median value of the poverty analysis sample).

[5] Pfaff et al. (2009) use distance to San Jose.

3 Methods

3.1 Estimator

Andam et al. (2008) and Andam et al. (2010) use matching techniques as identification strategies to estimate the average treatment effect on the treated (ATT).[6] Naturally, once an area is protected one is unable to observe what would have happened in this area had it not been protected (termed the fundamental problem of causal inference by Holland 1986). Matching therefore constructs an *ex post* counterfactual group of unprotected units that is observably similar to the group of protected units in terms of key covariates believed to affect both outcome and selection into treatment. The underlying goal is to achieve balance across the key covariates similar to that achieved by a randomized experiment. To achieve this balance, Andam et al. (2008) and Andam et al. (2010) use bias-adjusted nearest neighbor Mahalanobis matching.

Our study uses a quasi-experimental design to conduct subgroup analyses. We form an *ex post* control group, based on observable covariates, on which we conduct subgroup analyses with the ATT as the estimand of interest. Subgroup analyses are relatively rare in the program evaluation literature (Crump et al. 2008), but can provide valuable insight even when average treatment effects are not significantly different from zero (Crump et al. 2008; Imbens and Wooldridge 2009). Perhaps the most common method of subgroup analysis is the use of interaction terms in a regression framework. However, even if this type

[6] ATT is the appropriate estimand in these studies because the interest lies in the sample of areas that were protected as compared to areas that could have been protected (unprotected areas that are similar to protected areas based on key covariates). Alternatively, the average treatment effect (ATE) additionally imputes values for all control units (finds the best match from the treatment group). Given that there are many observational units that would never feasibly be selected for protection, using ATE as the estimand makes little sense.

of approach were preceded by matching (Ho et al. 2007) or trimming (Imbens 2004; Imbens and Wooldridge 2009), the subgroup treatment effect estimate is more similar to the Average Treatment Effect (ATE) than the ATT. Crump et al. (2008) suggest estimating separate regression functions (parametric or nonparametric) for treatment and control groups, and testing for differences in the coefficients on the subgroup variable. While this approach is more transparent, it too is an estimand that is more in-line with ATE than ATT.

We propose an estimator that uses regression-adjusted imputation methods (see Imbens 2004; Abadie et al. 2004; Imbens and Wooldridge 2009) and a general form matching-based variance estimator (Abadie and Imbens 2006; Imbens and Wooldridge 2009) to estimate subgroup effects in terms of ATT. The advantage of this approach is that it allows for the estimation of confidence intervals to compare point estimates across subgroup pairs, while still allowing transparent comparison of subgroup effects to the overall ATT.

Like nearly all estimators for treatment effect we use the form $\widehat{\tau} = \sum_{N} \lambda_i \cdot Y_i$ where Y_i is the outcome for unit i and λ_i is a known weight such that $\sum_{i:T_i=1} \lambda_i = 1$, $\sum_{i:T_i=0} \lambda_i = -1$, where T_i is the treatment indicator for unit i.[7] Letting s indicate the subgroup of interest, the subgroup ATT estimator is:

$$\widehat{\tau}^s = \sum_{N} \lambda_i^s \cdot Y_i^s. \tag{1}$$

where

$$Y_i^s = \begin{cases} Y_i^s & \text{if } T_i = 1 \\ \widehat{Y}_i^s = Y_{i:T=0}^s + \widehat{\mu}_0\left(X_{i:T=1}\right) - \widehat{\mu}_0\left(X_{i:T=0}\right) & \text{if } T_i = 0 \end{cases} \tag{2}$$

and $\widehat{\mu}_0(\cdot)$ represents the predicted values obtained from combining the coefficients from a

[7] The simplest example of weights would come from one-to-one matching without replacement. In this case $\lambda_{i:T=1} = -\lambda_{i:T=0} = 1/N_{T=1}$. In general the weight is based upon the estimation strategy (i.e., propensity score weighting, kernal matching etc.). For our purposes $\lambda_{i:T=1} = 1/N_{T=1}$, $\lambda_{i:T=0} = \#C/N_{T=0}$, where $\#C$ is the number of times an observation is used in the control group.

control group regression, of outcome on covariates, with the respective treated and control covariates.[8] Because we are interested in the ATT, our estimator is:

$$\widehat{\tau}^{s} = \sum_{N_{i:T=1}} \lambda_i^s \cdot Y_i^s + \sum_{N_{i:T=0}} \lambda_i^s \cdot \widehat{Y}_i^s. \tag{3}$$

3.2 Variance

Variances for these subgroup ATT estimates are calculated using a general method proposed by Imbens and Wooldridge (2009) which is related to the method proposed by Abadie and Imbens (2006). The method permits heteroskedasticity across treatment arms (protected, unprotected) and covariates. Matches are chosen, based on covariates, *within* treatment arms and the difference in outcome between these matches forms the basis for the variance estimation:

$$\widehat{\sigma}_i^2 (X_i) = (Y_i - Y_l)^2 / 2. \tag{4}$$

Where Y_l is the outcome of the nearest *within* treatment arm neighbor. This conditional variance estimate is then used to estimate the variance for the sample:

$$\widehat{V}(\widehat{\tau}) = \sum_{N} \lambda_i^2 \cdot \widehat{\sigma}_i^2 (X_i). \tag{5}$$

These variance estimates can then be used to form confidence intervals by which the point estimates of the differences between treated and control subgroups can be evaluated.[9]

[8] The imputations are calculated by plugging the covariates $X_{i:T=1}$ and $X_{i:T=0}$ into the vector of coefficients from the regression $Y_{i:T=0} = X_{i:T=0}\beta_0 + \varepsilon$ to obtain $\widehat{\mu}_0 (X_{i:T=1})$ and $\widehat{\mu}_0 (X_{i:T=0})$, repectively.

