Spring 5-7-2011

Essays on the Evaluation of Environmental Programs

Merlin M. Hanauer
Georgia State University

PERMISSION TO BORROW

In presenting this dissertation as a partial fulfillment of the requirements for an advanced degree from Georgia State University, I agree that the Library of the University shall make it available for inspection and circulation in accordance with its regulations governing materials of this type. I agree that permission to quote from, to copy from, or to publish this dissertation may be granted by the author or, in his or her absence, the professor under whose direction it was written or, in his or her absence, by the Dean of the Andrew Young School of Policy Studies. Such quoting, copying, or publishing must be solely for scholarly purposes and must not involve potential financial gain. It is understood that any copying from or publication of this dissertation which involves potential gain will not be allowed without written permission of the author.

Signature of Author

NOTICE TO BORROWERS

All dissertations deposited in the Georgia State University Library must be used only in accordance with the stipulations prescribed by the author in the preceding statement.

The author of this dissertation is:
Merlin Mack Hanauer
Box 515
Forks of Salmon, CA 96031

The director of the dissertation is:
Dr. Paul J. Ferraro
Andrew Young School of Policy Studies
Georgia State University
P. O. Box 3992
Atlanta, GA 30302-3992

Users of this dissertation not regularly enrolled as students at Georgia State University are required to attest acceptance of the preceding stipulations by signing below. Libraries borrowing this dissertation for the use of their patrons are required to see that each user records here the information requested.

Name of User	Address	Date	Type of use (Examination only or copying)

ESSAYS ON THE EVALUATION OF ENVIRONMENTAL PROGRAMS
BY
MERLIN MACK HANAUER

A Dissertation Submitted in Partial Fulfillment
of the Requirements for the Degree
of
Doctor of Philosophy
in the
Andrew Young School of Policy Studies
of
Georgia State University

GEORGIA STATE UNIVERSITY
2011

ACCEPTANCE

This dissertation was prepared under the direction of the candidate's Dissertation Committee. It has been approved and accepted by all members of that committee, and it has been accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Economics in the Andrew Young School of Policy Studies of Georgia State University.

Dissertation Chair:	Dr. Paul J. Ferraro
Committee:	Dr. H. Spencer Banzhaf
	Dr. Andrew R. Hanson
	Dr. Laura O. Taylor

Electronic Version Approved:
Mary Beth Walker, Dean
Andrew Young School of Policy Studies
Georgia State University
April 2011

ACKNOWLEDGMENTS

Completion of this dissertation was made possible by the direct and indirect support of many people. Indeed, I owe a debt of gratitude to all of the friends, family members and teachers that have facilitated my personal and academic growth over the years.

A vast majority of my professional development during my graduate career owes to the guidance of Paul Ferraro. As his research assistant and coauthor, I have gained tools that have greatly enhanced my ability to conduct and produce defensible research. I have probably benefited most profoundly from Paul's impeccable intuition, economic and otherwise. He has the ability to quickly capture the essence of research and asses the importance, failings, potential biases, etc. My research and intuition are the beneficiaries of the countless discussions in his office. I can say with a great deal of certainty that I would not be in the position that I am – beyond the completion of this dissertation – without Paul and for that I, and my family, thank him wholeheartedly.

I thank Spencer Banzhaf for his willing support and encouragement. He is a great thinker and teacher, with a human touch. My research benefited importantly from his suggestions. Others that have provided helpful comments and assistance on my research include Jim Cox, Kurt Schnier, Andrew Hanson, Laura Taylor, and participants in the classes, workshops, seminars, and conferences in which I have presented my work. I thank Felix Rioja who was an advocate for me in the early stages and facilitated a nice distraction, on the gridiron, from the floggings of graduate school.

A great deal of my happiness during the past five years owes to luck. In my cohort I have made friends that will remain for life. I thank Andrew Balthrop, Gustavo Canavire-Bacarezza, Zack Hawley and Kelly Wilkin, in particular. From studying for the comps to discussing econometrics over spirits to Friday volleyball, things would have been much more difficult, and much less fun, were it not for their presence.

I thank my parents – Gail, Betty Ann, Creek and Gerry – for their unfailing love and support. Although my trajectory has varied over the years, I have always received the loving encouragement and assurance necessary for success. I thank Jim and Mary Jane Schilz as well. For many years I have seen them as my third set of parents, and for that I count myself lucky.

The person that has played the most integral role in supporting and motivating me through graduate school and life, is my wife. Erin, I thank you for the countless facets of support that you provide, both tangible and intangible. You make me a better person and have enriched my life in so many ways – I love you. To my children, Lyla and Lars, you will not remember all the joy that you have brought me throughout this process, but I will.

This Dissertation is Dedicated to my Mother

Contents

List of Figures

Figure **Page**

List of Tables

ABSTRACT

ESSAYS ON THE EVALUATION OF ENVIRONMENTAL PROGRAMS

BY

MERLIN MACK HANAUER

April 2011

Committee Chair: Dr. Paul J. Ferraro

Major Department: Economics

This dissertation comprises four chapters. The unifying theme is the evaluation of environmental programs. Specifically, each chapter examines some facet of the impacts of protected areas.

The first chapter examines the heterogeneous environmental and economic impacts of protected areas in Costa Rica. Previous studies 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.

The second chapter explores in more detail the heterogeneous impacts of protected areas in Costa Rica and Thailand. In particular we investigate the potential for protected areas to act as a mechanism for poverty traps and use semiparametric models to identify the spatial congruence of environmental and economic outcomes. We find no evidence that protected areas trap historically poorer areas in poverty. In fact, we find that poorer areas at baseline appear to have

the greatest levels of poverty reduction as a result of protection. However, we do find that the spatial characteristics associated with the most poverty alleviation are not necessarily the characteristics associated with the most avoided deforestation. We demonstrate how an understanding of these spatially heterogeneous responses to protection can be used to generate suitability maps that identify locations in which both environmental and poverty alleviation goals are most likely to be achieved.

In the third chapter we address the mechanisms through which protected areas affect economic outcomes. Using recently developed quasi-experimental methods and rich biophysical and demographic data, we quantify the causal post-treatment mechanism impacts of tourism, infrastructure development and ecosystem services on poverty, due to the establishment of protected areas in Costa Rica prior to 1980. We find that nearly 50% of the poverty reduction estimated in a previous study can be attributed to tourism. In addition, although the mechanism estimates for the infrastructure and ecosystem services proxies are negligible, we argue that the results provide evidence that enhanced ecosystem services from the establishment of protected areas has likely helped to reduce poverty. The results provide additional information to policy makers that wish to enhance the future establishment of protected areas with complementary policy.

The final chapter studies the economic impacts of protected areas in Bolivia. We find that municipalities with at least 10% of their area occupied by a protected area between 1992 and 2000 exhibited differentially greater levels of poverty reduction between 1992 and 2001 compared to similar municipalities unaffected by protected areas. We find that the results are robust to a number of econometric specifications, spillover analyses and a placebo study. Although the overarching results that Bolivia's protected areas were associated with poverty reduction are similar to previous studies , the underlying results are subtly, but significantly, different. In previous studies it was found that controlling for key observable covariates lead to

fundamentally antithetical results compared to naïve estimates. Conversely, these results indicate that naïve estimates lead to an over-estimation of the poverty reducing impacts of protected areas. The results expose the heterogeneity of protected area impacts across countries and, therefore, underscore the importance of country-level impact evaluations in order to build the global knowledge base regarding the socioeconomic impacts of protected areas.

Chapter I

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

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 (WDPA 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 (Duffy-Deno 1998, Lewis et al. 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.

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

[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.

poverty between 1973 and 2000. Similar to Andam et al. (2008), the authors use matching techniques to form 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.

Data

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 pre-dates 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 A) 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 by protected land (results are robust to changes in this threshold definition).[3] Conversely, any census tract

[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`

[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%).

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

Variable	Description	Mean	Median	Standard Deviation	Range
Deforestation Covariates					
High Productivity Land	Land Use Capacity I, II or 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 or 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	Linear 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	
% High Productivity Land	Percent of census tract occupied by Land Use Capacity I, II or 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 (km)	Average linear distance 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

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 et al. (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 A 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.

Subgroup Variables

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

[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.

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

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

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 et al. 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).

Methods

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

[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.

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 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 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

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

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

[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.

[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 $\widehat{\mu}_0\left(X_{i:T=1}\right)$ and $\widehat{\mu}_0\left(X_{i:T=0}\right)$, repectively.

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

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

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

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

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]

Inference

There are two components of our estimator $\widehat{\tau}^s$, delineated by high baseline levels of the covariates mentioned in the Data Section, $\widehat{\tau}^H$, 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

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

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

[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.

$x_i^{s=H}$. 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 $\widehat{\tau}^H > \widehat{\tau}^L$.

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 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

[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).

[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

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) covariates listed in the Subgroup Variable Section. Estimates of subgroup ATT are made for each subgroup within each of the subgroup pairs according to the methods outlined in the Estimator Section. 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.

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

---

solution to: $\frac{1}{\alpha\cdot(\alpha-1)} = 2 \cdot \left[\frac{1}{p(X_i)-(1-p(X_i))} \middle| \alpha < p\left(X_i\right) < 1-\alpha\right]$. The estimate for our set is $\alpha = 0.027$.

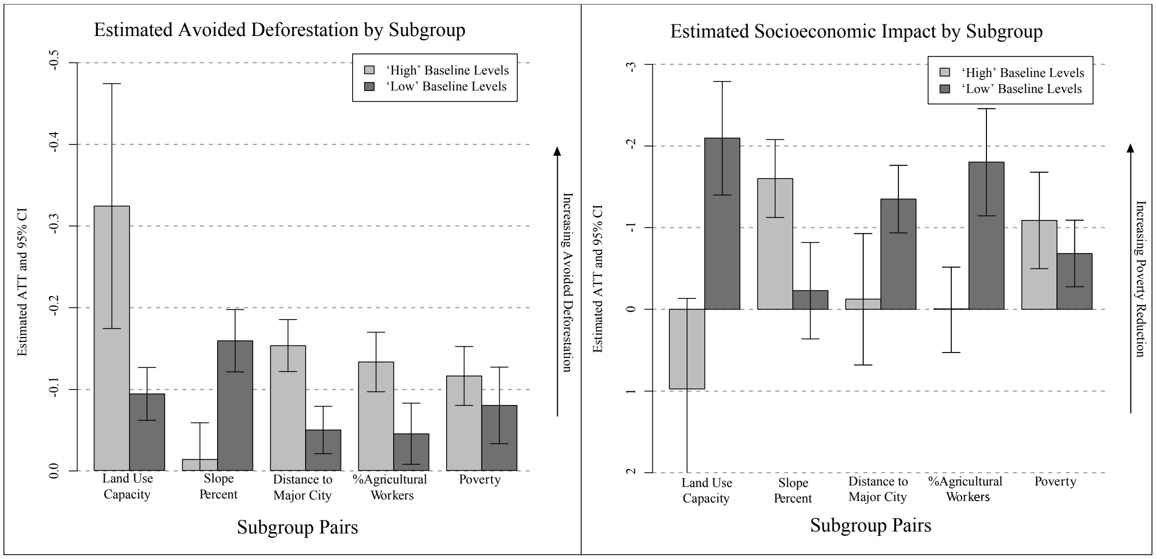

Figure 1: Estimated heterogeneous impacts of protection on avoided deforestation and poverty.

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.

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.

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 I and I 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.

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

Table 2: Estimated average treatment effect on the treated (ATT) by subgroup pair.

		Deforestation						Socioeconomic					
		High Baseline Levels			Low Baseline Levels			High Baseline Levels			Low Baseline Levels		
Subgroup Pair	Threshold	$Y_{T=1}$	$\widehat{Y}_{T=0}$	$\tau^{s=H}$	$Y_{T=1}$	$\widehat{Y}_{T=0}$	$\tau^{s=L}$	$Y_{T=1}$	$\widehat{Y}_{T=0}$	$\tau^{s=H}$	$Y_{T=1}$	$\widehat{Y}_{T=0}$	$\tau^{s=L}$
Land Use	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
Capacity		[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	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
Major City		[1418]	[1377]	(0.016)	[1388]	[1429]	(0.015)	[67]	[298]	(0.511)	[164]	[675]	(0.223)
%Agricultural	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
Workers		[1660]	[1676]	(0.019)	[1146]	[1130]	(.019)	[131]	[487]	(0.267)	[100]	[486]	(0.335)
Baseline	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
Poverty		[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.
$\widehat{Y}_{T=0}$ s imputed according to equation (2).
τ^s is the subgroup ATT calculated, $\tau = Y_{T=1} - \widehat{Y}_{T=0}$.
[Number of Observations in Subgroup]
(Standard Errors)

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.

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).

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.

Robustness to Subgroup Definitions

To define subgroups, we use median values of the relevant covariates (see the Subgroup Variable Section). 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]

[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.

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.

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-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

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

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 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

[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%.

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.

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 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 et al. (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.

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

Chapter II

Conditions Associated with Protected Area Success in Conservation and Poverty Reduction

Introduction

Protected areas are the dominant approach to protecting biodiversity and the supply of ecosystem services (MEA 2005). A fundamental concern surrounding the establishment of protected areas, particularly in developing countries, is that ecosystem conservation goals may conflict with poverty alleviation goals by reducing incomes or perpetuating poverty traps (Adams et al. 2004, Coad et al. 2008, Wilkie et al. 2006, WDPA 2009, Brockington et al. 2006). A poverty trap, as described in the introduction to this special issue, is a self-reinforcing mechanism that causes an area to remain poor. By restricting access to natural resources, protected areas might create new poverty traps or reinforce old ones.[16] Protected areas tend to be established away from major cities and on agriculturally undesirable land (Joppa and Pfaff 2009); characteristics also associated with high levels of poverty. We might therefore be concerned that protected areas would reinforce poverty traps. More optimistically, they might push local economies out of poverty traps by providing tourism business opportunities, improved infrastructure, or enhanced supplies of ecosystem services. For example, new evidence from Costa Rica and

[16]For example Robalino (2007) predicts that protected areas would place a greater burden on non-landowning workers, who are often the poor.

Thailand suggests that protected areas in these two countries have, on average, reduced local poverty (Andam et al. 2010, Sims 2010).

To fully understand protected area impacts, one should consider environmental and socioeconomic outcomes jointly and quantify the heterogeneity in impacts. Unfortunately, there is little scientific evidence on the nature of this heterogeneity or of the potential tradeoffs between environmental and socioeconomic outcomes (Coad et al. 2008, Joppa and Pfaff 2010). Retrospective causal analysis of the socioeconomic impacts of developing country protected areas is limited (Brockington et al. 2006, Andam et al. 2010, Sims 2010, Ferraro and Hanauer 2011, Bandyopadhyay and Tembo 2010). Only the work in Thailand and Costa Rica (Andam et al. 2010, Sims 2010, Ferraro and Hanauer 2011) also collects information on environmental outcomes. However, those previous studies do not include sufficiently detailed analysis of heterogeneity in impacts to assess potential tradeoffs between ecosystem protection and poverty alleviation (Manski 2005, Crump et al. 2008).

Using data from Costa Rica and Thailand, we examine the heterogeneity of protected area impacts as a function of baseline poverty and covariates that are likely to moderate how protection affects outcomes (Baron and Kenny 1986). We select these two nations because they have significant biodiversity, large protected area systems and reliable spatially explicit data. Unlike previous studies that explore heterogeneous impacts of protected areas (Sims 2010, Ferraro and Hanauer 2011, Pfaff et al. 2009), we examine impacts on both avoided deforestation and poverty reduction and use a nonparametric method of locally weighted scatterplot smoothing (LOESS) (Cleveland 1979, Cleveland and Devlin 1988) and a semiparametric partial linear differencing model (PLM) (Yatchew 1997, 1998). These models estimate more informative continuous relationships between observable characteristics and outcomes. We are thus able to identify covariate

ranges that are associated with high conservation and poverty reduction outcomes ('win-win'), low conservation and poverty exacerbation outcomes ('lose-lose'), or incongruence where one outcome is 'win' and the other is 'lose' ('win-lose').

The rapidly growing conservation planning literature focuses on how to target conservation investments conditional on observable environmental and economic characteristics (Margules and Pressey 2000, Naidoo et al. 2006). Planners interested in achieving both avoided deforestation and poverty reduction need to understand how these outcomes co-vary with observable characteristics. Such understanding allows for the development of conditional empirical success rules (see Manski (2005), p.75) that can be used to target interventions based on expected impacts as predicted by observable characteristics. We demonstrate how such rules can be visualized through suitability maps that identify locations associated with 'win-win', 'lose-lose', or 'win-lose' scenarios.

Data

For additional details on data, see Appendix B and Andam et al. (2008, 2010). Previous studies estimated that protected areas resulted in significant avoided deforestation and poverty reduction in Costa Rica and Thailand (Andam et al. 2010, Sims 2010, Andam et al. 2008). About 11% of the area protected in Costa Rica would have been deforested had it not been protected (25). Using similar methods, we estimate that about 15% of protected forest in Thailand would have been deforested in the absence of protection (see Appendix B). Protected areas in Costa Rica accounted for about 10% of the poverty decline around the areas. In Thailand, protected areas reduced poverty by about 30% (Andam et al. 2010). We use data from these studies to explore the heterogeneity of protected areas' impacts.

Poverty

Poverty measures are based on national census data of household characteristics and assets (see Appendix B for detail). Costa Rica analyses use 1973 and 2000 census tract poverty indices (Andam et al. 2010) from a principal components analysis ((Cavatassi et al. 2004); see Appendix B). Thailand analyses use the subdistrict poverty headcount ratio, which is the share of the population in 2000 with monthly household consumption below the poverty line and comes from a poverty mapping analysis (Healy and Jitsuchon 2007, Elbers et al. 2003). The sample comprises subdistricts in north and northeast Thailand, which is where the majority of protected forest areas are located. Larger values of both poverty measures imply greater poverty.

As in Andam et al. (2010), we define a census tract or subdistrict as protected if at least 10% of its area is protected prior to 1980 (Costa Rica) or 1985 (Thailand) (249 census tracts and 192 subdistricts).[17] With protection assigned 15 or more years before poverty outcomes are measured, longer-term impacts can be measured. Unprotected units, from which matched controls are selected, comprise units with less than 1% protected before 1980 or 1985 (4,164 census tracts and 3,479 subdistricts).[18] Protected areas comprise IUCN Categories I, II, IV and VI in Costa Rica and IUCN Categories I and II in Thailand.

Avoided Deforestation

As a proxy for conservation success, we estimate avoided deforestation from protected areas (we acknowledge this is not the only possible measure of success).

[17] Andam et al. (2010) select a 10% threshold because it reflects the call by the fourth World Congress on National Parks and Protected Areas to protect 10% of each of the world's major biomes by 2000, and by the Conference of Parties to the Convention on Biological Diversity to conserve 10% of each of the world's ecoregions. Andam et al. (2010) show that the estimated impacts are robust to changes in this threshold.

[18] Units with one to ten percent of their area protected are dropped from the analysis to avoid matching protected units to "marginally" protected units.

The unit of analysis for the deforestation data is a 3 hectare land parcel (20,000 randomly selected) drawn from forested areas at baseline (Costa Rica, 1960; Thailand, 1973). Each parcel is classified as deforested or forested by the end year (Costa Rica, 1997; Thailand, 2000). A parcel is defined as protected if it lies within a protected area that was established prior to 1980 (Costa Rica) or by 1985 (Thailand). Control parcels were never protected.

Covariates

For each country, multiple spatial layers are used to create covariates for each census tract, subdistrict or parcel (Tables 1 and 14).

Study Design

To estimate the impact of protection on the protected units, one must establish what would have happened in the absence of protection. Like the studies from which we obtain our data (Andam et al. 2010, 2008), we using matching to select unprotected control units that are similar at baseline to protected units. Preprocessing the data (Ho et al. 2007) through matching ensures that the distributions of key covariates believed to affect both outcome and selection into protection are balanced across protected and unprotected units (see Appendix B). The goal of matching, like standard regression techniques, is to control for differences in baseline characteristics that affect the designation of protected areas and poverty or deforestation (Ho et al. 2007, Imbens and Wooldridge 2009, Angrist and Pischke 2009). For example, protected areas are often placed on land less suited for agriculture (MEA 2005, Joppa and Pfaff 2009, Pfaff et al. 2009). The matching strategy assumes that, after matching, the expected outcomes of protected and matched control units in the absence of protection are the same. Thus the control group's outcome represents the protected group's counterfactual outcome. Although

there is no direct way to test this assumption, the previous studies in Costa Rica and Thailand found the estimates were robust to unobserved heterogeneity using our matching specifications (Andam et al. 2010, 2008). The Thailand results were also confirmed using an instrumental variable approach (Sims 2010). For the Costa Rica poverty sample and both deforestation samples, we use nearest neighbor Mahalanobis matching with replacement. For the Thailand poverty sample, we use propensity score matching with exact matching on district to control for baseline fixed effects. See Appendix B for details on the matching methods used and the covariates on which units are matched.

Post-matching, we use nonparametric LOESS (Cleveland 1979, Cleveland and Devlin 1988, Nelson and Chomitz 2009)to estimate impacts as a function of baseline poverty. LOESS allows us to assess whether or not protected areas contributed to poverty traps. We use LOESS because we are interested in how poor areas, including all the factors that make them poor, respond to protection.[19] To isolate the moderating effects on avoided deforestation and poverty from observable baseline characteristics net of other influences, we use semiparametric PLM (Cleveland 1979, Cleveland and Devlin 1988, Yatchew 1997, 1998) on the matched data. This two-stage estimator allows us to linearly control for other influencing covariates in the first stage and then estimate the outcome as a nonparametric function of the covariate of interest using LOESS in the second stage (see Appendix B). The benefit of this approach is that it allows us to conduct inference along a continuum of covariate values (e.g., distances from cities) while holding constant potentially complimentary or countervailing covariates (e.g., slope). The results from the PLMs are then used in the suitability mapping exercise (see Appendix B for more details).[20]

[19]See Appendix B for additional discussion of the choice of methods.

[20]As done in the studies from which we draw (Andam et al. 2010, 2008), we implement bias-adjustment techniques within all LOESS iterations (Imbens and Wooldridge 2009, Abadie et al. 2004).