[9] All ATT point estimates and associated variances were programmed in R v.2.9.1. The code is available from the authors upon request.

3.3 Inference

There are two components of our estimator $\widehat{\tau}^s$, delineated by high baseline levels of the covariates mentioned in Section 2, $\widehat{\tau}^H$, and low baseline levels, $\widehat{\tau}^L$. Protected and unprotected units are assigned to high and low subsets based on an established threshold $\Im$. Assignment to subgroup $s \in [L, H]$ is conducted according to the following rule:

$$s_i = \begin{cases} H & \text{if } x_i > \Im \\ L & \text{otherwise.} \end{cases} \tag{6}$$

Each subgroup pair is composed of units $x_i^{s=H}$ with corresponding estimator $\widehat{\tau}^H$ and units $x_i^{s=L}$ with corresponding estimator $\widehat{\tau}^L$. The estimator $\widehat{\tau}^H$ is therefore calculated by comparing the outcomes of protected and unprotected units for which $x_i^{s=H}$. Similarly, the estimator $\widehat{\tau}^L$ is calculated by comparing protected and unprotected units for which $x_i^{s=L}$. These estimators address how protected units with high baseline levels of a covariate, for instance, would have fared had they not been treated by comparing them to similar unprotected units with high baseline levels of the same covariate.

Greater interest lies in the comparison, within subgroup pairs, of the two components of the subgroup estimator than in the respective point estimates. We want to compare the ATT estimates of high-baseline units to the ATT of low-baseline units for each set of subgroup pairs. Specifically we want to know if $\widehat{\tau}^H \neq \widehat{\tau}^L$, which is an indication of heterogeneous subgroup response to treatment. Let $C^s\left(\widehat{\tau}^s, \widehat{V}\right) = \left[\widehat{\tau}^s - c \cdot \sqrt{\widehat{V}\left(\widehat{\tau}^s\right)}, \widehat{\tau}^s + c \cdot \sqrt{\widehat{V}\left(\widehat{\tau}^s\right)}\right]$ be the 95% confidence interval for subgroup s, where c is the appropriate critical value associated with the normal distribution. Let $\mathbb{C} = C^H \cap C^L$ be the intersection of the high and low-baseline covariate components of C^s. If $\mathbb{C} = \varnothing$ then there is a statistically significant difference between the point estimates of $\widehat{\tau}^H$ and $\widehat{\tau}^L$ within subgroup pairs. In other words, the absence of an intersection between the confidence intervals of two subgroup ATT point estimates provides evidence that the point estimates differ statistically. For instance, suppose that for some baseline covariate the subgroup pair deforestation outcomes have the relationship $\widehat{\tau}^H > \widehat{\tau}^L$ and $\mathbb{C} = \varnothing$. This supposition would indicate that those units with high baseline levels of the covariate exhibited statistically greater amounts of deforestation than those units with low baseline levels of the covariate. Conversely, if in the previous example $\mathbb{C} \neq \varnothing$ we cannot draw any statistically meaningful conclusions regarding heterogeneous treatment effects, in spite of the observed point estimates $\widehat{\tau}^H > \widehat{\tau}^L$.

3.4 Implementation

We begin by creating two counterfactual control groups for the deforestation and socioeconomic subgroup analyses. To ensure comparability, we follow the methods of Andam et al. (2008) and Andam et al. (2010) closely. There are two primary concerns in the formation of the counterfactual groups. The first is comparability across studies. We ensure comparability by drawing counterfactual groups that are similar to those used in previous studies.[10]

Our second concern is the precision of our estimates. Because subgroup analyses require the segmentation of the sample (or population), subgroup treatment effect estimates will generally have less precision than the overall sample (or population) treatment effect

[10] In the deforestation analysis our counterfactual group is slightly different for two reasons. First, we use an updated protected areas spatial layer which differs from the layer used by Andam et al. (2008). Second, we use only a single nearest neighbor match (Andam et al. (2008) uses the two nearest neighbors) because there are negligible gains to precision, whereas bias is minimized using only one match (Imbens and Wooldridge 2009).

estimates. In the deforestation sample, precision is not a concern. There are 2,806 protected land parcels in the sample and an equal number of unprotected parcels. However, because the unit of analysis in the socioeconomic analyses is the census tract, there are far fewer protected units (249) in the sample. Precision decreases when the sample is broken into subsets according to the observable characteristics of interest. To improve precision, we form the socioeconomic counterfactual group by combining propensity score and trimming methods (Imbens 2004; Imbens and Wooldridge 2009). We calculate propensity scores for the entire population of census tracts based on the covariates in Table 1. The population is then trimmed according to Crump et al. (2009) and Imbens and Wooldridge (2009) in order to remove extreme propensity score values which indicate that the units are not good comparison units for the treated sample.[11] After trimming, the remaining sample consists of 231 protected census tracts and 973 unprotected census tracts. By using this alternative method of forming our counterfactual group we face the concern that the estimates of ATT will differ significantly from the estimates obtained by Andam et al. (2010). It can be seen, however, that by using the same bias-adjustment techniques as those used by Andam et al. (2010), the estimated ATT of -1.39 is similar to that of original study. This gives us confidence that the subgroup estimates from this sample are indeed comparable to the average treatment effects from Andam et al. (2010).

To address potential heterogeneous deforestation and socioeconomic response by subgroup, we first break the deforestation and socioeconomic samples into subgroup pairs according to equation (6) using the threshold for each of the pretreatment (baseline) covari-

[11] This trimming method is based on the distribution of propensity scores. The trimmed set $Ţ = Ţ_\alpha = \{x \in X | \alpha \leq p(x) \leq 1 - \alpha\}$ where $p(x)$ is the estimated propensity score and α is the solution to: $\frac{1}{\alpha \cdot (\alpha - 1)} = 2 \cdot \left[\frac{1}{p(X_i) - (1 - p(X_i))} \middle| \alpha < p(X_i) < 1 - \alpha \right]$. The estimate for our set is $\alpha = 0.027$.

ates listed in section 2.2. Estimates of subgroup ATT are made for each subgroup within each of the subgroup pairs according to the methods outlined in section 3.1. This is done for both the deforestation and socioeconomic samples using the same threshold values to define subgroups. Using the same values allows us to compare how similar subgroups respond to protection in terms of deforestation and socioeconomic outcomes.

[Table 2 and Figure 1 about here]

4 Results

Table 2 presents the results. For each subgroup, it presents the average outcome for protected units, the imputed counterfactual values for these units and the ATT_s. Figure 1 graphically presents results from a statistical comparison of subgroup point estimates. Each major column represents a subgroup pair and contains two ATT sub-columns. The height of each bar represents the point estimate of ATT for the specified subgroup. The associated whisker represents the 95% confidence interval for each of point estimate. Figure 1 shows for which characteristics we find evidence of heterogeneous subgroup effects. If the whiskers of the two ATT estimates within a subgroup pair do not overlap, a statistical difference in subgroup effects exists.

4.1 Land Use Capacity

As an indicator of agricultural suitability we find that protected land parcels with high land use capacities display significantly higher levels of avoided deforestation (32.4%) than those with low capacities (9%). This result is consistent with the assumption that agricultural pressure increases the likelihood of deforestation. Table 2 indicates that even though deforestation was higher on protected land parcels with high land use capacity (21% were deforested as compared to 10% of low capacity protected parcels), the expected deforestation in the absence of protection was much higher (54% on high-capacity land as compared to 20% on low-capacity land). However, the results suggest that protection on high-capacity land may have exacerbated poverty (positive rather than negative ATT). In contrast, the poverty reduction impacts on low-capacity lands are quite large.

4.2 Slope

The results also indicate a significant difference in deforestation ATT for high-slope and low-slope land parcels. Avoided deforestation from protection on high-slope forest parcels is estimated to be 1.4%, which is significantly lower than the estimated avoided deforestation of 15.9% on low-slope parcels (these results are qualitatively similar to the estimates of Pfaff et al. 2009). However, as was the case using land use capacity to define subgroups, the impacts of protection on poverty are reversed: poverty alleviation associated with protection is greater on census tracts with high average slopes than those with low average slopes.

The results in 4.1 and 4.2 thus indicate that while the returns to protection in terms of avoided deforestation are higher on land with relatively higher potential returns to agriculture, protection assigned to such land leads to comparatively poorer socioeconomic outcomes.

4.2.1 Distance to Major City

We find that protected land parcels that are located further from one of Costa Rica's three major cities experience significantly higher levels of avoided deforestation (15.3%) than parcels that are closer (5%). These results are counterintuitive when distance to a major city is only viewed as a proxy for market access that increases the returns to agriculture. However, distance to a major city also serves as a measure of land-use law enforcement. There is less enforcement of existing land-use laws the further a land parcel is located from a city. This explanation is consistent with the estimated avoided deforestation values in Table 2: deforestation is higher on both treated and control parcels farther from major cities. The conditional impacts on poverty, however, are the opposite: although protection yields greater avoided deforestation when located farther from cities, it yields higher socioeconomic impacts when located near cities.