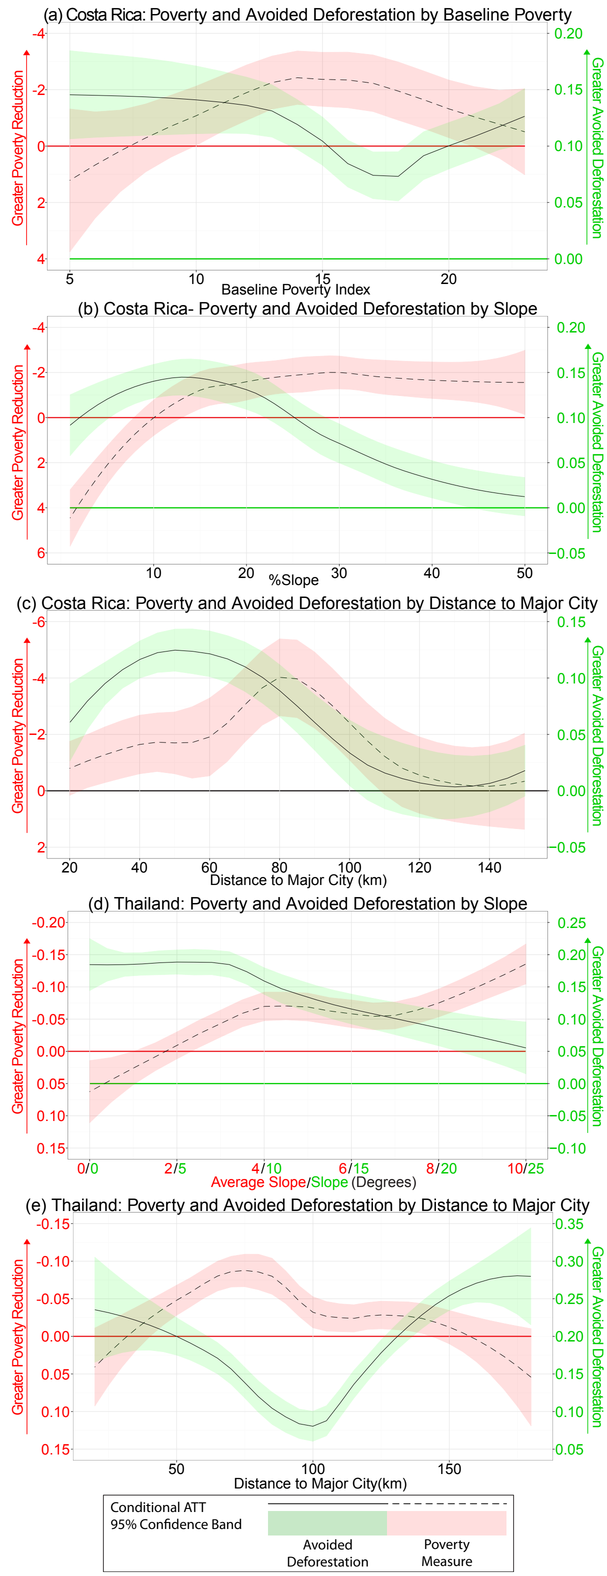

Figure 2: Heterogeneous responses to protection.

Results: Heterogeneous Impacts

Figure 2 presents the results. In each panel, the solid and dashed lines represent the estimated difference between protected and counterfactual units, i.e. the conditional average treatment effect on the treated (ATT) for avoided deforestation and poverty reduction, respectively. The green (red) shaded area around the solid (dashed) line represents the 95% pointwise confidence band for the avoided deforestation (poverty reduction) ATT estimate.[21] The solid green (red) horizontal line represents the zero line for the avoided deforestation (poverty reduction) estimate, at which there are no impacts from the establishment of protected areas.

Poverty Traps

We first test whether protected areas reinforced or exacerbated poverty traps in Costa Rica. If this were the case, we would expect to find that areas that were very poor at baseline would be negatively affected by protected areas. Based on theory, we would also expect that negative effects occur only when land-use restrictions are binding. Thus any exacerbation of poverty should be accompanied by avoided deforestation.

The results in Figure 2(a) confirm that avoided deforestation (solid line) in Costa Rica is positive across observed baseline poverty values. In other words, protected areas did impose binding land-use restrictions. Avoided deforestation is relatively constant along a majority of the baseline poverty range, although there is a dip between baseline poverty index values of 15 and 18. Poverty reduction (dashed line), however, appears to be U-shaped (inverted) as a function of baseline poverty. The estimates suggest that protected areas achieved significant poverty reduction for most of the range above the median baseline poverty level (poverty index = 12). At very high levels of poverty, these effects are not significantly different from zero. The

[21] See Figures 14 - 16 for more detailed illustration of all the impact heterogeneity results.

LOESS results therefore do not suggest that protected areas exacerbated poverty in the poorest populations. In fact, a majority of the poorest areas experienced poverty reduction compared to their estimated counterfactual poverty levels.[22]

Moderating Covariates

To better understand the nature of protected areas' impacts on poverty, we next consider two covariates that are highly related to poverty and, based on theory, are expected to moderate the impacts of protection: slope and distance to major cities. The primary driver of deforestation in Costa Rica and Thailand was agriculture (de Camino Velozo et al. 2000, Evans 1999, Cropper et al. 1999).[23] Slope is highly correlated with agricultural potential: the steeper the slopes, the less suitable the land is for agriculture. Steeper slopes are therefore associated with lower deforestation pressure and, therefore, lower opportunity costs of protection.[24] Slope and baseline poverty are also highly correlated: in Costa Rica the mean slope for land among the poorest quartile is 16.4 percent, whereas for the richest quartile it is only 3.8 percent.

Like slope, distance to city is also positively correlated with baseline poverty: in Costa Rica the mean distance of the poorest quartile is 70 km and of the richest quartile is 9 km. However, the distance to a major market city has a more complicated theoretical relationship with deforestation and protection. On one hand, being far from cities lowers agricultural returns and thus the returns to deforestation (because of, for example, higher transportation costs and poorer price

[22]As a robustness check we run a parametric quantile regression. These results are consistent with the LOESS results (see Appendix B).

[23]Logging was also an important source of deforestation during this time period and large-scale logging often cleared the way for conversion of previously forested land to agricultural use. Forest cover in logged areas tends to regenerate in these nations unless used for agriculture.

[24]Slope captures other deforestation pressures too, such as ease of logging (Pfaff et al. 2009), but agriculture is the key deforestation force in our study. In Costa Rica slope has been shown to be a good proxy for agricultural suitability (Ferraro and Hanauer 2011). Furthermore, the response functions conditional on slope and baseline labor force in agriculture exhibit similar trends (see Figure 15).

information). On the other hand, being far from cities also means one is likely to be far from the nodes of enforcement of land-use regulations inside and outside protected areas, thus increasing returns to agriculture. Finally, if one believes that tourism and associated infrastructure development is a key mechanism through which protection reduces poverty, then greater distance from cities implies less potential for poverty reduction. Thus the opportunity costs from protection can change nonlinearly as distance to cities increases.

Panels (b) and (c) of Figure 2 present the results of the analysis of the two moderating covariates in Costa Rica. Protection on low-sloped land is associated with significant tradeoffs in joint outcomes. We observe statistically significant poverty exacerbation up to an average slope of 10%, whereas the associated impact on avoided deforestation is relatively high along this range. Between approximately 15% and 40% slope, we observe 'win-win' outcomes of avoided deforestation and poverty reduction statistically different from zero. The results help to explain why we do not observe an association between protected areas and poverty traps despite evidence that land-use restrictions were binding. Figure 2(b) also indicates that the protection of low-sloped land is associated with significantly more avoided deforestation than the protection of steeply-sloped land.[25] As noted by Andam et al. (25), protected lands are rarely located on lands highly suitable for agriculture, and thus we can see why Andam et al. (2010, 2008) find a 'win-win' outcome, on average. These results suggest that protected areas are not serving as poverty traps partly because they tend to be sited in areas with low agricultural potential and thus low opportunity costs.

Figure 2(c) confirms the conjecture that distance to major cities captures countervailing forces and thus may generate nonlinear relationships between protection and the outcomes. The interval at which poverty reduction is greatest is

[25]This relationship arises largely because the amount of deforestation in the absence of protection decreases with slope (see Figure 15).

farther from cities than the interval at which avoided deforestation is greatest. Nevertheless, there is a substantial overlap of poverty reduction and avoided deforestation ('win-win') at intermediate distances (approximately 40km to 100km). These results provide indirect evidence that protected areas are not creating poverty traps partly because they tend to be sited in localities that can respond to opportunities afforded by tourism and associated infrastructure development. They also suggest that poor localities far from cities may not respond as well to protection as poor localities closer to cities.

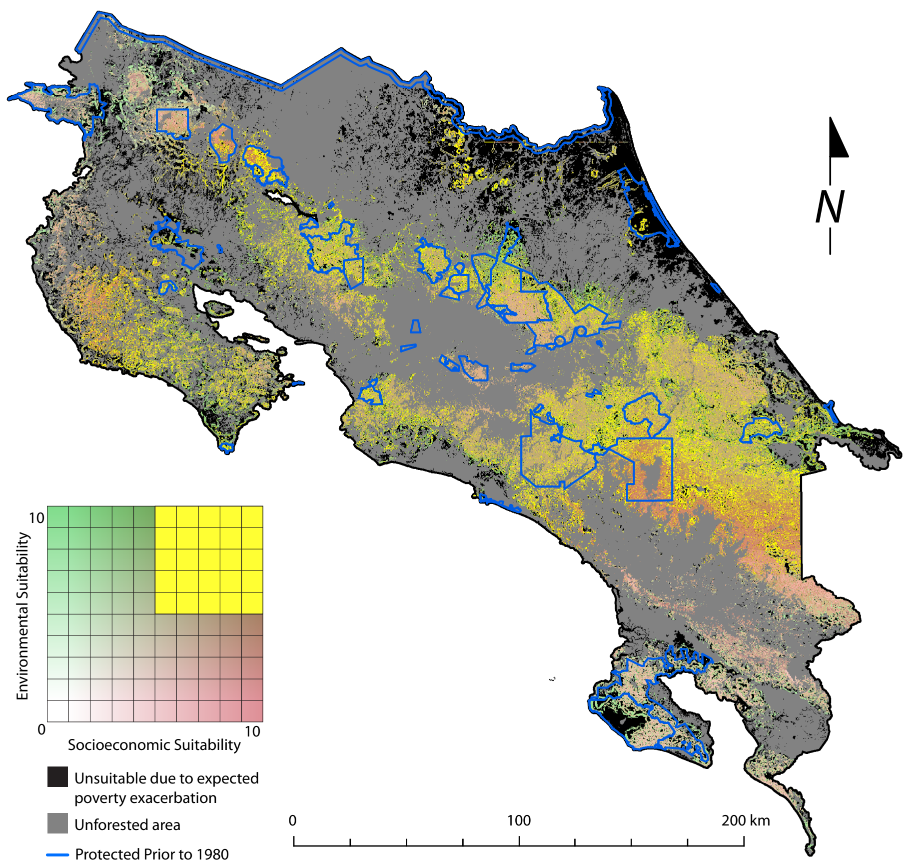

Figure 3: Costa Rica protected area suitability map.

In Thailand, we lack baseline poverty data,[26] but we can examine protection's impact on deforestation and poverty as a function of slope and distance to major cities. The shapes of the PLM graphs in Figure 2 panel (d) look remarkably similar to the shapes of the corresponding graphs for Costa Rica: slope is negatively related to avoided deforestation and positively related to poverty reduction. While there is a range over which 'win-win' outcomes are observed, the general trend of tradeoffs (more poverty reduction correlating with less avoided deforestation) is even more pronounced in Thailand. As in Figure 2(c), we observe in Figure 2(e) a nonlinear relationship between avoided deforestation and poverty impacts as a function of distance from major cities. The relationship with avoided deforestation in Thailand looks different from the relationship observed in Costa Rica (lower avoided deforestation at intermediate distances), but the relationship between poverty impact and distance from cities looks strikingly similar in both nations: the largest reductions in poverty are observed at intermediate distances from major cities.

Results: Suitability Mapping

Figure 2 suggests that the way in which areas respond to protected areas established in their midst will differ conditional on observable baseline characteristics. An understanding of these heterogeneous effects offers insights into how protected areas can be established in the future to manage tradeoffs between environmental and poverty reduction goals.

Suitability mapping allows one to visualize the joint outcomes spatially. We use the results from the previous section to create illustrative protected areas suitability maps for Costa Rica and Thailand. We break the regions into 3 hectare units and, based on results from PLM models, assign each unit a suitability score according to

[26] As did Andam et al. (2010), we address this lack of baseline poverty data by matching on a large number of baseline and time-invariant variables likely correlated with baseline poverty and by including district fixed effects.

the predicted impact on deforestation or poverty if the unit were protected (see Material and Methods and Appendix B). For example, based on historical impacts of protected areas in Costa Rica, a land parcel located on slopes of approximately 12% is highly suitable for protection in terms avoided deforestation, but only moderately suitable in terms of poverty reduction (recall that we are controlling for other parcel characteristics in the PLM estimation). By mapping underlying covariate relationships jointly with deforestation and poverty outcomes, we are able to identify areas of 'win-win', 'lose-lose' and 'win-lose'. These maps therefore are a type of graphical illustration of conditional empirical success rules.

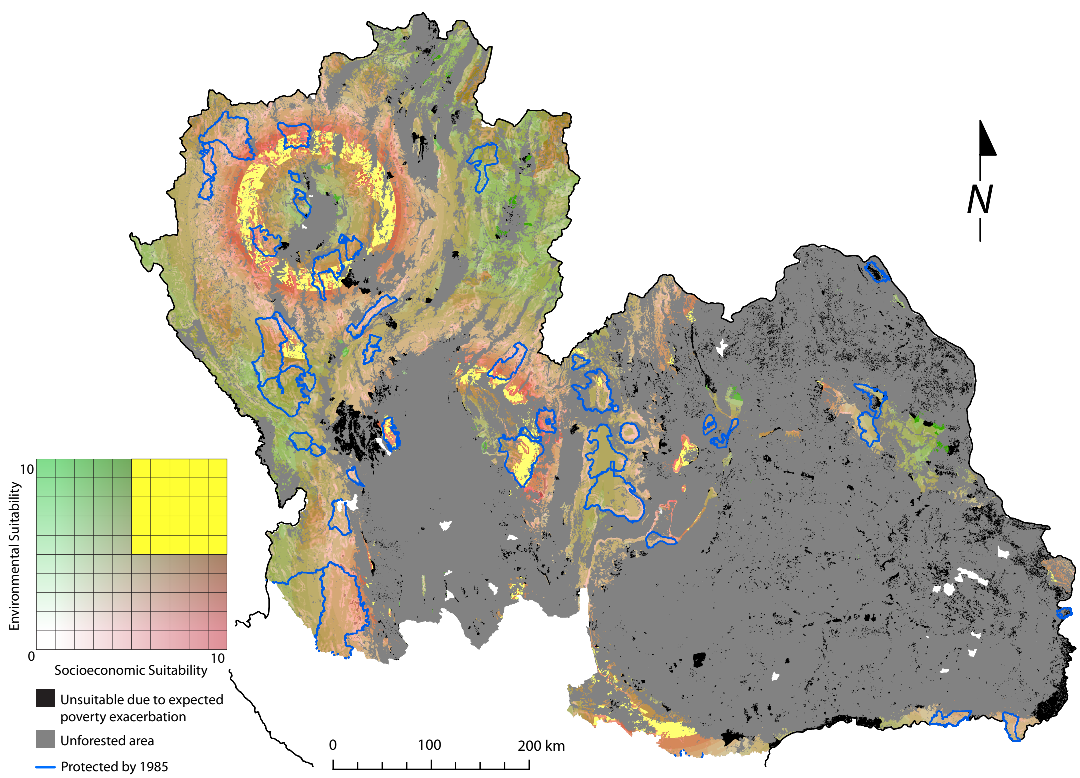

Figure 4: Thailand protected area suitability map.

We classify a land parcel's suitability for protection based on its slope and its distance from major cities, which are two time-invariant characteristics that are

typically available to decision makers (see Methods and Appendix B for details).[27] Because protection is assigned mainly to forested areas in the two nations, we limit our classification to parcels that were forested in the final period of our analyses: 1997 for Costa Rica and 2000 for Thailand.

Figures 3 and 4 display the illustrative suitability maps. The bivariate color grid represents increasing suitability for protected areas in terms of avoided deforestation (horizontal axis) and poverty reduction (vertical axis), based on historical impacts. Boundaries of the protected areas used to estimate the historical impacts of protection are in blue. In yellow, we highlight areas in the upper five deciles for both potential avoided deforestation and poverty reduction. These locations might be considered potential 'win-win' locations (see Methods and Appendix B). In Costa Rica, 324,156 hectares of forest in 1997 are classified as 'win-win' locations (14% of the total) with an average environmental (socioeconomic) suitability score of 7.25 (6.77). In Thailand, 662,013 hectares of forest in 2000 are classified as 'win-win' (5% of the total) with an average environmental (socioeconomic) suitability score of 6.17 (6.38). In black, we highlight areas that, based on historical responses, would likely experience poverty exacerbation and thus might be considered undesirable for establishing a protected area, regardless of environmental suitability. Because these expected poverty outcomes are driven by low-sloped land, all these areas are associated with positive avoided deforestation and therefore expected 'win-lose' outcomes. In Costa Rica, 659,730 hectares are classified as likely exacerbation locations (28% of the total forest area) with an average environmental suitability score of 5.2. In Thailand, 1,180,041 hectares are classified as likely exacerbation locations (10% of the total forest area) with an average environmental suitability score of 6.6 (see Appendix B).

[27] Slope data are often used in global protected area analyses (Joppa and Pfaff 2009, Nelson and Chomitz 2009).

These maps are meant to be illustrative and used in conjunction with other sources of data and expertise. Other baseline conditions are likely to be important in determining tradeoffs. In future applications, suitability maps would incorporate knowledge of other indicators of biological value (e.g., endemic species) and other forms of expert knowledge about local conditions into a more sophisticated optimization algorithm (see Naidoo et al. (2006) for examples of algorithms). Moreover, the maps are based on the assumption that past associations will hold for future outcomes, which may not be true in rapidly changing societies. Suitability maps present a static picture of expected relationships and do not capture potential general equilibrium effects: the protection of an area may fundamentally change the suitability of the remaining unprotected areas. Finally, future analyses should also incorporate an understanding of the differential impacts of protected area types (e.g., wildlife refuges versus national parks) and other characteristics determining economic opportunities.

Discussion

Debates over the effectiveness of protected areas in achieving conservation results and in affecting poverty are often based on little empirical evidence. Critics of protected areas highlight the role that protected areas can play in limiting agricultural development and exploitation of natural resources. They would thus predict that observable characteristics associated with high levels of avoided deforestation from protection would also be associated with poverty exacerbation. Proponents highlight the role that protected areas can play in supplying ecosystem services, promoting tourism and improving infrastructure. They would thus predict that characteristics associated high levels of avoided deforestation from protection would be associated with high levels of poverty reduction. Our results indicate that

the realities in Costa Rica and Thailand are more complicated than either of these two stereotypes.

Our results are not consistent with protected areas creating poverty traps. In fact, the results suggest that protection in areas associated with high poverty has, on average, reduced poverty while also reducing deforestation. Such 'win-win' outcomes were most commonly associated with locations at intermediate distances from major cities (40-80 km) and on land of moderate to poor agricultural potential. These patterns are consistent with a hypothesis that protected areas have reduced poverty by being placed on lands with little agricultural value that, by their proximity to major markets, can benefit from tourism and associated infrastructure development (thus offsetting any losses from foregone agriculture and forest resource exploitation). To support this hypothesis, more explicit analyses of mechanisms will be necessary (e.g., Imai et al. (2010)). Although we find no evidence that protection, on average, created poverty traps, our results do not imply that protection reduced poverty in all poor communities. poverty may have been exacerbated in some poor communities.

Despite the lack of evidence for poverty traps from protected areas, the results do suggest potential tradeoffs: the most avoided deforestation is found on low-sloped land with high agricultural value, but these lands are where poverty exacerbation is observed. Thus although protected areas did lead, on average, to moderate levels of avoided deforestation and poverty reduction in Costa Rica and Thailand, our analysis points to tradeoffs if decision makers desire higher levels of either outcome. The potential for tradeoffs underscores the importance of conditional empirical success rules, especially as practitioners attempt to better target protected area investments to increase conservation effectiveness and as policymakers look to protected areas as a means to obtain international financial transfers from reducing emissions from deforestation and forest degradation (REDD) programs.

Costa Rica and Thailand are middle-income countries, have made substantial investments in their protected area systems, and have relatively successful eco-tourism sectors. Whether our results would hold for other nations is an open question. Our approach can, and should be, replicated in other nations through cooperation between groups collecting spatially explicit data on poverty, protected areas, and land-use change. A greater understanding of heterogeneous impacts can improve conservation planning and offer insights into the potential tradeoffs between environmental and development goals in future efforts to reduce emissions from deforestation and degradation.

Chapter III

Causal Mechansisms of Protected Areas

Introduction

The proliferation of protected areas in recent decades has led to increased interest in understanding their economic impacts on surrounding populations. However, there have been few studies with the requisite data and methodologies to accurately estimate the socioeconomic impacts of protected areas (Andam et al. 2010, Coad et al. 2008). The few studies that satisfy the conditions for an impact study of high quality,[28] have found that the establishment of protected areas has been associated with poverty reductions in surrounding areas (Canavire-Bacarezza and Hanauer 2011, Andam et al. 2010, Sims 2010). Such results run counter to the conventional wisdom (Coad et al. 2008, Wilkie et al. 2006, Adams et al. 2004) and limited theory (Robalino 2007). Unfortunately the quasi-experimental methods employed in the previous protected area impact evaluations are not suitable for addressing the underlying mechanisms through which protected areas affect poverty. Therefore, the question of why protected areas have been found to be associated with reductions in poverty remains. An understanding of these mechanisms would help explain why impacts occur, rather than simply quantifying the impacts.

The establishment of protected areas has elicited concern from poverty advocates due to their associated land-use restrictions (Wilkie et al. 2006, Adams

[28]See Ferraro (2008) for a discussion of the necessary components.

et al. 2004). Coupled with the facts that, historically, protected areas have been placed on marginal lands (Joppa and Pfaff 2009),[29] and much of the remaining global biodiversity (land likely to be targeted for protection) lies in areas of high poverty (Sachs et al. 2009), land-use restrictions are expected to impose economic hardship on already imperiled populations. This concern is formalized in a Von Thunen model developed by Robalino (2007) in which the author shows that land-use restrictions associated with protected areas are predicted to negatively impact landless workers. Therefore, from a policy standpoint, identification of the mechanisms through which protected areas affect poverty is of particular interest; especially if negative channels can be mitigated, or positive channels bolstered, through social policy.

We use rich biophysical and socioeconomic data from Costa Rica, a developing country with a renowned protected area network, to identify and quantify the causal mechanisms through which protected areas established prior to 1980 impacted poverty between 1973 and 2000. Using recently developed quasi-experimental approaches to mechanism analysis which allow for both causal interpretation of mechanism effects and salient comparison to previous studies of Costa Rica (Andam et al. 2010, Ferraro and Hanauer 2011, Ferraro et al. 2011), we quantify the proportion of estimated poverty alleviation (Andam et al. 2010) from tourism, infrastructure development and ecosystem services due to the establishment of protected areas. By proxying for the respective mechanisms with park entrances, changes in road networks and changes in forest cover, we find that nearly half of the poverty alleviation associated with the establishment of protected areas is causally attributable to tourism. Conversely, infrastructure development accounts for a relatively small proportion of the estimated poverty alleviation. Finally, because we proxy for ecosystem services with avoided deforestation (which is associated with

[29] This is a concern because Andam et al. (2010) show that marginal land is correlated with poverty.

potentially negative poverty mechanisms), we argue that our findings of no mechanism affect due to the prevention of deforestation implies a positive impact on poverty due to the preservation of ecosystems associated with protected areas. In addition, we conduct several robustness checks which provide evidence that our general findings are likely not an artifact of our econometric strategy.