4.3 Agricultural Workers

We find a statistical difference in the efficacy of protected areas on deforestation outcomes according to the percentage of agricultural workers that reside in the census tract from which the land parcel is sampled. Avoided deforestation estimates are significantly higher on parcels that fall in census tracts with high percentages of agricultural workers (13.3%) compared to those in census tracts with lower percentages of agricultural workers (4.5%). Such a result is consistent with the conjecture that a higher proportion of agricultural workers in the population serves as a good measure of the amount of agricultural activity within the area, which is correlated with higher returns to avoided deforestation.

We find that census tracts with high percentages of agricultural workers exhibited significantly lower socioeconomic outcomes due to protection (0.008) than did census tracts with low percentages of agricultural workers (-1.802). These results provide evidence consistent with predictions that land restrictions associated with protected areas have a differential effect on agricultural workers (Robalino 2007).

4.4 Poverty

Although we find the point estimates of avoided deforestation due to protection to be higher on land parcels that fall within census tracts with high levels of baseline poverty, the difference between high (11.6%) and low (8%) subgroups is statistically insignificant. So too are the estimates of protections impact on socioeconomic outcomes for these subgroups. The point estimates indicate that protection was more beneficial in areas with high baseline poverty but the confidence intervals for these estimates clearly overlap. Statistical significance aside, the point estimates depict a desirable situation from many planners' perspectives. Although high-poverty areas fared no better, statistically, with protection than low-poverty areas, avoided deforestation and poverty alleviation in high-poverty areas were significantly different from zero. Thus placing protected areas in high-poverty areas can, on average, achieve environmental gains without exacerbating poverty. In fact, the evidence suggests that, if anything, protected areas have alleviated poverty in these areas.

4.5 Robustness to Subgroup Definitions

To define subgroups, we use median values of the relevant covariates (see section 2.2). We test the sensitivity of our results to a +/-10% change in these median threshold values. Our inferences are unchanged in all but two instances. In the analysis of protection's impact on poverty, the difference between subgroups near and far from major cities is no longer statistically significant at the 5% level for either a +10% or -10% change in the threshold value. The difference between subgroups with high and low-sloped land is no longer significant for a 10% increase in the threshold value. The ordinal relationships between the point estimates for each subgroup, however, remain qualitatively the same. In the slope subgroup analysis, the precision of the estimates changes when the threshold is increased because there are relatively few census tracts with a majority of land having very high slopes This problem does not arise when the threshold value is decreased (in fact, the qualitative and statistical relationships are the same using a threshold value that is 50% lower than the one used in our analyses).[12]

We run three additional robustness analyses. In the first two we define the threshold as the 40^{th} and 60^{th} percentile subgroup values. For the third analysis we drop any observation with a covariate value that lies between the 40^{th} and 60^{th} percentile and define the "low" group as any observation below the 40^{th} percentile and the "high" group as any observation above the 60^{th} percentile. The results from each of these analyses are qualitatively similar to the robustness analysis using a +/-10\% change in the median threshold values.

4.6 Unobserved Heterogeneity

Unobserved heterogeneity (hidden bias) is a concern in any non-experimental study. Consistent estimation of the average treatment effect on the treated depends on the untestable assumption that, after conditioning on baseline characteristics, the outcome under the no-

[12] The threshold value of 23% slope to separate the subgroups comes from the median slope of units in the deforestation sample. If one were to instead use the median slope of the census tract in the socioeconomic sample (16%), a relationship similar to that displayed by land use capacity is observed. High-slope areas show relatively high poverty alleviation, whereas low-slope areas are associated with poverty exacerbation.

treatment state is independent of treatment exposure. In our study, if the protected and matched unprotected units differ in some unobservable way that affects deforestation, our estimates will be biased. For example, consider how Andam et al. (2008) measure forest cover: a three-hectare plot is considered forested if its canopy cover was greater than 80%. If forested plots selected for protection systematically were to have more (less) canopy cover than the matched controls, our avoided deforestation estimates would be biased upward (downward). For example, say that mean baseline crown cover was 95% in protected plots and 85% in matched control plots. With similar levels of deforestation on protected and unprotected plots, unprotected forest plots would be more likely to pass the 80% threshold and be declared "deforested."[13]

To test the sensitivity of their results to hidden biases, Andam et al. (2008, 2010) use a sensitivity test recommended by Rosenbaum (2002). For example, in the avoided deforestation study of Andam et al. (2008), the authors examine the possibility that the protected plots may be unobservably less likely to be deforested than their matched controls. They posit the existence of a strong confounding factor that not only affects protection decisions, but also determines whether deforestation is more likely in protected plots or the matched controls. They find that the treatment effect estimate is highly robust to hidden bias: if an unobserved plot attribute caused the odds ratio of protection to differ between protected and unprotected plots by a factor of as much as 2.15, the 99% confidence interval of the estimate would exclude zero.