Background

Recent Studies

Only a handful of studies from developing nations have met the necessary data and methodological requirements for a protected area impact study of high quality. To properly account for changes in poverty due to the establishment of protected areas, a study must incorporate pre-protection, baseline measures of poverty. In addition, the non-random nature in which protected areas are established must be accounted for in the empirical strategy, which necessitates rich baseline measures of covariates that jointly determine the establishment of protected areas and poverty outcomes.

Recent studies, that meet the aforementioned requirements, from Bolivia (Canavire-Bacarezza and Hanauer 2011), Thailand (Sims 2010, Andam et al. 2010), and Costa Rica (Andam et al. 2010) have found the establishment of protected areas to be associated with subsequent reductions in poverty. Our study follows directly from Andam et al. (2010) in which the authors designate census tracts (segmentos) with 10% or more of their areas protected, as treated. They then use matching techniques to construct a counterfactual group that is similar along pretreatment dimensions to the treated census tracts. The authors' calculation of average treatment effect on the treated (ATT) provides evidence that census tracts with protected areas that were established prior to 1980 had differentially greater levels of poverty reduction between 1973 and 2000 than comparable unprotected census tracts.

The only paper to empirically address the potential mechanisms through which protected areas affect economic outcomes aims to quantify the effects of eco-tourism on local wages. Robalino and Villalobos-Fiatt (2010) explore how national parks affect local wages in Costa Rica and how these effects vary within different areas of a park and among different social groups. They use highly disaggregated geographic references, and find that parks' effects on wages vary according to economic activity and proximity to the entrance of the park. Workers close to entrances receive higher wages and are employed in higher-paid, non-agricultural activities.

The Argument for the Estimation of Mechanism Effects

The dearth of information regarding the mechanisms through which protected areas affect economic outcomes serves an archetype of the criticism leveled at experimental and quasi-experimental reduced form estimation (Deaton 2009, Heckman 2010).[30] Studies in which the estimate of interest is the effect of some non-random treatment have increasingly turned toward quasi-experimental methods. Such methods (e.g., matching) seek to mimic the identification mechanism of a randomized experiment through specific *ex post* manipulation of the data (e.g., trimming, weighting, etc.). Although quasi-experimental identification strategies tend to be transparent– if one can control all other influences on the outcome then all remaining differences are due to treatment assignment –they generally lack the ability to identify the causal mechanisms through which treatments work. There are three main limitations to the estimation of mechanism effects: (1) rich data that include intermediate observation of mechanism values are necessary; (2) post-treatment mechanisms are, by definition, affected by treatment (and, therefore, generally subject to selection bias) so simply controlling for post-treatment mechanisms within a regression framework will generally lead to biased estimates (Rosenbaum 1984); and (3) within

[30]Given the methodology used in this study I focus on quasi-experimental methods.

the quasi-experimental framework, the absence of theoretical models limit the identification of mechanistic channels (e.g., Heckman (2010)).

Despite these hurdles, an understanding of why protected areas affect poverty is of paramount importance. Given the recent movements toward increasing the global coverage of protected areas and the goal that the establishment of protected areas should at least do no economic harm (Adams et al. 2004), social planners need a deeper understanding of the interplay between protected areas and economic outcomes. This is a point that is highlighted by Ferraro and Hanauer (2011) and Ferraro et al. (2011). These studies provide evidence that protected areas (in Costa Rica and Thailand) have had heterogeneous economic (and conservation) impacts according to demographic and biophysical characteristics. They argue (a la Manski (2005)) that understanding how different subgroups respond to treatment can help planners optimize the placement of future protected areas. Similarly, understanding the post-treatment mechanisms through which protected areas affect economic outcomes will allow planners to optimize social policy concurrent with the establishment of protected areas. In conjunction, an understanding of both the heterogeneous impacts and mechanisms of protected areas might greatly improve the economic outcomes associated with the future establishment of protected areas.

Data

We use data from Andam et al. (2010) to identify and quantify the mechanisms through which protected areas affected poverty in Costa Rica. The unit of observation is the census tract. The 1973 census is used as the baseline year 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 by protected land.[31] Conversely, any census tract that contains less that 1% protected land is considered unprotected and a binary treatment indicator is assigned accordingly.[32] A poverty index is derived for each tract from census data following Cavatassi et al. (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 3).

Mechanisms

Mechanisms have received the most attention in the epidemiology (surrogate variables) and psychology (mediating variables) literatures (e.g., Imai et al. (2010), Rubin (2004), Frangakis and Rubin (2002), Baron and Kenny (1986)). Whereas, economics has seen relatively little estimation of mechanism effects (Flores and Flores-Lagunes 2011). Fundamentally a causal mechanism can be viewed as a variable which, once affected by treatment, impacts the outcome of interest. In causal Directed Acyclic Graphs (DAGs) developed by Pearl (2009) and highlighted

[31] 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%).

[32] 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.

Table 3: Summary Statistics of Matching Covariates and Mechanism Variables

Covariate	Description	Mean	Median	Std. Dev.	Range
		Matching 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 or 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 300m2 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 1 ha parcel within the census tract.	308,000	66,400	699,100	0.28-7,590,000
		Mechanism Variables			
Park Entrance	Binary indicator equal to 1 if census tract has at least 10% of its area occupied by a protected area with a park entrance	0.0276	0	0.164	0-1
Δ Roadless Volume	Change in roadless volume between 1969 and 1991	-8.76e+04	-1.75e+01	605191	-2.65e+07-6.32e+04
Δ Forest Cover	Percent change in forest cover between 1960 and 1986 within each census tract	0.0084	0	0.0926	0-0.75

by Morgan and Winship (2007), a mechanism (S) is drawn as causal pathway ($\rightarrow$) that links treatment (T) to outcome (Y), $T \rightarrow S \rightarrow Y$. Therefore, a causal mechanism is a variable whose quantity is directly effected by treatment, the result of which causes a direct change in the outcome of interest.

Mechanism variables

The putative mechanism through which protected areas achieve environmental outcomes (e.g., preventing deforestation, etc.) is land-use restriction. Such restrictions, which limit access, conversion and the exploitation of natural resources, would be expected to negatively impact economic conditions in surrounding areas. If land-use restrictions were the only mechanism (or the dominant mechanism) through which protected areas impact surrounding populations then we would expect poverty to have been exacerbated in these areas. In Costa Rica this has not been the case (Andam et al. 2010). We are, therefore, interested in investigating potential mechanisms through which protected areas have *positively* influenced economic conditions in surrounding populations.

Tourism. Tourism is widely cited (anecdotally, e.g., Wilkie et al. (2006), Adams et al. (2004) and empirically, e.g., Menkhaus and Lober (1996)) as a likely mechanism through which protected areas enhance local economies. Costa Rica's stable government and rich biodiversity make it a popular destination for so-called eco-tourists. Conjecture that tourism, catalyzed by the establishment of protected areas, is likely enhancing the welfare of surrounding communities is supported by the fact that approximately 54% of international tourists visit a protected area (ICT 2010). Further empirical support is offered from Ferraro et al. (2011). The authors find that reductions in poverty due to the establishment of protected areas are greatest at intermediate distances to cities; this range coincides with the location of a majority of Costa Rica's national parks (which receive the most tourists). Using

global positioning system (GPS) data from Robalino and Villalobos-Fiatt (2010) we proxy for tourism with the establishment of a park entrance. Of Costa Rica's 39 protected areas that were established prior to 1980, 19 received at least one park entrance prior to 2000 (total of 23 entrances). A protected census tract (see definition above) is considered affected by a park entrance if it is occupied by a protected area in which at least on entrance was established. According to this assignment rule, 122 census tracts are considered affected by a park entrance.

Infrastructure Development. Access to infrastructure can be expected to enhance economic outcomes (e.g., reduced production costs). Previous studies from Costa Rica and Thailand have shown a relationship between access to urban infrastructure and poverty (Andam et al. 2010). We proxy for infrastructure with road networks. Access to roads increases access to markets and other resources (reducing transportation costs, etc.). In addition, roads serve as a good indicator of the level of infrastructure development and urbanization. We are, therefore, interested in how differential levels of road development, due to the establishment of protected areas, has impacted poverty in surrounding communities. We use changes in roadless volume (Watts et al. 2007) between 1969 and 1991 to capture the impact of changes in access to infrastructure. Roadless volume is an aggregation of the euclidean distance to a road for each one-hectare land parcel within a census tract, adjusted for the size of the land parcel. Roadless volume is calculated by summing the product of the area of each land parcel (1 ha in this case) and the distance of that parcel to the nearest road (1969 and 1991). Therefore, higher measurements of roadless volume indicate fewer road networks within a municipality. Summary statistics for baseline roadless volume and changes in roadless volume can be found in Table 3.

Ecosystem Services. Since the seminal paper by Costanza et al. (1997) there has been great interest in quantifying the economic impacts of the services provided

by intact ecosystems. One of the arguments from protected area advocates for potential win-win outcomes is that the establishment of protected areas prevent ecosystem degradation (win) thereby providing a stream of economic benefits (win) to surrounding communities (in addition to the global benefits such as carbon sequestration). We proxy for maintenance of ecosystem services via avoided deforestation (the difference between observed and counterfactual levels of deforestation in protected census tracts). We are interested in how the causal reduction in deforestation (that would have occurred in the absence of protection) due to the establishment of protected areas has impacted surrounding communities. We measure the percentage of deforestation within each census tract using GIS and forest cover boundaries from 1960 and 1986 (see Table 3 for baseline and mechanism measurements of forest cover).[33]

Methods

Mechanism Concepts

In the context of a quasi-experimental design, in which one conditions on baseline characteristics to estimate effects of treatment on subsequent-stage outcomes, estimating the effect of a mechanism is confounded by the fact that the mechanism is necessarily observed post-treatment. As such, mechanisms are generally affected by treatment assignment (or selection) and, therefore, confounded. Thus, controlling for such concomitant variables generally leads to biased estimates (Rosenbaum 1984).[34]. This of course rules out the argument for the inclusion of a mechanism as a control. Therefore, precluding the estimate of mechanism effects via the difference between the estimates of a specification (e.g., regression or matching)

[33]We acknowledge that avoided deforestation is a coarse measure of maintained ecosystem services. In fact, preventing deforestation likely produces countervailing (to ecosystem services) mechanism effects. See the Summary of Results Section for a detailed discussion.

[34]The exception is when the concomitant variable is not effected by the treatment, in which case it can be considered a baseline covariate (Rosenbaum 1984)

with and without the mechanism variable. In order to estimate the effects of a mechanism, it must be treated, and controlled for, like an outcome (hence the concept of surrogate variables in the epidemiology literature, see e.g., Mealli and Rubin (2003)). We appeal to the concept of principal strata (see below) developed by Frangakis and Rubin (2002) to conceptualize and account for mechanisms within the potential outcomes framework.

Setup

To estimate the causal mechanisms through which protected areas have impacted economic outcomes we use an augmented potential outcomes framework (we follow the framework and much of the notation of Flores and Flores-Lagunes (2011). In the traditional potential outcome framework there are two potential outcomes, $Y_i(1)$ and $Y_i(0)$, for each individual $i \in N$ under treatment ($T = 1$) and control ($T = 0$), respectively. Intuitively this means that each individual would have one outcome if they were to receive treatment, and another if they were withheld treatment. Unfortunately, for any given individual i only one of the two potential outcomes is observed: $Y_i^{obs}(1)|T = 1$ or $Y_i^{obs}(0)|T = 0$. In practice either individual i's outcome under treatment is observed given they were treated or individual i's outcome in the absence of treatment is observed given they were in the control group. This is the fundamental problem for the estimation of causal effects because individual treatment effects are calculated $\tau_i = Y_i(1) - Y_i(0)$, for which only one of the rhs terms is observed. In order to calculate treatment effects in the absence of random assignment, it is necessary to invoke the conditional independence assumption[35]

Assumption 1 $Y_i(1), Y_i(0) \perp\!\!\!\perp T_i | X_i$,

which states that potential outcomes are independent ($\perp\!\!\!\perp$) of treatment given a set of covariates X that jointly determine outcomes and selection into treatment.

[35]Also known as ignorability, unconfoundedness or selection on observables.

Random assignment assures independence, without condition, due to the fact that each individual has and equal probability (or more generally, a probability known to the experimenter) of assignment to treatment. Therefore, all covariates (X) that influence outcomes are balanced across treatment and control groups, hence the independence of treatment and outcome. For conditional independence to hold under non-random assignment, one must condition on (e.g., matching) or control for (e.g., regression) all covariates (X), thus rendering any remaining differences in outcomes between groups a function of treatment.

Principal strata

Further complications arise when post-treatment mechanisms are introduced. Suppose S is a post-treatment mechanism that is measured at an intermediate period between administration of treatment and measurement of outcome.[36] Because, by definition, S is affected by treatment it is not unconditionally independent of treatment[37] and thus must be handled in a manner similar to the outcome of interest (Y). Therefore, as with Y, S has two potential outcomes $S_i(1)$ and $S_i(0)$ for each i, depending on assignment to treatment or control, respectively. This simply states that because mechanisms are affected by treatment, with the exception of some special cases, the mechanism outcome for each individual is dependent on the administered treatment. The implications, within the potential outcomes framework, are that four potential outcomes must now be considered for each individual: $(Y_i(1), Y_i(0), S_i(1), S_i(0))$.

There are now four compound potential outcomes of interest for i: $Y_i(1, S_i(1))$, the outcome under treatment when the mechanism is affected by treatment;[38]

[36] Note that the three mechanisms of interest are denoted formally as S_j, where $j = 1, 2, 3$. For ease of exposition throughout a majority of this discussion, the subscript is omitted.

[37] This is true under random assignment of treatment as well.

[38] $Y_i(1, S_i(1))$ represents the total effect of treatment and is equivalent to $Y_i(1)$ in the traditional potential outcomes framework.

$Y_i(1, S_i(0))$, the outcome under treatment when the mechanism is not affected by treatment (mechanism is blocked (Flores and Flores-Lagunes 2011)); $Y_i(0, S_i(0))$, the outcome under control and the mechanism is not affected by treatment;[39] and $Y_i(0, S_i(1))$, the outcome under control when the mechanism changes as if the individual was treated.[40]

To help conceptualize the joint potential outcomes and identify the casual mechanism effect we use the principal strata framework developed by Frangakis and Rubin (2002) (see also, Rubin (2004), Mealli and Rubin (2003)). Defining a principal stratum is similar to the concept of matching individuals (or groups of individuals) based on similar potential outcomes in a standard quasi-experimental setting. Two units from different treatment arms share a principal stratum if they share potential mechanism outcomes (formally a principal stratum is defined where $\{S(0) = s_0, S(1) = s_1\}$, see below).

To identify units from disparate treatment arms but similar principal strata an extension to the conditional independence of Assumption 1 is necessary

Assumption 2 $S_i(1), S_i(0) \perp\!\!\!\perp T_i | X_i$.

We term Assumption 2 conditional mechanism isolation. Morgan and Winship (2007) note that in order to estimate the effect of a mechanism on outcomes, the mechanism must be isolated from (independent of) confounding covariates. Assumption 2 states that potential mechanism outcomes are independent of treatment given a set of covariates (X) that jointly determine selection into treatment and mechanism outcomes, and, therefore, isolated from confounders. Under Assumption 2 we can identify units within similar principal strata: units from disparate treatment arms with similar values of X lie within common strata

[39] $Y_i(0, S_i(0))$ represents the outcome in the absence of treatment and is equivalent to $Y_i(0)$ in the traditional potential outcomes framework (implied that post-treatment mechanism is not affected in absence of treatment).

[40] Note that, in general, only $Y_i(1, S_i(1))$ and $Y_i(0, S_i(0))$ are observed in practice, leaving $Y_i(1, S_i(0))$ and $Y_i(0, S_i(1))$ as counterfactuals that necessitate estimation.

and, therefore, share similar potential mechanism outcomes. Assumptions 1 and 2 imply that potential outcomes and potential mechanism values are independent of treatment given covariates X. Combining Assumptions 1 and 2 we have

Assumption 3 $Y_i(1, S_i(1)), Y_i(1, S_i(0)), Y_i(0, S_i(0)), Y_i(0, S_i(1)), S_i(1), S_i(0) \perp\!\!\!\perp T_i | X_i$.

The necessary condition for Assumption 3 to hold, is that the covariates (X) must jointly determine selection into treatment, outcomes of interest *and* mechanism outcomes. Upon cursory examination Assumption 3 may seem somewhat untenable. However, when one considers jointly that a primary purpose of X is to control for the non-random process of selection into treatment and that treatment directly affects mechanisms, Assumption 3 seems more reasonable. Formally Assumption 3 allows

$$E\left[S_i(1)|X_i = x, T = 1\right] = E\left[S_i(1)|X_i = x, T = 0\right] \tag{7}$$

$$E\left[S_i(0)|X_i = x, T = 1\right] = E\left[S_i(0)|X_i = x, T = 0\right]. \tag{8}$$

Equations (7) and (8) state that the expected mechanism outcomes under treatment, for individuals that were treated, are equal to those in the control group, with similar values of X, *had they been treated*, and vice versa.[41] We present (7) and (2) for completeness, however, note that only (8) is necessary for our analyses.

Estimands

In the study from which we draw (Andam et al. 2010) the estimand of interest is the average treatment effect on the treated, ATT. Estimation of the ATT is akin to asking the question, "what would outcomes for treated units have been had they

[41]These are analogous to $E[Y(1)|X, T = 1] = E[(Y(1)|X, T = 0]$ and $E[Y(0)|X, T = 1] = E[(Y(0)|X, T = 0]$ which follow from Assumption 1. These equations are commonly used in the matching literature and demonstrate the equality (in expectation) of potential outcomes conditional on observable covariates (X) used to estimate average treatment effects.

not been treated?" Estimating the ATT is appropriate in a protected area impact evaluation because the supposition that all units have the potential to receive treatment is inappropriate (e.g., supposing that urban census tracts might contain a protected area). Therefore, the average treatment effect (ATE), which requires the estimation of counterfactual outcomes for all control units (had they been treated) is not the estimand of interest.

Given that the total treatment effect estimand of interest is the ATT, the mechanism treatment effect of interest is the Mechanism Average Treatment Effect on the Treated. Our estimands follow directly from the framework for mechanism average treatment effects (MATE) and net average treatment effects (NATE) developed by Flores and Flores-Lagunes (2011). Defining a principal strata as $\{S(0) = s_0, S(1) = s_1\}$,[42] the MATT can be written

$$MATT = E\left\{E\left[Y_i(1, S_i(1)) - Y(1, S_i(0)) | S_i(0) = s_0, S_i(1) = s_1, X_i = x, T = 1\right]\right\}. \tag{9}$$

To estimate the MATT one must ask, "what would outcomes for the treated have been, had they remained treated but treatment not affected the mechanism?" Estimation of the MATT answers this question by isolating the only source of variation in (9) to be the effect on outcomes due to a change in the mechanism (via blocking the effect of the mechanism on the outcome in the second term of (9)). A similar estimand of interest is the net average treatment effect on the treated (NATT) which isolates the effect on outcomes due to a change in treatment

$$NATT = E\left\{E\left[Y_i(1, S_i(0)) - Y_i(0, S_i(0)) | S_i(0) = s_0, S_i(1) = s_1, X_i = x, T = 1\right]\right\}, \tag{10}$$

[42] This states that individuals located within a common principal strata would have similar mechanism outcomes s_0 had they been in the control group ($S(0)$), or s_1 had they been treated ($S(1)$), independent of actual treatment received.

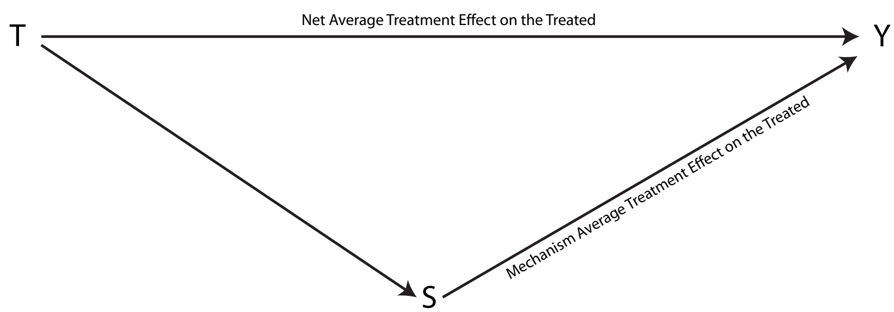

Figure 5: Causal Directed Acyclic Diagram (DAG) depicting the the concept of Mechanism Average Treatment Effect on the treated ($T \rightarrow S \rightarrow Y$) and Net Average Treatment Effect on the Treated ($T \rightarrow Y$) on outcome.

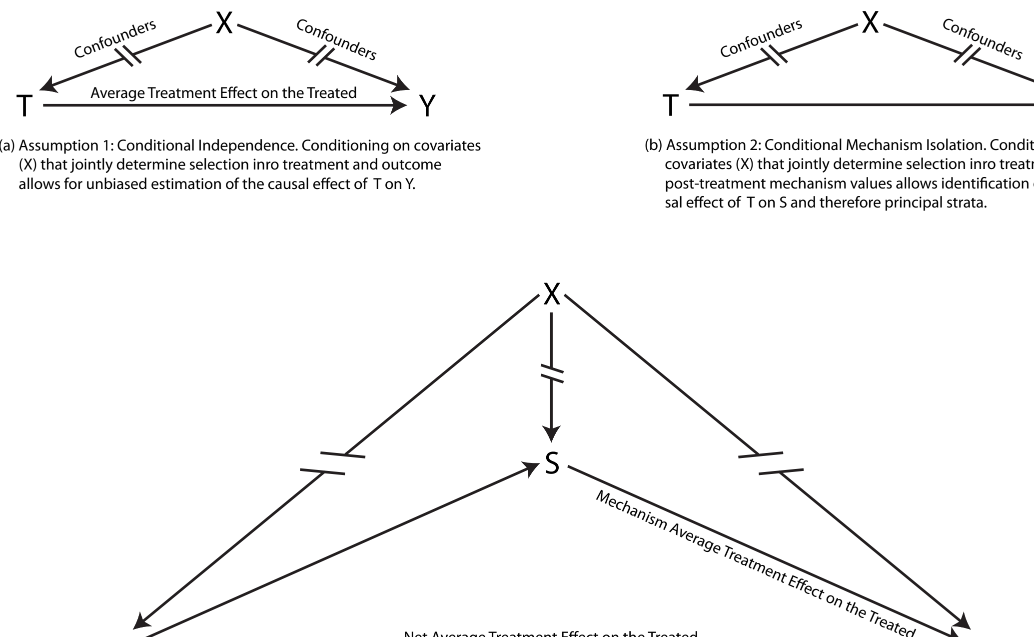

(a) Assumption 1: Conditional Independence. Conditioning on covariates (X) that jointly determine selection inro treatment and outcome allows for unbiased estimation of the causal effect of T on Y.

(b) Assumption 2: Conditional Mechanism Isolation. Conditioning on covariates (X) that jointly determine selection inro treatment and post-treatment mechanism values allows identification of the causal effect of T on S and therefore principal strata.

(c) Assumption 3: Sequential Ignorability. Conditioning on covariates that jointly determine selection into treatment, outcomes of interest, *and* mechanism outcomes allows for estimation of Mechanism Average Treatment Effects on the Treated (MATT) and Net Average Treatment Effects on the Treated (NATT).

Figure 6: Directed Acyclic Graphs (DAGs) demonstrating the assumptions necessary for the causal estimation of ATT, MATT and NATT. Each DAG shows how conditioning on observable covariates (X) breaks the confounding causal relationships ($T \leftarrow X \rightarrow Y$, $T \leftarrow X \rightarrow S$ and $S \leftarrow X \rightarrow Y$; represented by the broken single-headed arrows) and allows for estimation of ATT (a), the causal effect of treatment on mechanism outcomes (b) and MATT (and NATT)(c).

holding S at untreated levels. Estimation of NATT is akin to asking, “what would outcomes for the treated have been, had they not been treated but their mechanism values remained at levels realized under treatment?” An advantage of MATT and NATT (under Assumption 1) is that they decompose the ATT such that ATT=MATT+NATT.[43,44] This decomposition states that the average treatment effect on the treated is equal to the proportion of the of treatment effect that is due to a change in the mechanism (catalyzed by treatment), the MATT, and the proportion that is due solely to the effect of treatment (net of the effect of the mechanism), the NATT (see Figure 5). Therefore, once either MATT or NATT is estimated the complementary estimate falls out of the difference with ATT.

Estimation Strategy

Estimation of either MATT or NATT is confounded by the fact that $Y_i(1, S_i(0))$ is rarely observed.[45] We use matching in the first stage of the estimation to satisfy Assumption 3 (which encompasses Assumptions 1 and 2), see Figure 6. Post-matching we follow methods suggested by Flores and Flores-Lagunes (2011), using mechanism data from the matched control units, and a simple assumption about the way in which mechanisms affect outcomes within principal strata, to impute outcomes for treated units had treatment not affect the mechanism variables: $\widehat{Y}_i(1, \widehat{S}_i(0))$, the counterfactual of interest.

[43] Assumption 1 is necessary for this identity to hold. Morgan and Winship (2007) outline conditions under which $T \rightarrow Y$ can be estimated using a set of mechanisms (e.g., the set of mechanisms is exhaustive and isolated). However, one can measure the partial effect of $T \rightarrow Y$ using a non-exhaustive set of mechanisms, $\mathcal{S}$ (i.e., $\mathcal{S} \rightarrow Y$), which leads to an estimate of MATT. In conjunction with Assumption 1, under which unbiased estimates of the ATT can be estimated, the remaining difference between MATT and ATT can be attributed to the mechanisms not included in S.

[44] The full decomposition can be written: ATT=$E\left[Y(1, S(1)) - Y(1, S(0))\right] + E\left[Y(1, S(0)) - Y(0, S(0))\right]$, given principal strata $\{S(0) = s_0, S(1) = s_1\}$.

[45] In the case where a subgroup of treated units for which treatment did not affect mechanism values can be identified, Flores and Flores-Lagunes (2011) develop an estimand for the local average treatment effect (LNATE) which requires less restrictive assumptions. See the LNATT Section for an application of this methodology to our data.

First stage: matching

We use one-to-one Mahalanobis covariate matching with replacement and post-match bias-adjustment (Abadie et al. 2004, Abadie and Imbens 2006) to match control units to treated units. This approach serves two purposes. First, it provides an estimate of ATT, for comparison to MATT and NATT, which offers comparability to previous studies from Costa Rica (Andam et al. 2010, Ferraro and Hanauer 2011, Ferraro et al. 2011). Second, it provides a set of matched controls that, by Assumption 3 are within the same principal strata as the treated units to which they are matched. The latter purpose implies that the mechanism outcomes of the matched controls can be assumed to be the value observed by their treated counterparts, *had treatment not affected the mechanisms*. See Table 3 for a description of the covariates used for matching.

Second stage: estimate the influence of mechanisms

Flores and Flores-Lagunes (2011) suggest using a form of regression adjustment to impute outcomes for treated units had treatment not affected mechanisms, $\widehat{Y}_i(1, \widehat{S}_i(0))$. The necessary assumption for this approach (in addition to Assumption 3) is that the mechanism has a similar effect on potential outcomes $Y_i(1, S_i(1))$ and $Y_i(1, S_i(0))$, i.e., their conditional expectation functions share the same functional form (Flores and Flores-Lagunes 2011).

Assumption 4 *Suppose*

$$E\left[Y_i(1, S_i(1)) | S_i(1), X_i = x, T = 1\right] = a_1 + b_1 S_i(1) + c_1 X, \tag{11}$$

then,

$$E\left[Y_i(1, S_i(0)) | S_i(1), X_i = x, T = 1\right] = a_1 + b_1 S_i(0) + c_1 X_i. \tag{12}$$

Assumption 4 implies that the marginal effect of a change in the mechanism outcome has the same effect on units for whom exposure to treatment affects the mechanism as it does on units for whom exposure to treatment does not affect the mechanism.

In (11) and (12) of Assumption 4, b_1 represents the effect on outcome due to a change in the value of the mechanism S. The counterfactual of interest $(\widehat{Y}(1, \widehat{S}(0)))$ can be estimated by evaluating (12)), which uses the coefficients from (11), setting $S_i(0) = E[S_i(0)|T = 1] = [\widehat{S}_i(0)|T = 1]$ which, according to (8), is equal to the observed control mechanism values within the common principal stratum of each treated unit.

Empirical estimation of the counterfactual of interest $\left(\widehat{Y}_i(1, \widehat{S}_i(0))\right)$ is conducted by first running a regression of observed outcomes on covariate and mechanism values for treated units (as in (11)). Using the coefficients from this regression (a_1, b_1, c_1), we impute $\widehat{Y}_i(1, \widehat{S}_i(0))$ using the same treated unit covariates (as in (12)) *and the matched control unit mechanism outcomes* (where in (12) $S_i(0) = E[S_i(0)|T = 1] = S_i^{obs}(0)$ and $S_i^{obs}(0)$ is the observed mechanism outcome of each treated units respective matched control). Replacing the the second term in (9), the empirical form for MATT becomes

$$MATT = E\left\{E\left[Y_i^{obs}(1)|S_i^{obs}(1) = s_1, X_i = x, T = 1\right]\right\} - E\left[f_1(S_i(0), X_i)\right]. \quad (13)$$

Similarly, the empirical form of NATT becomes

$$NATT = E\left[f_1(S_i(0), X_i)\right] - E\left\{E\left[Y_i^{obs}(0)|S_i^{obs}(0) = s_0, X_i = x, T = 1\right]\right\}, \quad (14)$$

where in $f_1(S_i(0), X_i)$ in (13) and (14) is equal to $E\left[Y_i(1, S_i(0))|S_i(1), X_i = x, T = 1\right]$ from (12).

We again emphasize the intuition behind the counterfactual of interest, which can be used in the estimation of both MATT and NATT. The regression imputation methods presented in (11) and (12) allow us to address the question, "what would the outcomes for treated units have been had their respective covariates $\left(X_i^{obs}|T=1\right)$ and influences of these covariates on outcomes (b_1) remained the same, but their mechanism taken on the values that would have been observed had they not been treated $S(0)|T=1$?" We note that the difference between $S_i^{obs}(1)|T=1$ (the observed mechanism value of treated units) and $\widehat{S}_i(0)|T=1$ (the estimated counterfactual values of treated units, had they not been treated) represents the unit-level causal effect of treatment on mechanism outcomes $(T \rightarrow S)$.

Bias-adjusted mechanism outcomes

Abadie and Imbens (2006) and Abadie et al. (2004) suggest the use of post-match regression bias adjustments in the estimation of ATT to control for remaining imbalance in matched samples. We apply a similar method in the estimation of our counterfactual mechanism outcomes to control for imbalance in our matched sample.

Post-match bias-adjustment in estimation of ATT is conducted by first running a regression of outcomes on matching covariates $Y_{T=0} = X_{T=0}\beta_C + \epsilon$. This regression estimates the impact (β_C) of the matching covariates on outcomes for the matched control sample. To impute the ATT counterfactual of interest, β_C is combined with the covariates from the treated units $X_{T=1}$ to estimate $\widehat{Y}_{BA} = X_{T=1}\beta_C$: what treated unit outcomes would have been had their matching covariates had the same influence on outcomes as the control units. Note that if matching produces perfect balance across treated and matched control units then a counterfactual based on the observed values of the matched control outcomes $(Y_{i:T=0})$ will be identical to those estimated from the regression bias adjustment procedure $\left(\widehat{Y}_{i:BA}\right)$

The estimation of our counterfactual of interest in (12) is a function of both b_1 from (11) and $\widehat{S}_i(0)|T = 1$. By Assumption 2 we can use the mechanism outcomes of the matched controls as the counterfactual for treated units. However, if imbalance in the baseline mechanism covariates remains after matching, we may be concerned that counterfactual mechanism values will be biased.[46] We, therefore, estimate our counterfactual mechanism values

$$\left[\widehat{S}_i(0)|T = 1\right] = S^{obs}_{i:T=0} + \widehat{\mu}_0(X_{i:T=1}) - \widehat{\mu}_0(X_{i:T=0}) \tag{15}$$

where $\widehat{\mu}_0$ represents the predicted values obtained from combining the coefficients from a control group regression, of mechanism outcomes on covariates, with the respective treated ($\widehat{\mu}_0(X_{i:T=1})$) or control group ($\widehat{\mu}_0(X_{i:T=0})$) covariates. This procedure estimates the influence of baseline covariates on mechanism outcomes for control units and uses these estimated values to impute what the mechanism outcomes would have been had the control units been treated.

Standard errors

To calculate the precision of our MATT estimates we base our standard error estimator on the heteroskedasticity robust matching-based estimator suggested by Abadie and Imbens (2006).[47] Our estimator is calculated in two stages to allow for heteroskedastic variances within and across treatment arms. The variance for control units (for which comparison to MATT is not meaningful) is calculated using a within treatment arm matching estimator. The Mahalanobis weighting matrix from the original matching process (used to create the matched sample) is used to

[46] If mechanism outcomes are state dependent, then imbalance is a concern. For instance, if, after matching, unprotected tracts have lower baseline roadless volume, on average, than protected tracts, change in roadless volume may be less (in absolute terms) in unprotected tracts, simply because they started with larger road networks.

[47] A function that estimates the standard errors outlined in this section was programmed in `R` 2.11.1 and is available from the author upon request.

find the nearest within treatment arm (unprotected) neighbor to estimate unit-level variances

$$\widehat{\sigma}^2_{i:T=0}\left(X_i\right) = \left(Y_i - Y_l\right)^2/2, \tag{16}$$

where Y_l represents the outcome of the nearest neighbor to unit i. The treatment level variance is then calculated

$$\widehat{V}_{T=0}\left(\widehat{MATT}\right) = \sum_{N_{T=0}} \lambda_i^2 \cdot \widehat{\sigma}^2_i\left(X_i\right), \tag{17}$$

where $\lambda_i = \#C_i/N_{T=0}$, and $\#C_i$ is the number of times that control unit i occurs in the set (was used as a match in the original matching specification).

The individual level variance for protected units is based on unit level deviations from the estimated MATT

$$\widehat{\sigma}^2_{i:T=1}\left(X_i\right) = \left(Y_i - \widehat{Y}_i(1, \widehat{S}(0)) - \widehat{MATT}\right)^2. \tag{18}$$

These unit level variances are then aggregated to calculate the treatment level (protected) variance

$$\widehat{V}_{T=1}\left(\widehat{MATT}\right) = \frac{1}{N^2_{T=1}} \sum_{N_{T=1}} \widehat{\sigma}^2_{i:T=1}\left(X_i\right). \tag{19}$$

The final MATT standard error estimate is therefore

$$\widehat{\sigma}\left(\widehat{MATT}\right) = \sqrt{\left(\widehat{V}_{T=0} + \widehat{V}_{T=1}\right)}. \tag{20}$$

Results

Empirical Estimation of MATT

We conduct two distinct analyses to estimate the MATT of our mechanisms of interest. First, the mechanisms are considered separately and the procedure outlined in the preceding sections is performed for each mechanism. Second, the mechanisms are considered jointly in the estimation of each MATT via inclusion of all mechanisms in (11) and (12).

We begin by matching protected and unprotected census tracts using one-to-one Mahalanobis covariate matching with replacement. The resulting matched set (identical to the sample used by Andam et al. (2010)) comprises 249 protected and unprotected tracts, the covariate balance can be seen in Table 4. Using post-match regression bias-adjustment, the estimated ATT is -1.27, according to the poverty index. This result indicates that census tracts with at least 10% of their area occupied by a protected area prior to 1980 had differentially greater levels of poverty reduction (lower poverty index scores) between 1973 and 2000, on average, than comparable census tracts that remained unaffected by protected areas (see Andam et al. (2010) for full details).

Counterfactual mechanism values

The counterfactual of interest necessitates estimation of mechanism outcomes for treated units, had protection not affected the mechanism. For each mechanism, estimation of the counterfactual entails a two-step process. First, we estimate a matched unprotected group regression

$$S_{i:T=0} = X_{i:T=0}\beta_{1C} + \epsilon \tag{21}$$

Table 4: Balance Results for Matched Set. Mahalanobis one-to-one covariate matching with replacement.

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	15.05	5.376	9.673	0.769	9.687	
1973	Matched	15.05	15.240	-0.187	0.013	1.64	98.07%
% Forest 1960	Unmatched	0.523	0.117	0.406	0.734	0.405	
	Matched	0.523	0.488	0.035	0.054	0.035	91.38%
% Land Use	Unmatched	0.093	0.304	-0.211	0.315	0.212	
Capacity 1,2,3	Matched	0.093	0.12	-0.028	0.060	0.03	86.84%
% Land Use	Unmatched	0.209	0.453	-0.244	0.330	0.245	
Capacity 4	Matched	0.209	0.20	0.009	0.016	0.026	96.19%
% Land Use	Unmatched	0.233	0.196	0.036	0.056	0.102	
Capacity 5,6,7	Matched	0.233	0.243	-0.011	0.019	0.034	70.98%
Distance to	Unmatched	58.53	34.87	23.670	0.286	23.62	
Major City	Matched	58.53	57.56	0.968	0.01	5.282	95.91%
Roadless	Unmatched	1113000	66830	1046000	0.321	1035000	
Volume 1969	Matched	1113000	681500	431600	0.110	440900	58.75%

where $S_{i:T=0}$ and $X_{i:T=0}$ represent the observed mechanism and baseline covariate values, respectively, of matched *unprotected* census tracts. The coefficients from (21) are then used to impute counterfactual mechanism outcomes for each mechanism

$$\left[\widehat{S}_i(0)|T = 1\right] = X_{i:T=1}\widehat{\beta}_{1C} \tag{22}$$

where $X_{i:T=1}$ are the observed covariate values of the *protected* census tracts (empirical analog to equation (15)). Observed and counterfactual mechanism values for the protected census tracts can be seen in Table 5. The imputed counterfactual mechanism values from (22) are then used to calculate the counterfactual of interest: the outcomes for protected units, had protection not affected mechanisms ($\widehat{Y}_i(1, S_i(0))$).

Columns (i) and (ii) of Table 5 list the observed and estimated counterfactual mechanism values for the protected census tracts (see Table 6 for estimates of counterfactual mechanism values when bias-adjustment is not implemented). The counterfactual for our proxy for tourism is straight forward. Of the 122 census tracts that were impacted by a protected area with a park entrance, none would

have a park entrance in the absence of protection. The estimated counterfactual for change in forest cover is significantly different from observed values as well. The average deforestation in protected census tracts between 1960 and 1986 was only 6.7%. We estimate that, had protection not affected deforestation, deforestation would have been approximately 23% (i.e., avoided deforestation between 1960 and 1986 due to the establishment of protected areas was approximately 16.3%). Finally, we observe that there was greater infrastructure development (greater reduction in roadless volume) in protected census tracts between 1969 and 1991. However, the counterfactual measures of road networks are not significantly different from observed values.[48]

Single mechanism estimation

In the single mechanism estimation the following procedure is run on each mechanism of interest independently. We first estimate the influence of covariates and mechanism on outcomes using the protected census tracts

$$Y_{i:T=1} = X_{i:T=1}\beta_{1T} + S_{i:T=1}\beta_{2T} + \epsilon \tag{23}$$

where $Y_{i:T=1}$, $X_{i:T=1}$ and $S_{i:T=1}$ are the observed outcomes, matching covariates and mechanism values for the protected census tracts, respectively. The counterfactual of interest is then estimated by obtaining the fitted values from

$$\widetilde{Y}_{i:T=1} = X_{i:T=1}\widehat{\beta}_{1T} + \widehat{S}_i\widehat{\beta}_{2T} \tag{24}$$

where $\widehat{S}_i = \left[\widehat{S}_i(0)|T = 1\right]$ are the counterfactual mechanism values from (22), thus $\widetilde{Y}_{i:T=1} = \widehat{Y}_i(1, S_i(0))$. MATT for each mechanism is calculated by subtracting the

[48] Differences are significant when counterfactual mechanism values are estimated without bias-adjustment. See the Without Mechanism Bias-Adjustment Section for results without mechanism imputation.

mean of the fitted values ($\widetilde{Y}_{i:T=1}$) from mean of the observed protected tract outcomes ($Y_i(1, S(1)) = Y_i(1)$).

Results from the single mechanism estimation strategy can be found in Columns (iii) and (iv) of Table 5. Column (iii) lists the estimated marginal impact of each mechanism ($\widehat{\beta}_{2T}$ from (23)) on poverty. Concordant with conjecture that protected areas have a positive impact on poverty by attracting tourism, we find that protected census tracts that were impacted by parks with entrances had lower poverty (by 1.04 according to the poverty index) than similar protected tracts. Because no protected tract would have been influenced by a park entrance in the absence of protection, the estimated MATT (column (iv)) is -0.492. In other words, tourism, as measured by park entrances, accounted for approximately 40% of the poverty reduction associated with the establishment of protected areas.

The marginal impact of infrastructure development also has the expected sign (Column (iii)). Our results indicate that as road networks develop (roadless volume decreases) there is an associated reduction in poverty. We estimate that, had protection not affected road development in surrounding census tracts, there would have been less development in the absence of protection. However, the difference between observed and counterfactual values is relatively small. The slightly greater road development in protected census tracts accounts for a poverty reduction (MATT) of only -0.143 (approximately 11% of the total ATT).

The results for change in forest cover reflect the conflicting impacts underlying deforestation. There is a significant difference in observed and counterfactual deforestation in protected census tracts. We estimate that, had protection not affected deforestation, over 22% of the protected census tracts, on average, would have been deforested between 1960 and 1986 (compared to 6.7% observed deforestation). Despite this stark difference the MATT of deforestation is quite

small, 0.099, and indicates that the prevention of deforestation caused by the establishment of protected areas had essentially no impact on poverty.

Table 5: Mechanism Results Using Mechanism Imputation

			Single Mechanism		All Mechanisms	
	(i) Observed Mechanism	(ii) Counterfactual Mechanism	(iii) Mechanism Coefficient	(iv) MATT	(v) Mechanism Coefficient	(vi) MATT
Park Entrance	122	0	-1.004	-0.492 (0.439)	-1.345	-0.619 (0.448)
Δ Roadless Volume	-727,579	-674,147	2.694e-06	-0.143 (0.447)	2.790e-06	-0.148 (0.449)
Δ Forest Cover	-0.067	-0.23	0.627	0.103 (0.439)	0.124	0.02 (0.45)

(Heteroskedasticity robust standard errors)

Joint mechanism estimation

In the single mechanism estimation strategy each mechanism is considered independently. However, the estimated impact (according to $\widehat{\beta}_{2T}$) of a particular mechanism may be influenced by the inclusion or exclusion of additional mechanisms in (23). By including all of the mechanism variables in (23) we allow the coefficients for each mechanism to reflect the presence of the other mechanisms. For clarity we denote the park entrance, change in roadless volume and deforestation mechanism variables as E, R and F respectively. The joint mechanism estimation analog to (23) is

$$Y_{i:T=1} = X_{i:T=1}\beta_{1T} + E_{i:T=1}\beta_{2T} + R_{i:T=1}\beta_{3T} + F_{i:T=1}\beta_{4T} + \epsilon \tag{25}$$

where all variables represent the observed values from the protected census tracts. The counterfactual of interest for each mechanism is estimated in a series of three

imputations

$$\widetilde{Y}^{E}_{i:T=1} = X_{i:T=1}\widehat{\beta}_{1T} + \widehat{E}_i\widehat{\beta}_{2T} + R_i\widehat{\beta}_{3T} + F_i\widehat{\beta}_{4T} \quad (26)$$

$$\widetilde{Y}^{R}_{i:T=1} = X_{i:T=1}\widehat{\beta}_{1T} + E_i\widehat{\beta}_{2T} + \widehat{R}_i\widehat{\beta}_{3T} + F_i\widehat{\beta}_{4T} labeleq : mech21 \quad (27)$$

$$\widetilde{Y}^{F}_{i:T=1} = X_{i:T=1}\widehat{\beta}_{1T} + E_i\widehat{\beta}_{2T} + R_i\widehat{\beta}_{3T} + \widehat{F}_i\widehat{\beta}_{4T} \quad (28)$$

where $\widehat{E_i}$, $\widehat{R}_i$ and $\widehat{F}_i$ represent the imputed mechanism values from (22) (i.e., $\left[\widehat{S}_i(0)|T=1\right]$ for the respective mechanisms). Equations (26) - (28) show that the counterfactual of interest for each mechanism is estimated by substituting the imputed mechanism values (from (22) for the mechanism of interest) into the respective equation, while leaving the covariates and complementary mechanism values at observed levels.[49] For instance, the counterfactual of interest for change in roadless volume ($\widetilde{Y}^{R}_{i:T=1}$) is calculated by plugging in the imputed counterfactual values for change in roadless volume ($\widehat{R}_i$) into the coefficients from (25), while leaving covariates ($X_{i:T=1}$) and mechanism values for park entrances (E_i) and change in forest cover (F_i) at the observed levels of protected census tract.

Results for the joint mechanism estimation strategy can be found in Columns (v) and (vi) of Table 5. We find that inclusion of all mechanisms in (25) does change the estimated influence of each mechanism (compare to Column (iii)): the coefficient on the park entrance mechanism increases in absolute terms from -1.004 to -1.345 (indicating increased poverty reduction attributable to tourism, comparatively); the coefficient on the roadless volume mechanism increases from 2.694e-06 to 2.790e-06 (indicating increased poverty reduction attributable to infrastructure development, comparatively), and; the coefficient on the deforestation mechanism decreases from 0.627 to 0.124 (indicating reduced poverty exacerbation attributable to deforestation, comparatively).