Of course, a sensitivity test to hidden bias only quantifies and expresses the uncertainty

[13] We thank an anonymous referee for noting this particular potential source of hidden bias.

from hidden bias. It does not dispel that uncertainty.[14] Our study, however, focuses on the ordinal rankings of treatment effect estimates within subgroup pairs rather than on the level of the point estimates themselves. In other words, we are less interested in stating the avoided deforestation is X% in a particular subgroup, and more interested in saying that avoided deforestation in subgroup A is greater than in subgroup B. Unobserved heterogeneity would be a concern in our analyses only if it were to differentially affect the subgroup pairs such that it caused the ordering of subgroup estimates to switch. We cannot think of a simple story of systematic unobserved heterogeneity that would act differentially within subgroup pairs (e.g., on flat lands, decision makers systematically sought out sparse-canopies among forests observably similar on the dimensions we match, and on steep lands they systematically sought out thick-canopies). Thus, even if unobserved heterogeneity were to bias the underlying average treatment effect on the treated estimates of the original samples, it is unlikely to affect our estimated ordering of subgroup pairs.

5 Discussion

Recent studies have found what appears to be evidence of so-called "win-win" outcomes associated with protected areas in Costa Rica. Protection has been moderately effective, on

[14] To directly assess the potential source of bias from not using continuous crown cover data, we would need continuous baseline data, which we lack. However, we obtained such data for the period 1992 -1993 from the Global Land Cover Facility (Earth Science Data Interface). If we assume that any canopy cover bias in decisions to protect forests before 1980 would continue into the early 1990s, we can use these recent data to test whether canopy cover percentages were similar between protected and unprotected plots at baseline. We measure canopy cover inside and outside of protected areas established between 1991 and 1995. These data (measured at 1square km-level) range from 0-80%. Because there is no variation above 80% (the threshold for our binary indicator), we use the next quintile (60-80%). If forest canopy percentage affects selection into protection, we should observe a difference in the mean canopy cover for protected and unprotected units. We do not observe any meaningful difference: mean canopy cover percentage within protected areas is 69.75% and in unprotected areas is 69.6%.

average, in preventing deforestation (Andam et al. 2008) and in alleviating poverty (Andam et al. 2010). However, these impact estimates ignore the potential for heterogeneous responses to protection for different subgroups. Understanding heterogeneous treatment response is important from the perspective of a social planner because conditional assignment of protected areas can lead to greater average treatment response for the population (Manski 2005).

Using new quasi-experimental methods, we estimate the heterogeneous subgroup impacts of protected areas established prior to 1980 on deforestation and socioeconomic outcomes in Costa Rica. For nearly all the biophysical and demographic subgroups we define, we find statistically significant, and policy-relevant, evidence of heterogeneous responses to protected areas. Avoided deforestation is highest when protection is assigned to lands that are highly suitable for agriculture, are far from major cities and infrastructure, or where a high percentage of adults are employed in agriculture: about three times higher than on lands that exhibit the opposite characteristics. However, poverty alleviation is highest when protection is assigned to areas with the opposite characteristics. In other words, the characteristics associated with the *most* avoided deforestation are the characteristics associated with the *least* poverty alleviation.

Caution should be observed when using our results to guide future conservation planning in Costa Rica. We estimated the average treatment effects of protection on protected forests in each subgroup impacts. Thus extrapolation should only be made to areas that are observably similar to the protected ecosystems in this study. Given that the covariates associated with areas already protected are most likely very similar to areas that will be chosen for protection in the future, basing extrapolation on the counterfactual samples used in this study may not be unreasonable. Future analyses, however, should estimate the average treatment effect on the control (ATC) to provide insights into the way in which protection anywhere in Costa Rica that is currently unprotected would affect deforestation and poverty. As noted in Andam et al. (2010), future analyses should also focus on the impacts of alternative management strategies, such as community management (e.g., Somanathan, Prabhakar, and Mehta 2009), and on elucidating the mechanisms through which protection has reduced poverty (e.g., tourism, infrastructure development, ecosystem services). Our analysis provides a useful foundation for such analyses by highlighting the spatially heterogeneous impacts of protection.

Although historical treatment responses do not necessarily predict future ones, our results indicate that prudent conservation planning would pay special attention to covariates related to agriculture. For example, decisionmakers may wish to look at the composition of employment in the surrounding areas before assigning protective legislation to an ecosystem. If protecting ecosystems in areas with a large percentage of adults employed in agriculture cannot be avoided, additional interventions, such as performance payments for environmental services to local communities, may be warranted to contribute to poverty alleviation goals.

One of the goals set forth at the Fifth World Parks Congress in 2003 is that protected areas should do no economic harm to surrounding human populations (Adams et al. 2004). The results to date indicate that, on average, Costa Rica's protected area system achieved this goal. Equally important, the results support claims that protecting ecosystems in high-poverty areas can, on average, achieve environmental gains and alleviate poverty. Yet the amount of avoided deforestation generated by Costa Rica's protected area system was modest. As in other nations, Costa Rican protected areas tend to be assigned to ecosystems with low economic returns from conversion.[15] Our study shows that the same factors that have limited the conservation effectiveness of protected areas may have improved the social welfare impacts of these areas. This observation implies that 'win-win' efforts to protect ecosystems and alleviate poverty may be possible when policymakers are satisfied with low levels of each outcome, but tradeoffs exist when more of either outcome is desired. Without innovations in conservation technology, having more of one will imply having less of the other.