[49] A function that performs this iterative process was written in R 2.11.1 and is available from the author upon request.

Under the joint mechanism estimation we find that the MATT for the park entrance mechanism increases, in relative terms, to -0.619. This result implies that tourism associated with the establishment of protected areas accounts for approximately 49% of the estimated poverty reduction due to protection. Joint estimation also affects the MATT for the deforestation mechanism which falls to 0.02. In other words, reductions in deforestation due to the establishment of protected areas has almost no impact of poverty. Joint mechanism estimation has a trivial effect on the MATT of roadless volume which increases, in absolute terms, to -0.148.

Summary of Results

We estimate the MATT for each of our mechanisms using both a single, and joint estimation strategy. Our results indicate that, while there are some differences, the choice of strategy is not driving the results or underlying implications. However, *a priori*, we prefer the joint estimation strategy because each mechanism coefficient (and, therefore, each MATT) accounts for the presence of the other mechanisms.

Of the mechanisms we consider, tourism accounts for greatest MATT, in absolute terms, and the greatest proportion of total poverty reduction due to the establishment of protected areas (estimated in the ATT). Nearly half of the poverty reduction associated with the establishment of protected areas is accounted for by our proxy for tourism, the establishment of a park entrance within a protected area. These results are concordant with anecdotal evidence, conjecture, and findings from a previous study (Robalino and Villalobos-Fiatt 2010).

The development of infrastructure in protected census tracts has a strong poverty reducing influence as well (as measured by $\widehat{\beta}_{3T}$). However, because the establishment of protected areas did not substantially increase the road networks in

the affected census tracts, compared to our counterfactual estimates, the MATT on poverty was modest.

We find that the reduction in deforestation associated with the establishment of protected areas (compared to counterfactual levels) has essentially no impact on poverty, as measured by the MATT. By measuring the impact of reductions in deforestation on poverty, due to protection, we were hoping to capture the impact of preserving ecosystem services. However, as mentioned in the Introduction, there are likely countervailing mechanism effects of avoided deforestation, which we believe are highlighted by our results. Figure 7 presents a DAG that depicts two of the potential underlying impacts that avoided deforestation would likely have on poverty. The establishment of protected areas reduces deforestation $T \overset{(-)}{\rightarrow} F$. This causal reduction in deforestation has two impacts. First, we expect an increase in ecosystem services $F \overset{(+)}{\rightarrow} ES$ which would lead to a positive impact on poverty (poverty reduction), $ES \overset{(+)}{\rightarrow} Y$. Second, we expect the reduction in deforestation to decrease extraction profits $F \overset{(-)}{\rightarrow} EP$ which would lead to a negative impact on poverty (poverty exacerbation), $EP \overset{(-)}{\rightarrow} Y$.[50]

Considering the concepts from the DAG in Figure 7, our estimated MATT for change in forest cover and results from previous studies in Costa Rica (Ferraro and Hanauer 2011, Ferraro et al. 2011), we believe that ecosystem services can be shown to exhibit a positive MATT on poverty (reduces poverty). First, assume that the only causal mechanism relationship between protection, deforestation and poverty is $T{\rightarrow}F{\rightarrow}EP{\rightarrow}Y$ (i.e., ecosystem services either are not affected by avoided deforestation or do not impact economic outcomes). Our results show that the establishment of protected areas is associated with a causal reduction in

[50]This mechanism channel captures the concern the land-use restrictions associated with the establishment of protected areas may impose economic hardship by prohibiting extractive activities. Avoided deforestation provides an indication that land-use laws were binding and, therefore, $F \overset{(-)}{\rightarrow} EP \overset{(-)}{\rightarrow} Y$ is likely a valid channel.

deforestation $T \overset{(-)}{\rightarrow} F$. Ferraro and Hanauer (2011) and Ferraro et al. (2011) show that increases in avoided deforestation, associated with the establishment of protected areas, generally result in lower levels of poverty reduction or, in some cases, poverty exacerbation. This association is due to the fact that protected areas reduce profitable economic activities such as resource extraction. Therefore, by reducing deforestation we expect the establishment of protected areas to reduce extractive profits $F \overset{(-)}{\rightarrow} EP$ thereby causally reducing poverty outcomes (exacerbating poverty) $EP \overset{(-)}{\rightarrow} Y$. Taken together, if $T \overset{(-)}{\rightarrow} F \overset{(-)}{\rightarrow} EP \overset{(-)}{\rightarrow} Y$ was the sole deforestation mechanism process then we would expect the change in forest cover MATT to be positive (indicating poverty exacerbation to a higher magnitude).

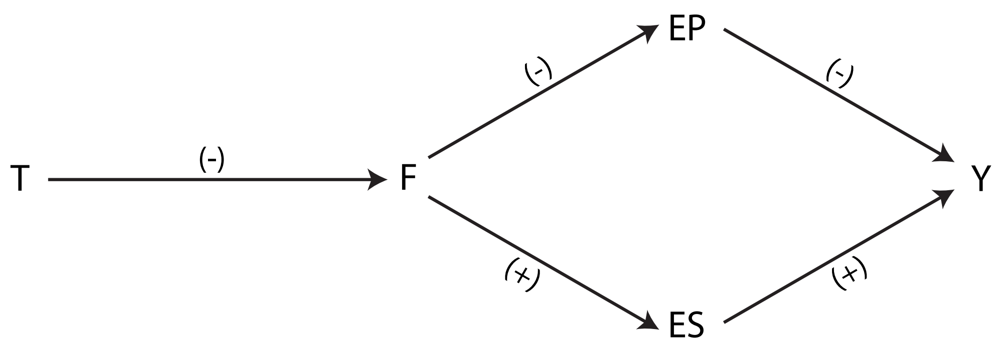

Figure 7: Directed Acyclic Diagram (DAG) depicting two of the potential underlying effects (and direction) of the change in forest cover mechanism. The the establishment of protected areas reduces deforestation $T \overset{(-)}{\rightarrow} F$. This causal reduction in deforestation has two impacts. First, we expect an increase in ecosystem services $F \overset{(+)}{\rightarrow} ES$ which would lead to a positive impact on poverty (poverty reduction), $ES \overset{(+)}{\rightarrow} Y$. Second, we expect the reduction in deforestation to decrease extraction profits $F \overset{(-)}{\rightarrow} EP$ which would lead to a negative impact on poverty (poverty exacerbation), $EP \overset{(-)}{\rightarrow} Y$. The relative magnitude of these countervailing effects determine the estimated MATT.

Incorporating the preceding logic into the model established in the DAG in Figure 7, our posited relationship $T \overset{(-)}{\rightarrow} F \overset{(+)}{\rightarrow} ES \overset{(+)}{\rightarrow} Y$ must hold. In other words, because preventing deforestation prevents economically beneficial extractive activities, the buttressing of ecosystem services associated with the establishment of protected areas must provide countervailing poverty reducing impacts.

Robustness

Without mechanism bias-adjustment

In the Counterfactual Mechanism Values Section we motivate and describe the use of bias-adjustment techniques to impute counterfactual mechanism values. We argue that, like the use of post-match regression bias-adjustments in the estimation of average treatment effects, this technique provides a (more) unbiased estimate of counterfactual mechanism values when imbalance persists (especially in baseline mechanism covariates) post-matching. We re-estimate all MATTs without using bias-adjustment, the results can be found in Table 6. As expected, given the purpose of the bias-adjustment procedure, we find little difference in counterfactual values for change in forest cover (compare Column (ii) in Tables 5 and 6), for which a high degree of balance in baseline measures is achieved (see Table 4). In addition, because of the binary nature of the park entrance mechanism, the counterfactual values are identical with, and without, bias-adjustment. However, the counterfactual values for the roadless volume mechanism differ substantially. Without bias-adjustment the estimated counterfactual change in roadless volume is only -447,024 (compared to -674,147 with bias-adjustment). In turn, there is a much larger difference between observed and counterfactual roadless volume mechanism values which, thus, leads to much larger, in absolute terms, estimate of the MATT (-0.7827). In other words, by not using bias-adjustment the estimated proportion of poverty reduction in the ATT attributable to protections causal effect on roadless volume, changes from approximately 11% to 61%. These results highlight that infrastructure development has a large influence on poverty reduction (as measured by $\widehat{\beta}_{3T}$). However, the magnitude of the associated MATT is determined by the counterfactual mechanism value, which we believe is best estimated using the

bias-adjustment framework that we outline in the Counterfactual Mechanism Values Section.

Table 6: Mechanism Results Without Mechanism Imputation

			Single Mechanism		All Mechanisms	
	(i) Observed Mechanism	(ii) Counterfactual Mechanism	(iii) Mechanism Coefficient	(iv) MATT	(v) Mechanism Coefficient	(vi) MATT
Park Entrance	122	0	-1.004	-0.492 (0.439)	-1.345	-0.619 (0.546)
Δ Roadless Volume	-727,579	-447,024	2.694e-06	-0.756 (0.54)	2.790e-06	-0.7827 (0.54)
Δ Forest Cover	-0.067	-0.223	0.627	0.099 (0.439)	0.124	0.019 (0.55)

(Heteroskedasticity robust standard errors)

LNATT

The estimation of MATT and NATT requires the imputation of counterfactual mechanism values which are, by definition, unobserved. Our data provide a unique opportunity to estimate the causal effects of protection *net* of tourism under less restrictive assumptions than those used in the main analyses. We exploit the fact that some protected census tracts are observed in the absence of a park entrance mechanism. For this subset of the data $S_i(1) = S_i(0)$ by definition. In other words, the potential park entrance mechanism value for protected units that did not receive an entrance is same under protection as it would have been in the absence of protection ($S_i(1) = S_i(0) = s_0$). Therefore, we can identify this principal stratum ($\{S_i(1) = S_i(0) = s_0\}$) without invoking Assumption 2 or 3. In addition, we observe $Y_i(1, S_i(0))$ for this subset of the data and, therefore, do not need Assumption 4 to impute the counterfactual of interest.

The local NATT (LNATT) can be estimated[51]

$$LNATT = E\{E[Y_i(1, S_i(0)) - Y_i(0, S_i(0))|S_i(1) = S_i(0) = s_0]\} \tag{29}$$

for the subset of data in the principal stratum $\{S_i(1) = S_i(0) = s_0\}$ (Flores and Flores-Lagunes 2011). The fact that we observe protected census tracts that were not affected by a park entrance means that we can take the simple difference in these protected tract outcomes $(Y_i(1, S_i(0)))$ and their matched controls $(Y_i(0, S_i(0)))$, both of which are observed in the data. We estimate the LNATT for this subgroup to be -0.6122. Flores and Flores-Lagunes (2011) note that the LNATT represents the local ATT (LATT) for this subgroup because there is no mechanism effect for this group so $Y_i(1,(S_i(0)) = Y_i(1)$. Therefore, under Assumption 1, LNATT=LATT=$E[Y_i(1) - Y_i(0)|X_i]$ for this subgroup.

We note that the estimated LNATT for park entrances is very close to the NATT from the main analysis (-0.6122 and -0.659, respectively). We believe that the similarity between the two estimates provides evidence that the assumptions and methods employed in the main analyses are providing unbiased estimates of the respective mechanism effects. We can make further comparisons to the MATT estimates using the estimated LNATT and an additional assumption of constant individual net effects (Flores and Flores-Lagunes 2011)

Assumption 5 $Y_i(1, S_i(0)) - Y_i(0, S_i(0)) = C$, *for all* i.

Under this assumption we can define LNATT=NATT and, therefore, estimate MATT=ATT-LNATT. Using this framework, the estimate of park entrance MATT (-0.6658) is very close to the estimate from the main analysis (-0.619).

[51] This framework follows directly from the framework for the local net average treatment effect (LNATE) established by Flores and Flores-Lagunes (2011).

Discussion

Recent impact evaluations from developing countries have found that the proliferation of protected areas in the past several decades has been associated with poverty reduction in surrounding communities (Canavire-Bacarezza and Hanauer 2011, Andam et al. 2010, Sims 2010). The findings are contrary to the expectations that the land-use restrictions, through which protected areas achieve their environmental goals, restrict economically beneficial activities. While the results from these studies are important to understanding the policy impacts, in the context of protected areas, it is arguably more important to understand through what mechanisms protected areas affect economic outcomes than it is to estimate the overall effect. Unfortunately, the methodologies employed in previous studies are not suited to identify or quantify the potential mechanisms through which protected areas affect economic outcomes.

Using recently developed quasi-experimental methods, and rich biophysical and demographic data from Andam et al. (2010), we quantify the mechanistic impacts of tourism, infrastructure development and ecosystem services on poverty due to the establishment of protected areas in Costa Rica prior to 1980. To capture the causal effects of our respective mechanisms we use the establishment of park entrances, changes in road networks and deforestation as proxies. Our results indicate that approximately 50% of the poverty reduction estimated by Andam et al. (2010) can be attributed to tourism. Conversely, infrastructure development played a negligible role in poverty reduction. Finally, although the mechanistic impact of avoided deforestation is near zero, we argue that, given the negative economic impacts associated with the prevention of deforestation, enhanced ecosystem services (due to the establishment of protected areas) likely had a positive effect of poverty reduction.

Our results offer insight to potential policy actions that might complement the establishment of protected areas. For instance, in support of common conjecture, our results suggest that the promotion of tourism concurrent with the establishment of protected areas may have beneficial poverty effects. Given that one of the goals set forth by the 5th World Parks Congress is that the establishment of protected areas should at least do no economic harm, greater understanding of the mechanisms through which protected areas affect poverty is needed. Future studies in Costa Rica should obtain less coarse measures of mechanism variables in order to more accurately measure the impacts. This will facilitate, for instance, a more precise measure of the impacts that enhanced ecosystem services, due to the establishment of protected areas, have on economic outcomes.

There is likely a great deal of heterogeneity in the overall impacts and the mechanism impacts of protected areas across countries. Therefore, caution should be taken in the extrapolation of these results to protected area networks elsewhere. Future studies should apply similar methodologies in other countries. Indeed, to truly understand the mechanisms through which protected areas affect poverty, the evidence base will need to built on a country-by-country basis.

Chapter IV

Estimating the Impacts of Protected Areas on Poverty in Bolivia

Introduction

Protected areas are an important tool for the global conservation of ecosystems and biodiversity (MEA 2005). Presently, approximately 13% of the world's terrestrial surface is covered by some form of protected area (WDPA 2009). The sheer scale of the global coverage of protected areas highlights the importance of understanding their underlying impacts. Unfortunately, there is little empirical evidence on the environmental impacts of protected areas and even less on the socioeconomic impacts of protected areas on surrounding communities (Coad et al. 2008). The socioeconomic implications of establishing protected areas are of particular interest given the high degree of overlap between areas of remaining biodiversity (i.e., areas likely to be targeted for protection) and poverty (Sachs et al. 2009). This raises concerns from poverty advocates that achieving environmental goals may come at the expense of the populations impacted by such policy (Coad et al. 2008, Adams et al. 2004).

The dearth of quality evidence on the impacts of protected areas fuels a general debate regarding the relationship between areas protected by environmental law and the socioeconomic outcomes in surrounding areas. Conservationists see the establishment of protected areas as essential to global environmental stability, whereas poverty advocates argue that, while the benefits from protecting these

areas are paid to all, the costs are borne only by those proximal to the areas (Coad et al. 2008, Wilkie et al. 2006, Adams et al. 2004). This argument concerns the land-use laws associated with protected areas that restrict economic development by preventing forms of profitable activities such as the exploitation of natural resources and agricultural cultivation (Coad et al. 2008, Fleck et al. 2006).

There have been few empirical studies with the proper data and methodologies to accurately estimate the socioeconomic impacts of protected areas on local economies, especially in developing nations (exceptions include: Ferraro and Hanauer (2011), Ferraro et al. (2011), Andam et al. (2010), Sims (2010), Robalino and Villalobos-Fiatt (2010)).[52] Most studies have either been *ex ante* estimates of future costs and benefits, or *ex post* studies of observed states of welfare (Andam et al. 2010, Wilkie et al. 2006). The ability to empirically measure the socioeconomic impacts of protected areas is complicated by the non-random nature in which areas are assigned to protection. The presence of such selection issues necessitates the use of sophisticated research design and methodologies, absent most previous studies. Further confounding the process is the fact that most developing nations do not have sufficiently rich data sets with which to measure pre- and post-treatment poverty outcomes.[53]

Bolivia is an apt setting for evaluating the impacts of protected areas on poverty in surrounding communities. Bolivia is one of the most biodiverse countries in the world.[54] Yet despite having a wealth of natural resources, Bolivia is one of the

[52]To our knowledge, there is no study that examines the relevance of protected areas on welfare in Bolivia.

[53]See Ferraro (2008) for a discussion on the components of a quality socioeconomic impact evaluation which include: 1) An appropriate measure of welfare; 2) observations on outcomes and pertinent covariates for both pre- and post-treatment; 3) relevant indicators for both treatment and control units, and; 4) observations of pretreatment covariates that affect both selection into treatment and socioeconomic outcomes.

[54]Bolivia is one of the 15 most biologically diverse countries in the world. It is recognized as one of the 11 nations with the greatest diversity of flora (about 20 thousand species) and one of the top 10 most abundant in terms of birds (1,400 species) and mammals (356 species). Information provided by the Protected Areas National Service of Bolivia SERNAP (2009). Moreover, according

poorest countries in Latin America, with poverty levels upwards of 60% (UDAPE 2006). Bolivia also has an extensive protected area network that was made effective by an identifiable restructuring of the existing system in 1992, followed by a surge in proliferation of new protected areas throughout the 90s. Moreover, Bolivia has rich biophysical and socioeconomic data that predate the effective establishment of protected areas.

Using rich biophysical and socioeconomic data, and quasi-experimental methods, we ask, "what would poverty outcomes in Bolivian communities affected by protected areas have been had protected areas not been established?" We find no evidence that communities affected by protected areas established between 1992 and 2000 fared any worse, between 1992 and 2001, than similar communities that remained unaffected by protected areas. In fact, all of our point estimates indicate that protected communities had differentially greater levels of poverty reduction. We find that these results are robust to a number of econometric specifications, sensitivity analyses, spillover analyses and placebo studies. Our results are concordant with findings of poverty alleviation due to the establishment of protected areas in Costa Rica (Andam et al. 2010) and Thailand (Andam et al. 2010, Sims 2010). However, unlike previous studies (Andam et al. 2010, Sims 2010) our results indicate that naïve (uncontrolled) comparisons of protected and unprotected communities leads to an overestimation of the poverty alleviation associated with the establishment of protected areas. Accordingly, our results underscore the fact that protected area impacts are likely not generalizable and, therefore, the importance of country-level protected area impact evaluations. To that end, our results add to the scientific body of knowledge on the socioeconomic impacts of protected areas in developing countries, which is exceedingly sparse.

to an UNESCO report, Bolivia has the largest water reserves in Latin America and ranks 6^{th} in the world in terms of tropical moist forest resources (the third in the continent after Brazil and Mexico).

Background

Protected areas in bolivia

The evolution of the establishment and enforcement of protected areas in Bolivia is complex but can be defined, coarsely, by two periods: pre- and post-1992. A non-trivial amount of Bolivia's area was designated for protection between the late 30's and early 90's.[55] However, the criteria for establishing these areas were not uniform, or systematic. Most were established with little technical background and absent the participation of local actors (SERNAP, 2007). Furthermore, there was a lack of recognition and requisite enforcement within these areas, a phenomenon commonly referred to as, "paper parks" (e.g., Bruner et al. (2001)).

The progression of the theme of conservation and the consequent international commitments undertaken by countries in the early 90's, after the Rio Conference, led to the development of policy and institutional foundations related to the then new paradigm of sustainable development and environmental care. Thus, under the Environment Law (Law 1333), the National System of Protected Areas in Bolivia (SNAP) was created in 1992, defined as natural and cultural heritage of the State and public and social interest.

The Law 1333, defines protected areas in Bolivia as "natural areas with or without human intervention, declared under state protection by law, in order to protect and preserve the flora and fauna, genetic resources, natural ecosystems, watersheds and values of scientific, aesthetic, historical, economic and social interest, in order to conserve and protect natural and cultural heritage of the country."[56]

[55] The 10 protected areas that were established prior to the 1990's cover 5,917,638 ha which is approximately 1/3 of the total protected area.

[56] One of the most important characteristics of Bolivia's protected areas is their compatibility with the existence of traditional indigenous people (Environment Law 1333, Section 60-65). Since its initiation in 1992, the National System of Protected Areas in Bolivia has been designed with a participatory approach, recognizing that the areas are occupied and are ancestral territories of indigenous populations. Therefore, the participation of local people is a fundamental in the main aspect of the system.

The consequence of Bolivia's history of protection and Law 1333 is that, even for areas designated prior to 1992, the effective establishment date of Bolivia's protected area system was 1992 (and thereafter). The identification of this administrative recognition and enforcement allows us to use 1992 as a baseline, pre-treatment year, after which the intervention impact of protected areas can be measured.

Related literature

There have been only a handful of studies on the socioeconomic impacts of protected areas that properly control for their non-random establishment. Further, there is no formal literature regarding the impacts of protected areas on socioeconomic outcomes in Bolivia. The closest study to this kind for Bolivia (Yáñez 2006) examines the potential effects of three protected areas on poverty in Bolivia. Based on household surveys carried out near these protected areas, the author finds a small positive effect of protected areas on poverty. Nevertheless, this study not only has a small sample of protected areas, but also presents some methodological drawbacks such as selection bias and sample selection.

The most comparable study to ours is a quasi-experimental analysis of protected areas in Costa Rica (Andam et al. 2010). The authors designate census tracts (segmentos) with 10% or more of their areas protected, as treated. They then use matching techniques to construct a counterfactual group that is similar along pretreatment dimensions to the treated census tracts. The authors' calculation of average treatment effect on the treated (ATT) provides evidence that census tracts with protected areas that were established prior to 1980 had differentially greater levels of poverty reduction between 1973 and 2000 than comparable unprotected census tracts.

In a similar study Sims (2010) uses a continuous measure of the percent of land area protected within Thailand sub-districts to measure the marginal effect of

protected areas (IUCN category I & II) on a poverty headcount ratio. The author compiles an extensive set of pre- and post-treatment biophysical covariates. However, the outcome variable is only available for the contemporaneous period, therefore, the baseline levels of poverty are unknown. The results of the study show that when baseline geographic and development variables are controlled for, sub-districts with more protected area displayed lower poverty levels than comparison districts.

Robalino and Villalobos-Fiatt (2010) explore how national parks affect local wages in Costa Rica and how these effects vary within different areas of a park and among different social groups. They use highly disaggregated geographic references, and find that parks' effects on wages vary according to economic activity and proximity to the entrance of the park. Workers close to entrances receive higher wages and are employed in higher-paid, non-agricultural activities.