Acknowledgements: The authors thank the sources for the data used in Andam et al. 2008 and 2010: Arturo Sanchez-Azofeifa and the Earth Observation Systems Laboratory of the University of Alberta for the forest cover layers; Douglas Guell Vargas, Roger Gutierrez Moraga, Allan Ramirez Villalobos, and Hugo Lopz Arauz from the Instituto Nacional de Estadsticas y Censos (INEC) for digitizing the Costa Rica census segment GIS; Patricia Solano and Maria Elena Gonzalez (INEC) for providing access to the Costa Rica census data; Margaret Holland for calculating the poverty indices, cleaning the census boundary layers, and for sharing the protected area data layer; and Juan Robalino for sharing the slope digital elevation model and, with Alexander Pfaff, other data layers. We also thank Kwaw Andam for sharing the data samples from the original studies and for assistance in interpreting the data files.

[15] The Millennium Ecosystem Assessment (2005, pp. 130) reports that "many protected areas were specifically chosen because they were not suitable for human use."

6 Appendix

6.1 Areal Interpolation

Costa Rica's census tract boundaries are not spatially consistent across time. The number of census tracts increases from 4,694 in 1973 to 17,625 in 2000. Furthermore, the addition of census tracts over time did not follow any discernible pattern, the newer subdivided census tracts do not necessarily fall within the boundaries of the old census tracts. This poses a problem for the comparability of the demographic data over time. In order to make the 2000 data comparable to the 1973 data, the geographic method of Areal Interpolation (Reibel 2007) is implemented.

Areal interpolation is a GIS method by which demographic variables are made comparable across time given changes in political boundaries. For our analyses the 1973 census tracts are used as baselines. Therefore, areal interpolation assigns weights (assuming a uniform population distribution) based upon the amount that the 2000 census tracts overlap with the 1973 census tracts. These weights are used to interpolate the 2000 populations that reside within the 1973 census tract boundaries. The resulting data set contains the original 1973 demographic data according to its native boundaries and the 2000 demographic data distributed as if the census tract boundaries had not changed since 1973.

6.2 Poverty Index

Costa Rica does not have properly disaggregated income data that date back to 1973 (Gindling and Terrell 2004). To measure the socioeconomic impacts of protected areas an alternative metric is necessary. Cavatassi et al. (2004) suggest the use of principal components analysis to form a poverty index. The method uses indicators from the respective censuses that are believed to affect poverty to create a measure that is spatially and temporally comparable. The variables included in the poverty index are: (* indicates a percentage): *men in total population*, families who cook with coal or wood*, families without washing machine*, families without refrigerator*, people who are employed and get a salary as job remuneration*, illiterate population aged 12 or more*, household dwellings without connection to private or public water system*, household dwellings without sewers*, household dwellings without electricity*, household dwellings without telephone*, dwellings with earth floor*, dwellings in bad condition*, dwellings without bathroom*, dwellings without access to hot water*, dependency ratio, average number of occupants per bedroom, average years of education per adult.* A similar measure was employed by the Mexican government in the analysis of the PROGRESA program (Cavatassi et al. 2004).

References

[1] Abadie A, Drukker D, Herr J, Imbens G (2004) Implementing matching estimators for average treatment effects in Stata. SJ 4: 290-311

[2] Abadie A, Imbens G (2006) Large sample properties of matching estimators for average treatment effects. Econometrica 74(1): 235-267

[3] Adams W, Aveling R, Brockington D, Dickson B, Elliott J, Hutton J, Roe D, Vira B, Wolmer W (2004) Biodiversity conservation and the eradication of poverty. Science 306(5699): 1146-1149

[4] Andam KS, Ferraro PJ, Pfaff A, Sanchez-Azofeifa G, Robalino JA (2008) Measuring the effectiveness of protected area networks in reducing deforestation. Proc Natl Acad Sci USA 105(42): 16089-16094

[5] Andam K, Ferraro P, Sims K, Healy A, Holland M (2010) Protected areas reduced poverty in Costa Rica and Thailand. Proc Natl Acad Sci USA 107(22): 9996

[6] Cavatassi R, Davis B, Lipper L. 2004. Estimating poverty over time and space: construction of a time-variant poverty index for Costa Rica. ESA Working Paper

[7] Coad L, Campbell A, Miles L, Humphries K (2008) The Costs and Benefits of Protected Areas for Local Livelihoods: a review of the current literature. World Conservation Monitoring Centre Working Paper, UNEP

[8] Cropper M, Puri J, Griffiths C (2001) Predicting the location of deforestation: The role of roads and protected areas in North Thailand. Land Econ 77(2): 172

[9] Crump R, Hotz V, Imbens G, Mitnik O (2008) Nonparametric tests for treatment effect heterogeneity. Rev Econ Statist 90(3): 389-405

[10] _________ (2009) Dealing with limited overlap in estimation of average treatment effects. Biometrika 96(1): 187

[11] de Camino Velozo R, de Camino R, Segura O, Arias L, Perez I (2000) Costa Rica: forest strategy and the evolution of land use. World Bank Publications, Washington, D.C.