Several studies from the United States have shown no effect of protected areas on economic outcomes at the county level. In two studies, (Lewis et al. 2002, 2003) use a simultaneous equations framework to examine the county level effects of protected areas (publicly owned land designated for preservation and multiple use) in the Northern Forest Region of the United States on migration, on employment and wage composition. A broader study by (Duffy-Deno 1998) uses a cross-section of intermountain western counties of the United States to determine the effect of protected areas (Wilderness, Forest Service and BLM land) on population and employment densities. The author finds no significant effect on either outcome of interest. However, all of these studies suffer from the lack of a true baseline, given that all of the protected areas were designated decades prior to the first census observations.

Data and Methods

Data

We employ three categories of spatial and demographic data in this study: (1) temporally distinct boundary mappings of terrestrial protected areas, (2) boundary mappings of municipalities for the 1992 and 2001 censuses and the underlying demography, and; (3) key biophysical characteristics believed to jointly affect the establishment of protected areas and poverty.

The 1992 and 2001 census data were obtained from the Bolivian National Statistical Office (INE). The census provides information that allows us to estimate socioeconomic indicators at municipal level, such as structural poverty measures, education, employment, housing, indigenous populations and health. Information regarding Bolivia's protected areas and their boundaries was obtained from Servicion Nacional de Areas Protegidas (SERNAP) and the World Database on Protected Areas (WDPA). Further geographic data (e.g., road networks, digital elevation models, cities, forest cover, etc.) were obtained from NASA, Conservation International and Bolivian forest regulation office (Superintendencia Forestal).

Unit of analysis

The unit of analysis for our study is the municipality. The municipality is the penultimate political boundary in terms of disaggregation, second to the comunidad. Bolivia comprises 326 municipalities with an average area of 325,083 ha (about the size of Rhode Island; range: 28,928 – 1,298,121 ha). The maps in Figures reffig:bol1 and reffig:bol2 show that the municipalities in the mountainous and altiplano regions (southwest) tend to be smaller than those in the lowlands (east northeast; see Figure reffig:bol3 for the topography of Bolivia).

Table 7: Baseline Characteristics for Protected and Unprotected Municipalities

Variable	Description	Status	Mean	Median	Std	Min	Max
Poverty Index 1992	Asset-based poverty index	Unprotected	0.74	1.20	2.08	-5.89	3.97
	for 1992	Protected	0.68	1.27	1.96	-5.08	3.56
Poverty Index 2001	Asset-based poverty index	Unprotected	-0.45	-0.05	2.01	-6.59	2.94
	for 2001	Protected	-0.94	-0.98	1.90	-5.47	2.35
NBI 1992	% population with unmet	Unprotected	91.05	95.73	11.60	44.21	100.00
	basic needs 1992	Protected	87.52	93.28	14.07	45.76	99.67
NBI 2001	% population with unmet	Unprotected	86.03	93.82	17.00	19.08	100.00
	basic needs 2001	Protected	76.70	83.31	21.60	23.83	99.84
Area (m)	Municipality area, square	Unprotected	2.179E+09	9.313E+08	3.652E+09	1.235E+07	3.511E+10
	meters	Protected	6.937E+09	2.214E+09	1.270E+10	2.179E+08	7.136E+10
Change in Forest Cover	% deforestation in municipality	Unprotected	-0.01	0.00	0.05	-0.31	0.00
1976-1991	1976-91	Protected	-0.02	0.00	0.05	-0.26	0.00
Percent Forest	% of municipality under	Unprotected	0.19	0.00	0.31	0.00	0.99
Cover 1991	forest cover 1991	Protected	0.48	0.58	0.34	0.00	0.93
Average Distance	Avg. dist. to city from each 1ha	Unprotected	107999	94104	71562	6828	384000
to a City (m)	parcel within each municipality	Protected	142790	104553	131397	8400	589100
Average Elevation (m)	Average elevation of each 1ha	Unprotected	2713	3265	1470	143	4478
	parcel within each municipality	Protected	1825	1695	1357	149	4589
Average Slope (pct)	Average slope of each 1ha parcel	Unprotected	18.97	16.58	14.82	0.91	59.35
	within each municipality	Protected	23.89	27.38	14.31	1.29	55.74
Roadless Volume (m)	Sum of the product of area and	Unprotected	8.350E+13	9.801E+12	2.604E+14	1.342E+10	2.53E+15
	dist. to road, each 1ha parcel	Protected	2.979E+14	2.544E+13	8.477E+14	6.059E+11	5.61E+15

Notes: Sample includes 56 protected and 252 unprotected units. 21 “marginally” protected units (those with between [0.01,0.1) of their area occupied by a protected area) are removed from the sample.
*The municipalities are segmented into square parcels with sides of 100m. Measurements are made for each parcel then averaged within municipalities.

Protected areas and treatment assignment

Bolivia has 23 protected areas that cover a total area of 17,131,507.48 ha or ˜16 percent of Bolivia's terrestrial area (GIS calculations). The average size of a protected area is 475,875 ha (range: 221 – 2,919,143 ha). To determine the socioeconomic impacts of protected areas we must identify the municipalities that are spatially influenced by protected areas. We use GIS to determine the proportion (percentage) of each municipality that is occupied by a protected area established between 1992 and 2000. Of the 87 municipalities that intersect with one or more protected area, the average area designated as protected is 29.3% (range: 0.00007 – 100%). In order to assign municipalities a binary indicator of treatment we must establish a threshold (percentage), above which municipalities are considered protected. Our initial threshold is established at 10%.[57] In order to reduce potential bias to our estimates we need to ensure that we are not comparing protected municipalities with marginally protected municipalities (doing so would likely serve to weaken estimates of treatment effect). We, therefore, drop municipalities with percentage overlap along the interval [0.1, 10). According to our assignment rule, there are 56 municipalities that are considered protected. The percentage of overlap within these protected municipalities ranges from 10.26 – 100% (mean = 43.9%, median = 39.14%). We are left with 256 municipalities that are considered unprotected.

[57] We use the 10% threshold in accordance with previous studies (Andam et al. 2010, Ferraro and Hanauer 2011, Ferraro et al. 2011). A 10% threshold was chosen because one the goals of set forth by 4^{th} World Congress on National Parks and Protected Areas was to protect 10% of the worlds ecosystems (Andam et al. 2010). We test the robustness of our results to this protection threshold by defining alternative thresholds at 5, 20, 30 and 50%. We find that our results are robust to these alternative thresholds (see Appendix C for full threshold results).

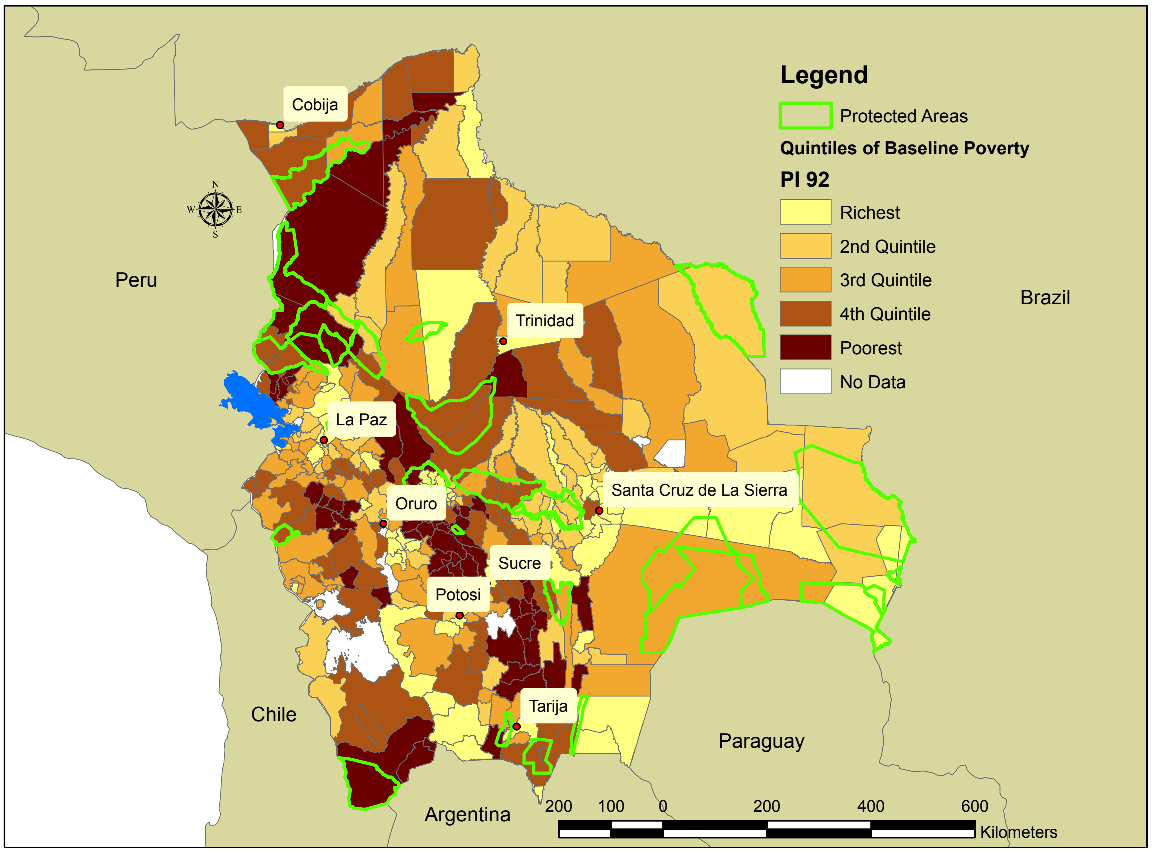

Figure 8: Map of protected areas, major cities and quintiles of poverty in 1992 according to poverty index.

Covariates of interest

To isolate the causal effect of protected areas on poverty, we compile a set of observable covariates that jointly affect the establishment of protected areas and poverty outcomes (and assume that all unobservables do not exhibit joint influence). These covariates are used in our analyses to control for the observable differences between protected and unprotected municipalities, therefore, isolating the impact of protection.

Distance to major city. Cities tend to be the nodes of major markets, economic activity and opportunity. Protected areas are often located distant from major cities, where the opportunity cost of land is lower (Sims 2010, Joppa and Pfaff 2009). We calculate the average distance from each municipality to the nearest city (each municipality is broken into 1 ha parcels and the average euclidean distance from the set of parcels within each municipality to the nearest city is calculated using GIS). Cities included in the measurement are the state capitals: La Paz, Sucre, Cochabamba, Cobija, Trinidad, Oruro, Potosi and Santa Cruz. Table 7 shows that protected municipalities are significantly farther from cities, on average, than unprotected municipalities. This is consistent with previous findings (e.g., Andam et al. 2010, Sims 2010, Joppa and Pfaff 2009).

Roadless volume. Access to roads increases access to markets and other resources (reducing transportation costs, etc.). In addition, roads serve as a good indicator of the level of infrastructure development and urbanization. Previous country-level studies have found that protected areas tend to be located in areas with sparse road networks (Andam et al. 2010, 2008). To control for baseline measures of these influences we calculate roadless volume (Watts et al. 2007). Roadless volume is an aggregation of the euclidean distance to a road for each land parcel within a municipality, adjusted for the size of the land parcel. Roadless volume is calculated by summing the product of the area of each land parcel (1 ha

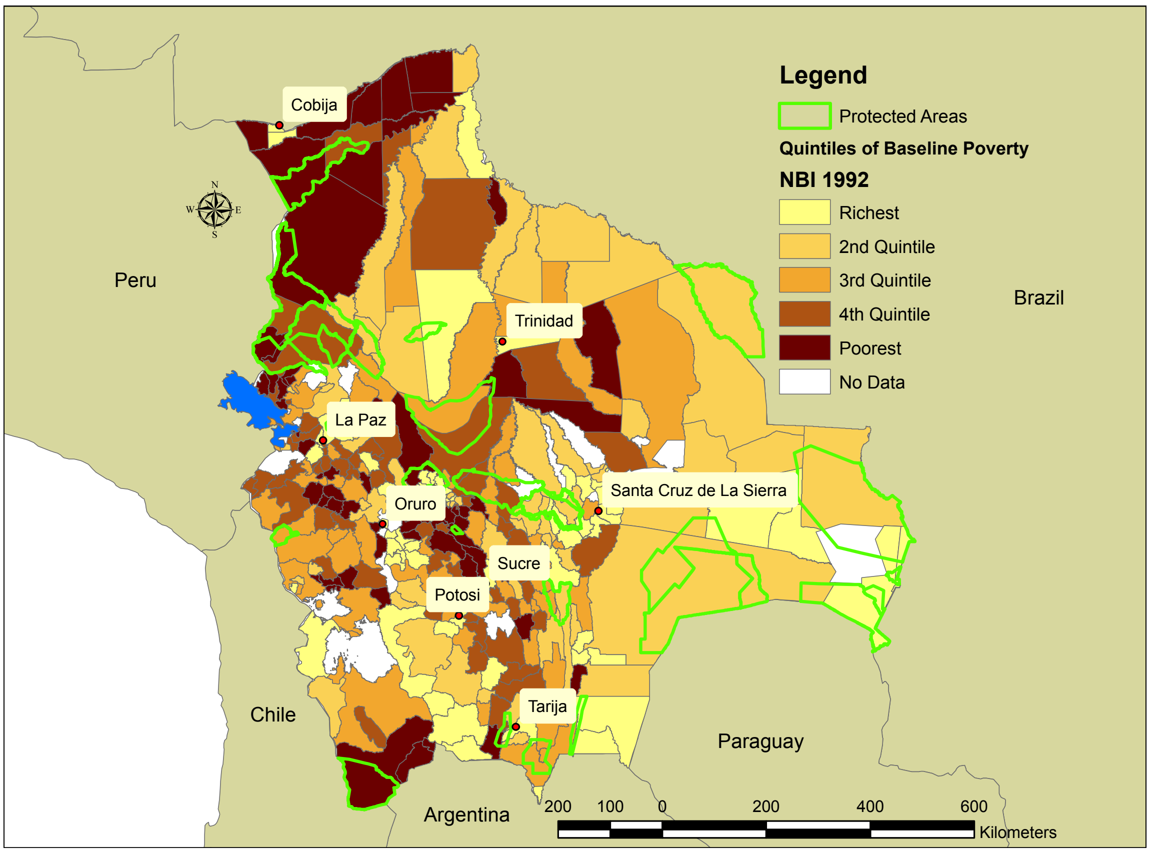

Figure 9: Map of protected areas, major cities and quintiles of poverty in 1992 according to NBI.

in this case) and the distance of that parcel to the nearest road (1992). Therefore, higher measurements of roadless volume indicate fewer road networks within a municipality. Table 7 shows that roadless volume is greater within protected municipalities, which is consistent with previous studies (e.g., Andam et al. 2010).[58]

Elevation and slope. Productivity of land, especially related to agricultural productivity, plays a large role in economic development. Low slope, low elevation land tends to be more suitable for agriculture and general development (lower production, extraction and development costs). Previous studies, both globally (Joppa and Pfaff 2009) and at the country-level (Sims 2010), have found that protected areas tend to be placed on land that is relatively steep and high elevation. It is, therefore, important to control for the average slope and elevation of municipalities. Bolivia presents somewhat of a unique case, however. Bolivia is characterized by a dichotomous landscape in that the country is defined (in general) by the highlands and lowlands. Table 7 shows that the slope in protected municipalities is greater in unprotected municipalities (expected), however the average elevation is lower (on average) within these protected municipalities (unexpected).

Forest cover. Protected areas tend to be placed on forested lands (Andam et al. 2010, Sims 2010). In addition, forests represent potential for economic opportunities (timber, fuelwood, etc.). We therefore calculate the percentage of each municipality covered by forest in 1991. Table 7 shows that protected municipalities contained significantly more forested areas at baseline.

[58]One may be concerned that disparate municipality areas across protected and unprotected units might confound the estimates of protected area impacts (e.g., may be correlated with urbanization or other unobservable). However, roadless volume is nearly perfectly correlated with the area of respective municipalities (Pearson correlation coefficient is 0.905), mitigating such concerns.

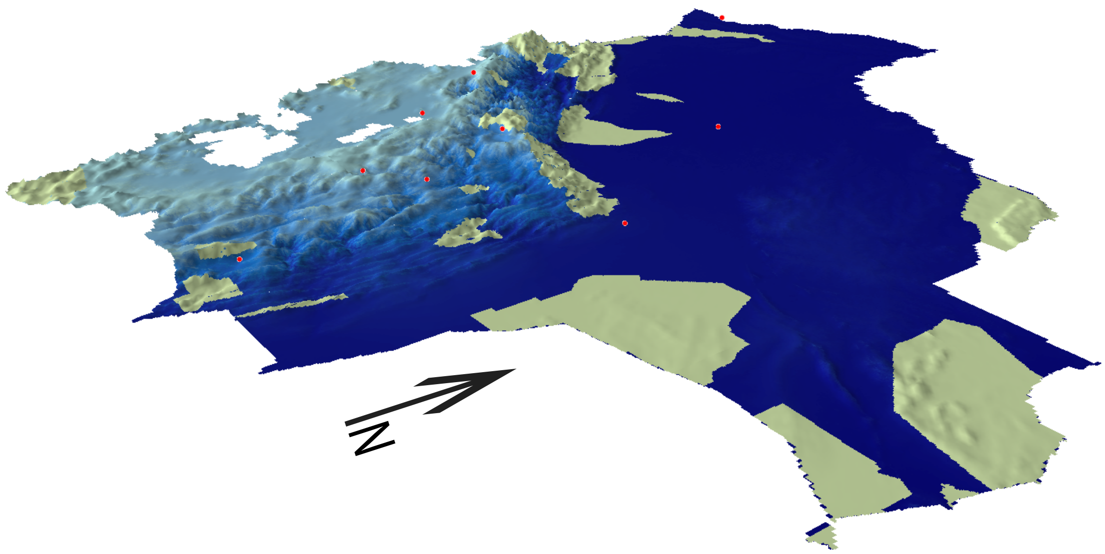

Figure 10: Digital elevation model of Bolivia, with protected areas draped (green). While a majority of the designated protected areas lie in relatively mountainous regions, there is a large area protected in the east that accounts for the relatively low average slope of protected municipalities. The red points represent the location of major cities listed in Figures 8 and 9.

Poverty

In order to estimate the impact of protected areas on poverty, an outcome that objectively measures levels of socioeconomic welfare is necessary. In the absence of a universal metric such as income, we are tasked with developing metrics that adequately capture socioeconomic outcomes. We employ two measures of poverty in our analyses: an asset based poverty index (PI) and Necesidades Basicas Insatisfechas (NBI). Both poverty indicators are measured during the 1992 and 2001 census years and serve two purposes. First, the PI and NBI in 2001 serve as the outcome of interest to measure differences in poverty between protected and unprotected municipalities. Second, the PI and NBI in 1992 serve as controls for baseline states of poverty. By ensuring that we compare protected and unprotected municipalities with similar baseline poverty characteristics, we improve the probability that these units share similar poverty trajectories prior to the establishment of protected areas.

Table 8: Eigenvectors from principal component analysis

Eigenvectors, Pooled Variable	EigenV
Adult men in population*	-0.02836
Households without bathroom*	0.34984
Households that use fuelwood for cooking*	0.39719
Households with dirt floors*	0.39399
Low-quality houses*	0.33074
Households without electricity*	0.46972
Illiterate population*	0.17906
Population employed with salary*	-0.06499
Average persons per bedroom	0.00384
Households without access to public water*	0.39091
Households without sewer or septic*	0.20994
Average years of education	-0.02452

Notes: Census data from 1992 and 2001 are pooled to measure average influence of assets across time.
* Indicates that variable is measured as a percentage.

Table 9: Mean Asset Values by Deciles of 1992 Poverty Index

	Richer			Deciles of PI 1992					Poorer	
Variable	1	2	3	4	5	6	7	8	9	10
pct.men.92	57.77	55.96	55.78	54.74	54.45	56.5	55.1	54.56	54.3	57.9
pct.wo.bath.92	43.96	66.88	61.13	68.52	77.72	86.13	82.14	85.15	92.49	99.26
pct.fuelwood.92	25.59	46.61	70.06	69.85	65.79	74.67	82.98	89.47	91.4	99.2
pct.dirt.floor.92	28.77	49.74	63.72	72.77	76.89	79.72	80.52	81.61	92.19	91.7
pct.low.qual.house.92	15.23	23.07	37.02	41.09	49.08	47.7	57.57	58.63	68.46	82.57
pct.wo.elct.92	26.35	45.1	63.42	73.31	82.13	86.89	94.17	94.65	97.19	99.87
illit.92	12.5	19.32	23.11	27.04	24.7	33.12	30.24	37.22	47.19	67.91
emp.92	95.8	91.28	90.17	88.66	89.6	88.09	85.62	84.93	81.45	61.62
avg.person.room.92	3.28	3.26	3.48	3.52	3.42	3.34	3.53	3.57	3.77	4.09
pct.nowater.92	27.13	46.12	59.33	69.38	75.83	75.63	77.57	83.98	90.18	97.56
pct.nosewer.92	65.05	87.19	90.97	93.19	96.16	96.98	97.34	98.5	99.03	99.87
avg.edu.92	6.77	5.1	4.66	4.2	4.42	3.68	3.77	3.36	2.72	1.45

We use two separate measures of poverty to buttress the robustness of our analyses. Although the PI and NBI are both designed as measures of poverty, they capture different aspects that contribute to poverty. Given that we do not have a more direct metric for poverty, such as income, we feel that it is important not to limit our analyses to a single proxy.

Poverty Index (PI).[59] Our PI is an asset based poverty index founded on household responses to the 1992 and 2001 censuses. The index is constructed using a principal component analysis (PCA). The primary purpose of the PCA is to measure the influence of a vector of variables on a latent outcome, poverty. The relative influence of each component variable is measured by the eigenvectors (factor loading) calculated from the variance/covariance matrix underlying the component variables. The eigenvectors are combined with the relative municipal-level variation in assets to calculate a municipal-level PI.

Table 8 lists the variables used in the construction of the poverty index, and the eigenvectors associated with each asset.[60] These eigenvectors (from the first

[59]A similar asset based poverty index was developed for Costa Rica (Andam et al. 2010, Cavatassi et al. 2004) and used by the Mexican government in the analyses of the PROGRESA program (cited by Cavatassi et al. (2004).

[60]Our analyses are designed to measure changes in poverty over time. To ensure that our poverty indexes are comparable across time we pool the asset data from the 1992 and 2001 census data similar to Cavatassi et al. (2004) (see also (Filmer and Pritchett 2001)). By pooling the data for the PCA we estimate the mean influence of each asset across time, allowing the variation in assets between time periods to drive the estimated changes in poverty.

component) account for over 60% of the variation in the asset variables and provide the factor scores F_j for asset $j \in \{1, 2, ..., J\}$ which indicate the weight and direction of the influence each asset a_j exerts on the PI. These factor scores are combined with observed asset levels to formulate the PI for municipality $i \in \{1, 2, ..., N\}$,

$$PI_i = \sum_{1}^{J} F_j \left[\frac{a_{ij} - \overline{a}_j}{s_j} \right], \qquad (30)$$

where a_{ij} is the observed level of asset j in municipality i, $\overline{a}_j$ is the mean of asset j across all municipalities, and s_j is the standard deviation of asset j across municipalities. The intuition underlying our PI is that there are a number of household assets and characteristics that explain variation in unobserved poverty outcomes. By understanding of how these assets co-vary (and by how much), we can infer, from the composition of these assets across municipalities, how relative poverty levels vary across municipalities.