[12] Duffy-Deno K (1998) The effect of federal wilderness on county growth in the intermountain western United States. J Reg Sci 38(1): 109-136

[13] Gindling T, Terrell K (2004) Legal minimum wages and the wages of formal and informal sector workers in Costa Rica. William Davidson Institute Working Papers Series, University of Michigan

[14] Ho D, Imai K, King G, Stuart E (2007) Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Polit Anal 15: 199-236

[15] Holland P (1986) Statistics and causal inference. J Amer Statistical Assoc 81(396): 945-960

[16] Imbens G (2004) Nonparametric estimation of average treatment effects under exogeneity: A review. Rev Econ Statist 86(1): 4-29

[17] Imbens G, Wooldridge J (2009) Recent developments in the econometrics of program evaluation. J Econ Lit 47(1): 5-86

[18] Joppa L, Pfaff A (2010) Reassessing the forest impacts of protection: The challenge of nonrandom location and a corrective measure. Ann N Y Acad Sci 1185: 135-139

[19] Lewis D, Hunt G, Plantinga A (2002) Public conservation land and employment growth in the northern forest region. Land Econ 78(2): 245

[20] _________ (2003) Does public lands policy affect local wage growth? Growth Change 34(1): 64-86

[21] Manski C (2005) Social choice with partial knowledge of treatment response. Springer, New York

[22] MEA (2005) Ecosystems and human well-being. Millennium Ecosystem Assessment. Island Press

[23] Pfaff A, Robalino J, Sanchez-Azofeifa G, Andam K, Ferraro P (2009) Park Location Affects Forest Protection: Land Characteristics Cause Differences in Park Impacts across Costa Rica. BE J Econ Anal & Pol 9(2): 5

[24] Reibel M (2007) Geographic information systems and spatial data processing in demography: A review. Population Res Pol Rev 26(5): 601-618

[25] Robalino J (2007) Land conservation policies and income distribution: who bears the burden of our environmental efforts? Environ Devel Econ 12(04): 521-533

[26] Rubin B, Hyman G (2000) The Extent and Economic Impacts of Soil Erosion in Costa Rica. In: C.A.S Hall, C Perez and G Leclerc (ed) Quantifying Sustainable Development: The Future of Tropical Economies. Academic Press

[27] Sims KRE (2010) Conservation and Development: Evidence from Thai Protected Areas. J Environ Econ Manage Forthcoming

[28] Somanathan E, Prabhakar R, Mehta B (2009) Decentralization for cost-effective conservation. Proc Natl Acad Sci USA 106(11): 4143

[29] Wilkie D, Morelli G, Demmer J, Starkey M, Telfer P, Steil M (2006) Parks and People: Assessing the Human Welfare Effects of Establishing Protected Areas for Biodiversity Conservation. Conserv Biol 20(1): 247

[30] WPDA (2009) World Database on Protected Areas: 2009 Annual Release http://www.wdpa.org/AnnualRelease.aspx

	Variable	Description	Mean	Median	Standard Deviation	Range
Deforestation Covariates	High Productivity Land	Land Use Capacity I, II & III Land suitable for agricultural production. May require special land and crop management (classes II & III).	0.008	0	0.09	0-1
	Medium-High Productivity Land	Land Use Capacity IV Moderately suitable for agricultural production; permanent of semi-permanent crops	0.0289	0	0.167	0-1
	Medium-Low Productivity Land	Land Use Capacity V, VI & VII Strong limiting factors on agricultural production.	0.0802	0	0.272	0-1
	Distance to Forest Edge	Distance (km) to the edge of the forest in 1960	2.79	2.35	2.19	0.0001-11.2
	Distance to Road	Distance (km) to nearest road in 1969.	16.99	14.28	11.62	0.04-53.31
	Distance to Major City	Distance (km) to nearest major city: Limon, Puntarenas or San Jose.	77.4	56.9	49.53	9-180.5
Socioeconomic Covariates	Baseline Poverty	Poverty index measured in 1973.	14.9	15.8	6.43	-6.4-28.9
	Forest Cover	Percentage of census tract occupied by forest in 1960.	0.412	0.383	0.342	0-1
	% High Productivity Land	Percent of census tract occupied by Land Use Capacity I, II & III land.	0.118	0	0.22	0-1
	%Medium-High Productivity Land	Percent of census tract occupied by Land Use Capacity IV land.	0.295	0.04	0.377	0-1
	%Medium-Low Productivity Land	Percent of census tract occupied by Land Use Capacity VI, VII or VIII land.	0.347	0.156	0.387	0-1
	Distance to Major City	Average distance (km) from each 300m^2 land plot within a census tract to nearest major city: Limon, Puntarenas or San Jose.	57.3	49.7	41.28	0.0037-208
	Roadless Volume	The sum of the product of area and distance to nearest road (1969) for every square with side length 100m within the census tract.	308,000	66,400	699,100	0.28-7,590,000

Table 1. Summary statistics and description of covariates used as controls to form counterfactual samples.