We confirm the validity of our PI both internally and externally. The factor scores from the eigenvectors in Table 8 provide evidence that our PI is internally coherent. A positive factor score indicates that the asset variable contributes (adds) to poverty, and *vice versa*. The factor score of each asset variable carries the expected sign.[61] We provide further evidence of the internal validity of our PI in Table 9, in which we list the mean values of each asset within the deciles of the 1992 PI. The trends in asset levels as the PI increases (increasing deciles) are similar to what we would expect to see as wealth decreases, indicating that the PI is likely capturing poverty. As an external (to the PI) measure of our poverty index's validity we measure the correlation between the PI and NBI. Although the PI and NBI capture different aspects of poverty, the two measures should be correlated. We

[61] According to the manner in which the poverty index was constructed, poverty is decreasing in the negative orthant.

find that this is the case, for instance, the correlation between NBI and PI in 2001 is over 0.88.[62]

NBI. The NBI measures the percentage of the population within a municipality with unsatisfied basic human needs. It captures poverty by measuring the goods or services that a household possesses that are associated with well-being and then comparing these municipality-level values to a norm (or ideal). The Bolivian NBI was estimated by the INE in coordination with UDAPE. It comprises a set of factors such as housing, basic services, education and health.[63] The housing component aims to isolate the household environment, in terms of providing protection from the outdoors and other external factors such as animals and insects that transmit diseases. It also includes living spaces inside the household in order to consider social environment, privacy and comfort. The basic services component considers basic sanitation in terms of the need for good quality water for food and hygiene, and the availability of health services that allow privacy, sanitation and hygiene. In addition, NBI considers energy availability and cooking sources. The education portion includes the years of schooling, school attainment and literacy. Finally, the health component relates to the capabilities of people, and good health that allows the proper development within the social environment.

The individual household components are compared to a norm which is used to determine if the household's basic needs are met. The compilation of each equally weighted component allows for the identification of the poverty condition of each

[62] This correlation can be observed spatially in the maps in Figures 8 and 9.

[63] The health component of the Bolivian NBI is not fully comparable between censuses as the questions have changed. The full methodology can be found at `http://www.ine.gob.bo/pdf/Metodologias2004/NBI.doc`

household.[64] Accordingly, a higher measure of the NBI indicates greater poverty within the associated municipality.

Baseline covariate distributions

Previous studies at the global and country level have found significant differences in the biophysical (Andam et al. 2010, Sims 2010, Joppa and Pfaff 2009, Pfaff et al. 2009) and socioeconomic (Andam et al. 2010) characteristics underlying protected and unprotected areas. The differences underscore the non-random nature in which protected areas are allocated. Globally it has been shown that protected areas tend to be located distant from cities (markets) and on agriculturally unsuitable land (high slope, high elevation), so-called “high and far” or “rocks and ice” bias (Joppa and Pfaff 2009). In addition, a similar study in Costa Rica demonstrated that communities affected by protected areas had significantly higher levels of baseline poverty than unaffected communities (Andam et al. 2010).

Table 7 shows that Bolivia shares many of the characteristics associated with protected areas that are observed globally. The geographic characteristics associated with access to markets, infrastructure and urbanization differ significantly between protected and unprotected areas. The average distance to a major city, roadless volume and percent baseline forest cover are greater in protected municipalities, indicating that protected areas tend to be established in more rural areas. In addition, it can be seen that one of the primary indicators of agricultural suitability, slope, is greater (indicating lower suitability) in protected municipalities. However, contrary to global trends, we observe that the average elevation in protected

[64]The methods used to formulate the NBI present some limitations related to the weight of the components that are included in the index. All the factors included have the same weight, in addition, the method require some norms to which indicators are compared. These norms are, to some extent, arbitrary. Also, a household is considered poor if at least one of the NBI components are not satisfied. Moreover, NBI does not consider explicitly the demographic structure of the household and prioritizes the housing indicators. There is one final, practical, limitation: there are 13 municipalities (four of which are considered protected) for which NBI was not calculated in 1992.

municipalities is lower than in unprotected municipalities (see Figure 10). Most interesting is the fact that, according to both poverty measures, baseline poverty is slightly *lower* in protected municipalities. This finding is contrary to findings from Costa Rica (Andam et al. 2010) and common wisdom.

Methods

The underlying differences in covariate values between protected and unprotected municipalities indicates the importance of controlling for such differences in the estimation of the impacts of protected areas. The selection issue that we must address is that protected areas are not established randomly across the landscape. Non-random allocation leads to the observed imbalance across these key covariates, that jointly determine selection into protection and socioeconomic outcomes (see Table 7), which may lead to biased estimates of the impacts of protected areas under naïve comparisons of protected and unprotected municipalities. To reduce the bias associated with our estimates of the socioeconomic impacts of protected areas, we use matching as our primary strategy to control for this confounding imbalance.

Matching

To measure the impact of protected areas on poverty in surrounding municipalities we use matching to estimate the average treatment effect on the treated (ATT). Estimation of the ATT is implied in our research question, "what would poverty outcomes in protected municipalities have been had they not been protected?" Answering such a question implies the estimation of a counterfactual, and because there are municipalities for which it is implausible to suppose the establishment of a protected area, the estimation of an ATT is most appropriate.[65]

[65]Estimation of average treatment effects (ATE), for instance, entails the estimation of an additional counterfactual: outcomes for all unprotected units had they been protected. We argue that it is implausible for many of Bolivia's municipalities to have been protected and, therefore, the estimation of ATT is more appropriate that ATE.

The key to matching as an identification strategy to estimate ATT is the balancing of covariate distributions across treatment arms (protected and unprotected) thus mimicking the identification strategy of a randomized experiment. This covariate balance is achieved in expectation through randomization. Covariate balance is implicit under randomization because each unit of the experimental sample has an equal probability (or more generally, a probability that is known to the experimenter) of being assigned to treatment or control. Therefore, treatment is assigned independent of potential outcomes $Y(1)$ and $Y(0)$ under treatment ($T = 1$) and control ($T = 0$), respectively. In the absence of a treatment, one would expect similar average outcomes from both groups. Similarly, if both groups were to receive (the same) treatment, one would expect similar average outcomes from both groups. In the statistics, epidemiology and social science literature this assumption is termed ignorability of treatment, independence of treatment or unconfoundedness. Stated formally,

$$E[Y(1)|T = 1)] = E[Y(1)|T = 0)] = E[Y(1)] \tag{31}$$

$$E[Y(0)|T = 1)] = E[Y(0)|T = 0)] = E[Y(0)]\,. \tag{32}$$

In words, (2) simply states that average potential outcome for the treatment group under treatment, $E[Y(1)|T = 1)]$, is equal to the average potential outcome of the control group *had they been treated*, $E[Y(1)|T = 0)]$. Similarly, (3) states that the average potential outcome for the treated group *had they not been treated*, $E[Y(0)|T = 1)]$, is equal to the average potential outcome of the control group in the absence of treatment, $E[Y(0)|T = 0)]$. In (2) and (3), the terms $E[Y(1)|T = 0)]$ and $E[Y(0)|T = 1)]$ are termed counterfactual outcomes. The fundamental problem for causal inference (Holland 1986) is the fact that counterfactual outcomes are not observed. However, with treatment assigned at random (and thus independent of

potential outcomes), the average outcome for control units can act as the counterfactual for treatment units, and *vice versa*.

Protected areas in Bolivia were not established randomly. Matching seeks to mimic the identification of randomization by balancing key covariates that jointly determine selection into treatment and outcomes. Balance, conditional on key covariates, leads to conditional ignorability or conditional independence. However, because our primary estimand of interest is the ATT we only need to estimate one counterfactual. Therefore, it is only necessary for us to invoke the conditional independence assumption (CIA) for (2). This more restrictive assumption can be stated formally as the analog to (2),

$$E[Y(0)|T=1,X] = E[Y(0)|T=0,X)] = E[Y(0)|X]\,. \tag{33}$$

Equation (4) states that, conditional on similar covariate distributions across treatment arms, the average outcomes for the matched control units, $E[Y(0)|X,T=0)]$, can be used as the counterfactual for treatment units. In other words, by ensuring that the distributions of key covariates are balanced across treatment and control groups, similar methods to those used in randomized experiments can be used to estimate ATT on matched datasets.[66] By ensuring that units are comparable across treatment and control groups, we make the CIA, which is necessary for causal inference, more defensible (Angrist and Pischke 2009).

Primary estimator

There are many matching metrics from which to choose. Our final specification is determined by the metric that provides the best balance across our covariates of

[66]Similarly, by additionally invoking CIA for equation (3) (i.e., E[Y(1)—T=1,X]=E[Y(1)—T=0,X)]= E[Y(1)—X]), average treatment effects can be measured.

interest.[67] We find that, given our relatively small sample size, genetic matching provides the best balance and, therefore, is most likely to satisfy the CIA. Genetic matching (Sekhon 2007) conducts an algorithmic search across potential weighting matrices in order to optimize the weighting matrix to best satisfy covariate balance. We conduct a series of robustness checks on the estimates stemming from the genetic matching, including adding calipers, calculating the Rosenbaum bounds (Rosenbaum 2002), and various regression based estimators (see the Robustness Section below).

Our primary genetic matching specification uses the single nearest neighbor (in terms of covariate distribution) to each treated unit to act as the counterfactual for each treated unit. We allow for replacement (which generally reduces bias but can increase the variance (Imbens and Wooldridge 2009, Dehejia and Wahba 2002)) during matching, use a post-match regression bias adjustment (Imbens and Wooldridge 2009, Abadie and Imbens 2006, Abadie et al. 2004),[68] and calculate so-called Abadie and Imbens (2006) heteroskedasticity robust standard errors. Our matching specification seeks to find an unprotected municipality that is observably similar to each protected municipality, isolating the only remaining variation between treatment arms to be the establishment of protected areas, thereby allowing the unbiased estimation of ATT.

[67] During the process of selecting a matching metric we tested the balancing properties of many different metrics (e.g., Mahalanobis, propensity score and inverse covariance). Outcomes and ATT estimates were omitted while inspecting balance across different specifications to prevent the estimates from potentially affecting the selection of a metric.

[68] Because our matched samples are relatively balanced, the post-match regression bias adjustment has relatively little effect on the point estimates.

Results

Primary Results

Figure 11 and Table 10 present the results from our primary, and ancillary robustness, analyses (Tables 11 and 12 provide balancing results from our primary matching specifications). In this subsection we focus on the second column of the respective right, and left, hand bar charts and the respective row in Table 10, which present the ATT estimates stemming from the genetic matching algorithm.

For both our PI and NBI poverty indicators we find no evidence to suggest that the establishment of protected areas in Bolivia had deleterious effects on poverty. To the contrary, all of our point estimates indicate that protected areas were likely associated with poverty alleviation. In other words, after controlling for covariates that jointly influence the establishment of protected areas and poverty, we find that there was differentially greater poverty reduction between 1992 and 2001 in municipalities that had at least 10% of their area occupied by a protected area. The point estimates from the primary specification are statistically significant (at any standard level) when the PI is used as the outcome of interest, but is insignificant when NBI is used as the outcome of interest (though the point estimates are concurrent with those of the PI).

An attractive feature of our matching-based estimator is its transparency in terms of allowing for the identification of mean poverty outcomes across treatment arms of the matched sample, which represent the components of the ATT. Table 10 highlights the underlying difference between the naïve and genetic matching estimates, which stem from the formulation of the counterfactual sample. When the counterfactual comprises all unprotected municipalities other than those marginally protected, the counterfactual poverty outcome (the poverty level that would have been observed in protected municipalities, had they not been protected) is -0.451

(85.61) according to the PI (NBI). Under our genetic matching specification we are left with 56 and 53 counterfactual unprotected municipalities for the PI and NBI analyses, respectively (41 and 38 of which are unique in the respective analyses). Because the counterfactual unprotected units under the genetic matching specification are observably similar (at least more so than the naïve counterfactual, see Tables 11 and 12), the counterfactual outcome estimates are more similar to the treated sample (-0.805 for the PI and 84.16 for NBI), resulting in a more modest estimate of the poverty reduction associated with the establishment of protected areas for both poverty metrics.

Our results are concordant with previous studies from Costa Rica (Andam et al. 2010) and Thailand (Andam et al. 2010, Sims 2010): protected areas are associated with poverty reduction. However, our results differ fundamentally from these previous studies. Andam et al. (2010) and Sims (2010) find that naïve estimates of protected area impacts indicate that protected areas exacerbated poverty. When key covariates are controlled for, however, their results reverse. In contrast, our results indicate that a failure to control for key covariates leads to the over estimation of the impacts of protected areas on poverty reduction.

Robustness

We test the robustness of our primary estimates in several ways. First, we test the sensitivity of our matching estimator to unobserved heterogeneity between protected and unprotected units. The purpose is to identify by how much the groups would have to differ (unobservably) in order to nullify our results of statistically significant poverty reduction. Second, we test the robustness of our matching specifications by comparing our primary specification to a number of matching- and regression-based econometric specifications.

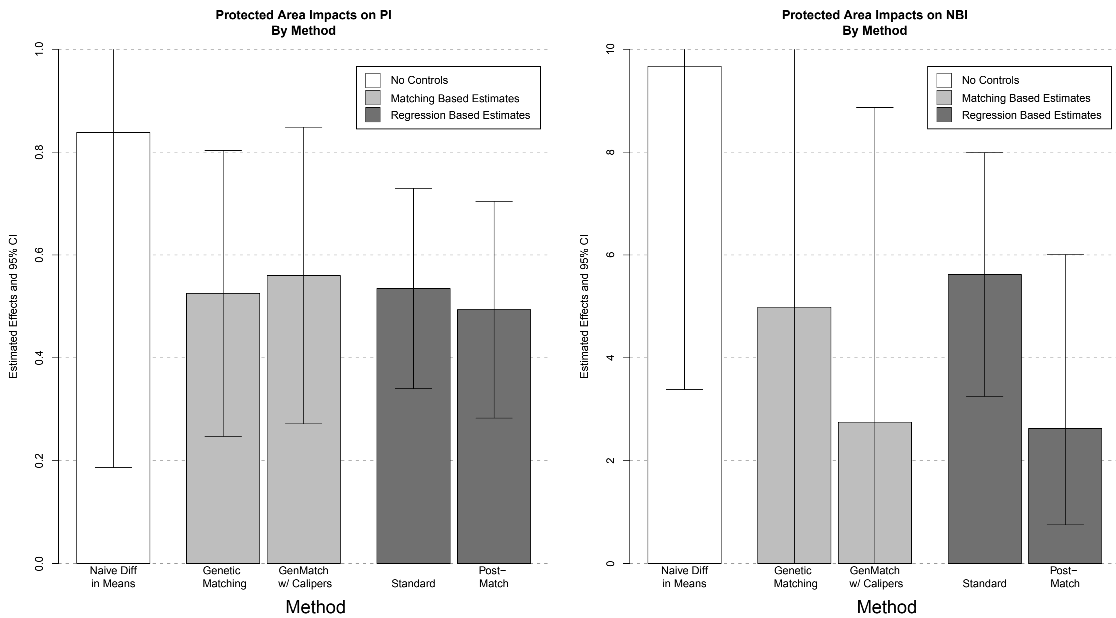

Figure 11: Primary estimates of the impacts of protected areas on poverty in Bolivia. The left hand (right hand) group of bars represent the impact according to poverty index (NBI) across a number of econometric specifications. The results from the primary genetic matching specification described in Methods can be found in the second bar of the respective bar groups. The whiskers represent the 95% confidence interval for the corresponding point estimates underlying each bar.

Internal robustness of matching specification

In any observational study it must be acknowledged that the ability to eliminate bias associated with non-random selection is limited by one's understanding of the underlying selection process (Meyer 1995), moreover, from a practical standpoint, by the pertinent characteristics of selection that one can actually observe and obtain. If the selection process and outcomes are systematically determined only by observable characteristics (for which one controls) then a treatment effect estimate derived from a matching algorithm that provides balance will be unbiased and consistent. However, if there are unobservable characteristics that also contribute to determining selection and outcomes, then treatment effect estimates, even for a well balanced matched sample, may be biased. We believe our data are rich enough to provide sufficient covariates with which to control, therefore mitigating unobserved heterogeneity. However, we test the sensitivity of our ATT estimates to unobserved heterogeneity/bias using Rosenbaum bounds (Rosenbaum 2002).

The Rosenbaum bounds sensitivity analysis measures the amount of unobserved heterogeneity necessary to undermine the statistical results from our matching process. If a great (small) amount of unobserved heterogeneity is necessary to weaken the significance of our results then the results are relatively robust (sensitive).[69] Table 13 indicates the level of unobserved heterogeneity (unaccounted for in our matching process) that would be necessary to nullify our findings of statistically significant poverty reduction according to the PI. Our results are robust (at the 5% level) to unobserved heterogeneity that affects the odds of selection into protection by a factor of 2.3. In other words, these results are highly robust to potential unobserved bias.

[69]See Appendix C for full details on Rosenbaum bounds.

Robustness of matching specification to alternative econometric specifications

To ensure that our results are not sensitive to the choice of econometric specification, we conduct a series of ancillary matching and regression based analyses. The results of these analyses can be found in Figure 11 and Table 10.

Genetic matching with calipers. We use one-to-one nearest neighbor matching based on the genetic matching algorithm in our primary econometric specification. Although we achieve a high level of balance in expectation across treatment arms, there are invariably a number of treated units that do not obtain a well-matched control/counterfactual unit. To ensure that a few poorly matched units are not biasing or driving the results, we impose calipers equal to one standard deviation on our primary matching specification. In other words, we use the identical genetic weighting matrix, however, we remove from the sample any matched pair that differ by more than one standard deviation across covariate values.[70]

Figure 11 and Table 10 show that the results for the PI are relatively robust to the introduction of calipers, i.e., there is only a marginal absolute increase in the point estimate of ATT. Seven protected municipalities are dropped from the analysis, resulting in a bilateral increase in PI outcomes across matched protected and unprotected municipalities (from -1.33 to -1.07 and -0.805 to -0.511, respectively). The variance of the resulting ATT changes little and the point estimate is significant at any conventional level.

Results for NBI are not as robust. Six protected municipalities are dropped from the sample resulting in an absolute reduction in the ATT from -4.99 to -2.47. Table 10 indicates that this change is due to opposite movements in average poverty

[70]The variance is measured according to the scalar value assigned to each unit after taking the product of the covariate values of each unit and the genetic weighting matrix. This scalar, like a propensity score, mitigates the so-called curse of dimensionality associated with multivariate matching.

Table 10: Results from Primary and Ancillary Analyses

	Poverty Index			NBI		
Method	Y(T=1)	Y(T=0)	Treatment	Y(T=1)	Y(T=0)	Treatment
Nave Difference in Means	-1.33 [56]	-0.451 [268]	-0.838 {0.014}	76.7 [53]	85.61 [242]	-8.92 {0.005}
Regression on Raw Data	NA [56]	NA [268]	-0.502 (0.98)	NA [53]	NA [258]	-5.52 (1.17)
Regression Dropping Marginal	NA [56]	NA [252]	-0.535 (0.099)	NA [53]	NA [242]	-5.62 (1.2)
Post-Match Frequency Weighted Regression	NA [56]	NA [41]	-0.494 (0.106)	NA [53]	NA [45]	-2.63 (1.7)
Genetic Matching	-1.33 [56]	-0.805 [56]	-0.525 (0.142)	76.18 [53]	81.16 [53]	-4.99 (3.67)
Genetic Matching, Calipers=1sd	-1.07 [49]	-0.511 [49]	-0.56 (0.147)	79.04 [47]	81.51 [47]	-2.47 (1.55)

[Number of observations]
(Standard errors)
{P-value}

across the protected and unprotected samples. Average NBI in protected municipalities increased from 76.18 to 79.04 while average NBI in unprotected municipalities decreased from 84.16 to 81.51. Although the variance in ATT decreases after calipers are imposed, the resulting ATT estimate remains insignificant at conventional levels.

Regression-based specifications. We run several regression-based econometric specifications to ensure that our results are not driven by the use of a matching-based estimator. The results of these specifications are found in Table 10 (see Appendix C Table 34 for full regression results) and we highlight these specifications in last two columns of each bar group of Figure 11. Although there is slightly greater heterogeneity in the specifications for which NBI is the outcome, the central results from these specifications are that: (1) protected areas are associated with significant poverty reductions according to both the PI and NBI, and; (2) the results do not differ significantly from the primary matching-based estimates.

Table 11: Balance Results for Primary GenMatch Specification- PI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	0.267	0.744	-0.477	0.141	0.436	
1992	Matched	0.267	0.319	-0.052	0.015	0.164	89.12%
% Forest 1991	Unmatched	0.485	0.194	0.291	0.476	0.292	
	Matched	0.485	0.466	0.019	0.029	0.037	93.40%
Distance to	Unmatched	142800	108000	34790	0.173	35950	
Major City	Matched	142800	136600	6151	0.029	32940	82.32%
Average	Unmatched	1825	2713	-888	0.307	884	
Elevation	Matched	1825	1794	30.94	0.011	118.9	96.52%
Average Slope	Unmatched	23.89	18.97	4.922	0.162	5.34	
	Matched	23.89	23.89	0.004	0.000	1.628	99.92%
Roadless	Unmatched	2.98E+14	8.35E+13	2.14E+14	0.194	1.96E+14	
Volume 1992	Matched	2.98E+14	2.53E+14	4.53E+13	0.030	1.27E+14	78.89%

Table 12: Balance Results for Primary GenMatch Specification- NBI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
NBI 1992	Unmatched	87.52	91.05	-3.527	0.078	3.15	
	Matched	87.52	87.56	-0.044	0.001	0.97	98.74%
% Forest 1991	Unmatched	0.464	0.187	0.277	0.453	0.277	
	Matched	0.464	0.459	0.005	0.007	0.03	98.22%
Distance to	Unmatched	136400	108600	27790	0.140	29580	
Major City	Matched	136400	132600	3730	0.017	25160	86.58%
Average	Unmatched	1888	2751	-863.1	0.299	862.8	
Elevation	Matched	1888	1884	4.489	0.002	106.7	99.48%
Average Slope	Unmatched	24.01	19.45	4.555	0.149	5.055	
	Matched	24.01	23.67	0.334	0.010	3.979	92.66%
Roadless	Unmatched	3.06E+14	8.53E+13	2.20E+14	0.195	1.99E+14	
Volume 1992	Matched	3.06E+14	2.58E+14	4.80E+13	0.033	1.06E+14	78.22%

Bar columns 4 and 5 of each bar group in Figure 11 present the results of a standard regression (marginally protected units are dropped from the sample and all covariates are included as controls), and a post-match frequency weighted regression,[71] respectively.