		Deforestation						Socioeconomic					
		High Baseline Levels			Low Baseline Levels			High Baseline Levels			Low Baseline Levels		
Subgroup Pair	Threshold	$Y_{T=1}$	$\hat{Y}_{T=0}$	$\tau^{s=H}$	$Y_{T=1}$	$\hat{Y}_{T=0}$	$\tau^{s=L}$	$Y_{T=1}$	$\hat{Y}_{T=0}$	$\tau^{s=H}$	$Y_{T=1}$	$\hat{Y}_{T=0}$	$\tau^{s=L}$
Land Use Capacity	High	0.212	0.535	-0.324	0.108	0.202	-0.094	1.62	0.003	1.617	-2.22	-0.528	-1.693
		[104]	[104]	(0.077)	[2702]	[2702]	(0.017)	[22]	[301]	(0.663)	[209]	[672]	(0.359)
Slope	23%	0.098	0.112	-0.014	0.132	0.291	-0.159	-3.9	-2.3	-1.62	1.03	1.25	-0.228
		[1624]	[1133]	(0.023)	[1139]	[1656]	(0.019)	[135]	[284]	(0.244)	[96]	[689]	(0.301)
Distance To Major City	57km	0.141	0.294	-0.153	0.081	0.131	-0.05	2.86	2.81	0.053	-3.82	-2.58	-1.247
		[1418]	[1377]	(0.016)	[1388]	[1429]	(0.015)	[67]	[298]	(0.511)	[164]	[675]	(0.223)
%Agricultural Workers	13%	0.107	0.24	-0.133	0.119	0.164	-0.045	-1.41	-1.41	0.008	-2.45	-0.643	-1.802
		[1660]	[1676]	(0.019)	[1146]	[1130]	(.019)	[131]	[487]	(0.267)	[100]	[486]	(0.335)
Initial Poverty	15	0.123	0.239	-0.116	0.082	0.162	-0.08	0.968	2.06	-1.088	-4.51	-3.83	-0.684
		[2002]	[2002]	(0.018)	[804]	[804]	(0.024)	[112]	[564]	(0.301)	[119]	[409]	(0.208)

Notes: Y denotes the outcome (deforestation, poverty index), $T=1$ denotes protected units; $T=0$ denotes matched unprotected units.

$\hat{Y}_{T=0}$ is imputed according to equation (2).

τ^{s} is the subgroup ATT calculated, $\tau = Y_{T=1} - \hat{Y}_{T=0}$.

[Number of Obsaervations in Subgroup]

(Standard Errors)

Table 2. Estimated average treatment effect on the treated (ATT) by subgroup pair for covariates listed in section 2.2.

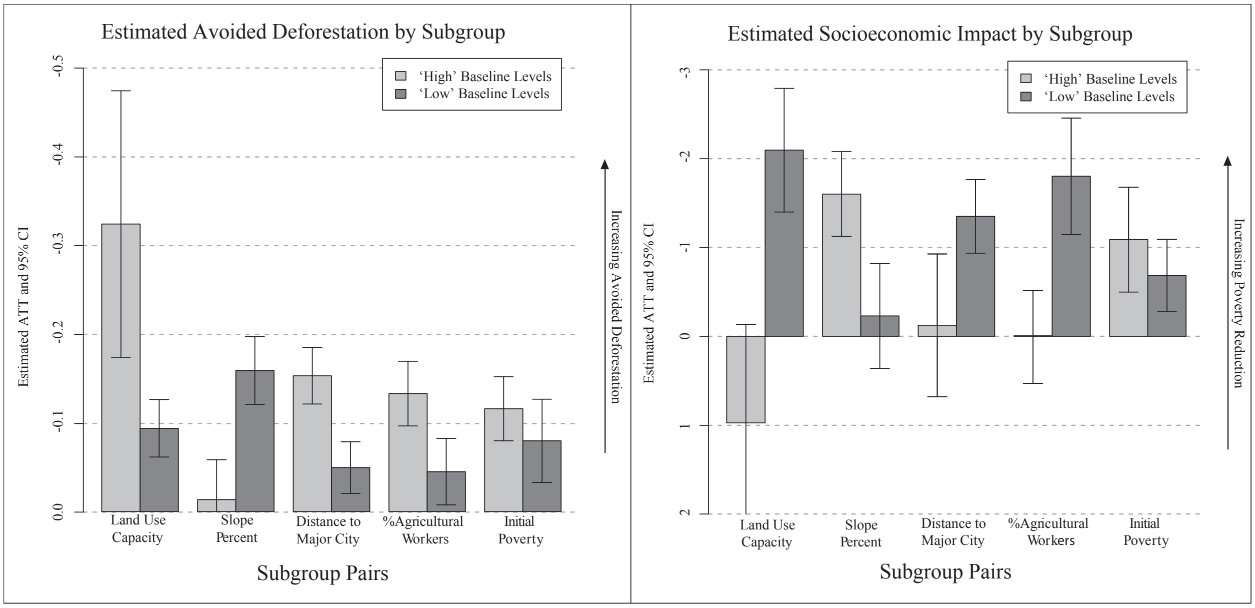

Figure 1. Estimated heterogeneous impacts of protection on avoided deforestation and poverty.