Both the regression-based analyses for which the PI is the outcome of interest (bars 4 and 5 on the left hand side of Figure 11) return estimates that are strikingly similar to one another, and to the primary matching specification. The confidence intervals are slightly tighter than those from the matching based estimators (expected given the efficiency properties of OLS and the fact that the standard specification contains a larger sample of unprotected units). There is more heterogeneity in the impact estimates in which NBI is the outcome of interest. However, one of the more interesting results from the NBI regression specifications is that, although the estimated impacts of protected areas on poverty do not differ significantly from the primary matching specification, the impacts of protected areas on NBI are all significantly different from zero.

Results Summary

The central finding in our results is that it does not appear as if protected areas have had any negative effect on poverty in surrounding communities. Rather, our results indicate that protected areas were likely associated with reductions in poverty. Though these results are relatively robust across specifications, a couple of questions linger.

[71] The post-match frequency weight regression is conducted on the resulting matched sample from the primary genetic matching specification. To correct for potential overstatement in the precision of coefficient estimates (due to repeat unprotected matched observations) we drop all duplicate observations from the unprotected sample and then weight each unprotected unit by the number of times it was used as a match for a protected unit, to ensure unbiased coefficient estimates. This so-called “double robust” estimation strategy is promoted by Ho et al. (2007) because the second stage regression helps to eliminate any residual differences across treatment arms that remain after matching.

PI or NBI?

Although the point estimates in all the primary and ancillary analyses indicate poverty reductions associated with protected areas, the results from analyses in which the PI is the outcome of interest are more consistent, and remain statistically significant. This begs the question, in which poverty measure we should place more stock? From a policy standpoint we argue that it matters little. There is no evidence that protected areas exacerbated poverty, by any measure. From a technical standpoint, however, we believe that the PI is more appropriate in our study.

The NBI has some technical and practical shortcomings, mentioned previously. First, the NBI weights all socioeconomic components equally, unlike the PI in which weights are determined empirically via the PCA. Second, the NBI measure the percentage of households that lie below a norm which is (somewhat) arbitrarily established. The PI, on the other hand, is based on deviations from municipality-level means (in either direction). Finally, and from a practical standpoint, analyses in which the NBI is used are limited by the fact that we do not have baseline NBI measurement for 14 municipalities (four of which are considered protected).

Why are the matching-based and regression-based estimates so similar?

The goals of matching and regression in causal analysis are the same: achieve plausible conditional independence of the treatment. However, the two methods go about doing so in different manners. Regression isolates the causal effect of a treatment by establishing a functional relationship between treatment, covariates and outcome, then isolates the causal effect of treatment by partialing out the effects of the covariates of interest. Instead of controlling for the differences across treatment arms via the imposition of a functional form, matching uses a weighting

scheme[72] to create balance in expectation across the covariates, thereby "controlling" for their influence.

Typically in an analysis in which the ATT is the estimand of interest we would expect greater differences in matching- and regression-based results. The reason relates to the idea of propensity for treatment. In a study like ours, where there are control units that would never plausibly be treated, the overlap in propensity for treatment is generally sparse (or completely lacking) at the tails of the distributions (of propensity scores and the underlying covariates) between protected and unprotected units. This scenario can lead to bias in regression results because the coefficient estimates can be heavily (and inappropriately) influenced by outlying control units (e.g., Ho et al. (2007)).

The reason that our results across matching- and regression-based estimates are so similar can be seen in Figure 22. We plot the distributions of propensity scores across treatment arms for pre- and post-matched samples. It can be seen that there is a high degree of overlap across the range of propensity scores even prior to matching (left panel of Figure 22). Therefore, it makes sense that the matching- and regression-based results are comparable because, even within the full sample, the regression results are not plagued by "out of sample" predictions.

Discussion

Protected areas have played an increasingly important role in the global conservation of biodiversity and ecosystem services over the past several decades. However, there is little empirical evidence of the environmental impacts of protected areas and even less evidence of their socioeconomic impacts on surrounding communities. Given the high degree of overlap between remaining global

[72] For example, with one-to-one matching without replacement, control units are excluded from the sample by receiving an effective weight of 0 whereas remaining units receive a weight of 1.

Table 13: Rosenbaum Upper Bound on P-Value at Given Levels of Γ for Primary Matching Analysis- PI

Γ	Upper Bound P-value
1	0
1.1	0.0001
1.2	0.0003
1.3	0.0007
1.4	0.0015
1.5	0.0029
1.6	0.005
1.7	0.0081
1.8	0.0124
1.9	0.018
2	0.0253
2.1	0.0342
2.2	0.0448
2.3	0.0572

biodiversity and poverty, it is of paramount importance to understand how the establishment of protected areas impacts poverty.

We use rich biophysical and socioeconomic data, and a myriad of econometric specifications to estimate the impact of protected areas established in Bolivia between 1992 and 2000 on poverty between 1992 and 2001. Contrary to the concerns of poverty advocates, that the land-use restrictions associated with protected areas impart economic hardship on surrounding communities, our results do not indicate that protected municipalities were adversely affected by the establishment of protected areas. In fact, we find evidence that municipalities with at least 10% of their areas occupied by a protected area had differentially greater poverty reduction than those unaffected by protected areas. We employ two separate measures of poverty in our analyses and find that the point estimates of poverty reduction are robust across our econometric specification.

Although our overarching results that Bolivia's protected areas were associated with poverty reduction are similar to previous studies from Costa Rica (Andam et al. 2010) and Thailand (Andam et al. 2010, Sims 2010), our underlying results

are subtly, but significantly, different. In those studies the authors found that controlling for key observable covariates lead to fundamentally antithetical results compared to naïve estimates. Conversely, our results indicate that naïve estimates lead to an over estimation of the poverty reducing impacts of protected areas.

The implications of our results are twofold. First, our results add to a growing body of literature on the impacts of protected ares on poverty. More importantly our findings add support to this literature that environmental conservation policies do not necessarily run in opposition to development goals. On the contrary, our results indicate that environmental goals can be complementary to social poverty goals. Second, our results underscore the importance of country-level analyses of the socioeconomic impacts of protected areas. Protected areas in Bolivia exhibit many of the characteristics observed globally (i.e., located relatively distant from major cities, roads and on steeper slopes), however, some of the key drivers of poverty differ in important ways from global and previous country-level observations. Importantly, we key find differences from previous country-level and global studies which indicate that evidence from single country or global studies are likely not generalizable across countries.

The fact that our results exhibit subtle differences to previous results implies that the external validity of our, and other studies of this ilk, is likely limited. Indeed, we believe that comprehensive understanding of the socioeconomic impacts for protected areas requires that the scientific body of knowledge be built on a country-by-country basis.

Further studies in Bolivia and elsewhere should strive to identify and quantify the mechanisms through which protected areas affect poverty (e.g., Hanauer (2011), Robalino and Villalobos-Fiatt (2010)). Although studies such as ours are important for building an understanding of the global impacts of protected areas, only by understanding *how* protected areas affect poverty (especially in terms of alleviating

poverty) can social policies be designed to enhance (mitigate) the positive (negative) impacts of protected areas. In addition, because the theme of protection in Bolivia has been toward integrated management and recognition of indigenous populations, future studies should account for differences in protected area management practices and baseline populations.

Appendix A

Appendix to Chapter I

A.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.

A.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).

Appendix B

Appendix to Chapter II

B.1 Data

B.1.1 Costa Rica

For full details on data, see Andam et al. (2008, 2010). Digital layers of protected areas (source: National System of Conservation Area Office, Ministry of Environment and Energy, 2006) were provided by the Earth Observation Systems Laboratory, University of Alberta. Other GIS layers are land use capacity (source: Ministry of Agriculture) and roads digitized from hard copy maps for 1969 (source: Instituto Geográfico Nacional, Ministerio Obras Publicas y Transporte). Summary statistics of the data are presented in Table 1.

For the deforestation analyses, digital forest cover layers are created from either aerial photographs (baseline, 1960) or Landsat Thematic Mapper satellite images.[73] The units of analysis are land parcel pixels that measure three hectares, the minimum mappable area. We use Andam et al. (2008) data set, which comprises 20,000 randomly selected land parcels that were forested (80% or more canopy cover) in 1960. Forested parcels in a given year receive a value of 1; deforested parcels receive a value of 0. The outcome of interest is change in forest cover between 1960 and 1997. Given all sample parcels were forested in 1960, the outcome measure equals 0 if the parcel is still forested in 1997 and 1 if it is deforested.

[73] Earth Observation Systems Laboratory, University of Alberta, Edmonton, AB.

Spatial layers of protected status (IUCN categories Ia, I, II, IV and VI are used in the analyses) and other geographic characteristics are used to create a set of covariates for each land parcel (Table 14). For various reasons (e.g., cloud cover), 4,737 land parcels are dropped prior to analysis, leaving 15,283 land parcels, of which, 2,809 were protected prior to 1980. We remove parcels that were protected after 1980 (2,183), leaving 10,291 unprotected land parcels from which matches can be drawn.

For the poverty analyses, data come from the population and housing censuses conducted by the Instituto Nacional de Estadistica y Censos (INEC) in 1973 and 2000. Digitized GIS census segment boundaries for 1973 and 2000 were provided by the Cartography Department at INEC. The unit of analysis is the census tract (segmento). In 1973 Costa Rica contained 4,694 census tracts with an average size of 8.82km (range: 0.00466-836 km). The 1973 census is used as the baseline year and all census data are geocoded to their respective census tracts. Between 1973 and 2000 there was a great deal of segmentation of census tracts, with few of the segmented tracts being proper subsets (or sharing major borders) with the original 1973 census tracts. Through the method of areal interpolation ((Andam et al. 2010, Reibel 2007)); see below), the 2000 census data are aggregated to fit to the 1973 census tract boundaries so that the data are spatially and temporally comparable.

The poverty measure (poverty index) builds on recent efforts to develop a census-based poverty index for Costa Rica (Cavatassi et al. 2004), which uses principal components analysis to formulate a temporally comparable index based on variables believed to influence poverty.

B.1.2 Thailand

For full details on data see Andam et al. (2010) and Sims (2010). Digital layers of protected area boundaries are from the IUCN World Database on Protected Areas

(accessed 3/2007; IUCN categories I and II were used in the analyses). Other GIS data and the source layers from which they are derived are slope and elevation (NIMA's Digital Terrain Elevation Data- USGS Global GIS database, 1999); distance to major cities (ESRI World Cities, 2000); distance to roads in 1962 (digitized East Asia Road Map, U.S. Map Service 1964, data from 1962); distance to rail lines, distance to major rivers, proximity to watershed boundaries, distance to mineral deposits, distance to Thai border, and ecoregions (USGS Global GIS database, 1999), average monthly temperature and rainfall (Marc Souris, IRD).

The deforestation analysis is based on two classified layers from 1973 and 2000. The 1973 data are based on Landsat MSS images interpreted by the Tropical Rain Forest Information Center (Michigan State University) and the 2000 data on Landsat TM images interpreted by the Thai Royal Forestry Department (courtesy of Marc Souris). The units of analysis are points which are spaced so as to represent the centroid of a three hectare parcel. The data set is created in a similar manner to the Costa Rica deforestation data set and comprises 20,000 randomly selected points which were forested in 1973. Forested points in a given year receive a value of 1; deforested points receive a value of 0. The outcome of interest is change in forest cover between 1973 and 2000. Given all sample points were forested in 1973, the outcome measure equals 0 if the point is still forested in 2000 and 1 if it is deforested. Spatial layers of protected status and other geographic characteristics are used to create a set of covariates corresponding to each sample point (Table 14).

For the poverty analysis the unit of analysis is a subdistrict (tambon). In descending order of size, Thailand has administrative units of "province," "district," "subdistrict," and "village." The sample consists of subdistricts in the North and Northeast regions, where the majority of protected forest areas are located. We exclude subdistricts that are less than 10 km away from a major city (population

100,000; all of these cities had been established by the 1960's). The average size of a subdistrict in the sample is 74 sq km; the average population is 5043.

The poverty measure for Thailand (poverty headcount ratio) is the share of the population with consumption below the poverty line. This outcome is derived from a poverty mapping analysis by Healy and Jitsuchon (2007), applying the poverty mapping methodology developed by Elbers et al. (2003).

B.2 Preprocessing

We preprocess the data (Ho et al. 2007, Imbens and Wooldridge 2009) using matching techniques prior to performing any of the LOESS or PLM analyses. Our primary motivation for matching is not the estimation of an overall average treatment effect on the treated. With the exception of an analysis of Thailand deforestation at the scale used in our study, these impacts have already been estimated (Andam et al. 2008, 2010). We use matching to preprocess the data so that we can estimate conditional average treatment effects on the treated. To ensure that our analyses are as comparable as possible to the studies from which we draw (Andam et al. 2008, 2010), we use the same matching methods to create the same matched data sets as those studies. These methods were chosen in these studies because they generated the best covariate balance.

The key to matching as an identification strategy to estimate average treatment effects on the treated is the balancing of key covariate distributions across treatment arms (protected and unprotected). This covariate balance is achieved in expectation through randomization. Covariate balance is implicit under randomization because each unit of the experimental sample has an equal probability (or more generally, a probability that is known to the experimenter) of being assigned to treatment or control. Therefore, treatment is assigned independent of potential outcomes $Y(1)$ and $Y(0)$ under treatment ($T = 1$) and control ($T = 0$), respectively. In the absence

of a treatment, one would expect similar average outcomes from both groups. Similarly, if both groups were to receive (the same) treatment, one would expect similar average outcomes from both groups. In the statistics, epidemiology and social science literature this assumption is termed ignorability of treatment, independence of treatment or unconfoundedness. Stated formally,

$$E[Y(1)|T=1)] = E[Y(1)|T=0)] = E[Y(1)] \tag{34}$$

$$E[Y(0)|T=1)] = E[Y(0)|T=0)] = E[Y(0)]\,. \tag{35}$$

In words, (34) simply states that average potential outcome for the treatment group under treatment, $E[Y(1)|T=1)]$, is equal to the average potential outcome of the control group *had they been treated*, $E[Y(1)|T=0)]$. Similarly, (35) states that the average potential outcome for the treated group *had they not been treated*, $E[Y(0)|T=1)]$, is equal to the average potential outcome of the control group in the absence of treatment, $E[Y(0)|T=0)]$. In (34) and (35), the terms $E[Y(1)|T=0)]$ and $E[Y(0)|T=1)]$ are termed counterfactual outcomes. The fundamental problem for causal inference (Holland 1986) is the fact that counterfactual outcomes are not observed. However, with treatment assigned at random (and thus independent of potential outcomes), the average outcome for control units can act as the counterfactual for treatment units, and *vice versa*.

Protected areas in Costa Rica and Thailand were not established randomly. Matching seeks to mimic the identification of randomization by balancing key covariates that jointly determine selection into treatment and outcomes. Balance, conditional on key covariates, leads to conditional ignorability or conditional independence. These more restrictive assumptions can be stated formally as the

analogs to (34) and (35),

$$E[Y(1)|T = 1, X] = E[Y(1)|T = 0, X)] = E[Y(1)|X] \tag{36}$$

$$E[Y(0)|T = 1, X] = E[Y(0)|T = 0, X)] = E[Y(0)|X]. \tag{37}$$

Equations (36) and (37) state that, conditional on similar covariate distributions across treatment arms, the average outcomes for the matched control units, $E[Y(0)|X, T = 0)]$, can be used as the counterfactual for treatment units, and *vice versa*. In other words, by ensuring that the distributions of key covariates are balanced across treatment and control groups, similar methods to those used in randomized experiments can be used to estimate average treatment effects on matched datasets. We present (34)-(37) for completeness; however, we focus on the estimation of conditional average treatment effects on the treated, for which only (35) and (37) are necessary.

By ensuring that units are comparable across treatment and control groups, we make the conditional independence assumption (CIA), which is necessary for causal inference, more defensible (Angrist and Pischke 2009). We extend the CIA by assuming that if average treatment effect on the treated estimates are unbiased, conditional on balance across key covariates, comparisons of subgroups within these balanced sets are also unbiased. This allows for causal inference to be drawn from the LOESS and PLM analyses.

As mentioned in Chapter 2, matching can only account for heterogeneity in observable covariates. If the selection process and outcomes are systematically determined only by observable characteristics (for which one controls) then a treatment effect estimate derived from a matching algorithm that provides balance will be unbiased and consistent. However, if there are unobservable characteristics that also contribute to determining selection and outcomes, then treatment effect

estimates, even for a well balanced matched sample, may be biased. There is no way to formally test the conditional independence assumption, however Andam et al. (2008, 2010) test the robustness of their estimates (which are derived from the same matched sets used in our study) to unobserved heterogeneity.

B.2.1 Matched Datasets

For the Costa Rica data, we use nearest neighbor Mahalanobis covariate matching with replacement to preprocess the socioeconomic and deforestation data. We use the same algorithm and covariates (Table 1) as Andam et al. (2008, 2010), and thus our resulting matched datasets are nearly identical to those used in their analyses.[74] The resulting socioeconomic matched set comprises 249 protected (prior to 1980) and unprotected census tracts. The resulting deforestation matched set comprises 2,809 protected (prior to 1980) and unprotected land parcels. See Table 1 for description and summary statistics of the covariates used in each Costa Rica matching specification.

For the Thailand socioeconomic data we use propensity score matching with exact matching on district in order to control for baseline fixed effects associated with poverty. This is the same specification and matched set used in (Andam et al. 2010) which comprises 197 protected (prior to 1985) and unprotected subdistricts. For the Thailand deforestation data we use Mahalanobis covariate matching, with exact matching on district, to create a dataset that is similar to the Costa Rica deforestation analysis (see Tables 16 and 17 for estimates of ATT and balancing results). The resulting matched set comprises 2,808 protected (prior to 1985) and

[74] The socioeconomic matched set is identical to the final data set in Andam et al. (2010). The deforestation matched sets would be exact, but we use a slightly updated protected areas database resulting in slightly more protected observations. The average treatment effect on the treated estimates, however, are not different between the two datasets. We present the balancing results in Table 15.

unprotected land parcels. See Table 14 for description and summary statistics of the covariates used in each Thailand matching specification.

B.2.2 Thailand Deforestation Analysis

To ensure methodological comparability across countries, we perform a similar deforestation analysis to that of (Andam et al. 2008) for Thailand. Our primary interest was to create a dataset, comparable to the Costa Rica deforestation dataset, with which to perform the heterogeneity analyses. As a point of departure, however, we perform sample average treatment effect on the treated calculations similar to those done in (Andam et al. 2008). There are two benefits to this approach. First it offers a comparison to the original Costa Rica deforestation analysis (Andam et al. 2008). Second, it provides an average treatment effect on the treated (ATT) estimate to which we can contrast our heterogeneity analyses.

In creating our deforestation dataset for Thailand we follow the methodology of (Andam et al. 2008); see their SI Text), all geoprocessing is done in ArcGIS 9.x. We begin by selecting 20,000 random points, spaced so as to represent 3 ha land parcels, from the areas of Thailand that were forested in 1973, our baseline year. Using spatial overlays, we create indicators for parcels that were protected by 1985 (2,808) and parcels that were protected after 1985 (3,423). The analysis is designed to estimate the impact of protected areas that were established prior to 1985 on deforestation outcomes between 1973 and 2000. Therefore, we remove from the pool of potential controls, any parcel that was protected *after* 1985. As a result, our potential pool of controls comprises 13,609 parcels that were never protected prior to 2000.[75] We run a series of overlay analyses on the remaining parcels to assign a value for each of the covariates listed in upper panel of Table 14.

[75] Due to incongruence in spatial layers, 160 parcels are dropped prior to analysis.

Using these data, we implement regression bias-adjusted nearest neighbor Mahalanobis matching with replacement (Imbens and Wooldridge 2009, Abadie et al. 2004) to estimate the ATT. Point estimates and balancing results can be found in Tables 15and 16, respectively.[76] Similar to (Andam et al. 2008), we find that the naive difference in means overestimates the amount of avoided deforestation attributable to the establishment of protected areas. As noted in Chapter 2, this is a finding that is consistent with the general observation that protected areas tend to be placed on land that is less desirable for agriculture, and therefore less likely to be deforested in the absence of protection. The resulting matched dataset is used for the Thailand deforestation heterogeneity analyses described in Chapter 2.

B.3 LOESS

Three LOESS estimators (Cleveland 1979, Cleveland and Devlin 1988) are performed for each of the covariates in the heterogeneous response to protection analyses: (1) on the protected units only; (2) on the imputed counterfactual control units only, and; (3) on the difference between protected and counterfactual unprotected units, the Average Treatment Effect on the Treated (ATT).

In LOESS the data of interest are the doubles (Y_i, X_i) representing the outcome and covariate values for observation $i \in \{1, 2, ..., N\}$, where N is the number of observations in the dataset. The data are first ordered according to X such that $X_1 \leq ... \leq X_N$. Beginning with the first observation ($i^* = 1$) in this ordered set, fitted values $\left(\widehat{Y}\right)$are predicted via a local quadratic regression

$$\widehat{Y}_{i \in s} = \widehat{\beta}_0 + \widehat{\beta}_1 X_{i \in s} + \widehat{\beta}_2 X_{i \in s}^2, \tag{38}$$

[76]In addition to the covariates listed in Tables 14 and 16, matching is required to be performed within districts (i.e., exact matching on district ID) to control for regional heterogeneity.

where the vector $\widehat{\beta}$ is estimated from

$$Y_{i \in s} = \beta_0 + \beta_1 X_{i \in s} + \beta_2 X_{i \in s}^2 + \epsilon_i, \tag{39}$$

and only observations that lie within span (s) are used. The total number of observations used for each imputation is therefore $j = sN$. Moving stepwise through the ordered data set, N local regressions are estimated.

For each of these local regressions all of the j observations are assigned a weight (w_d) using the tricubic function

$$w_d = \begin{cases} \left(1 - |d_i|^3\right)^3 & \text{for } 0 \leq |d_i| < 1 \\ 0 & \text{otherwise} \end{cases}, \tag{40}$$

where d_i is a cardinal distance ratio

$$d_i = \frac{|X_{i^*} - X_i|}{\max\left(|X_{i^*} - X_i|\right)}. \tag{41}$$

Here X_{i^*} represents the covariate value of the observation for which we are imputing $\widehat{Y}$. The weight w_d reduces the influence of observations according to their disparity in covariate value as compared to the observation being evaluated. The LOESS estimation moves stepwise repeating (38)-(41) for each (ith) observation, "re-centering" the span s to include an equal number j observations about the ith observation. The result of these N local regressions is N local fit values $\left(\widehat{Y}_i\right)$ and their corresponding standard errors of the fit which can be used to form confidence intervals about each fit value. This standard LOESS process is run on the protected units for each analysis (dash-dot line in Figures 14 - 19.

We extend the LOESS methodology in order to offer comparability to the studies from which we draw (Andam et al. 2008, 2010, Ferraro and Hanauer 2011) by

including local bias-adjusted imputation of counterfactual (unprotected) outcomes. This type of method is used in the matching literature (e.g., Imbens and Wooldridge (2009), Abadie et al. (2004)) to impute counterfactual values by plugging the values of treated unit covariates into the coefficients estimated from a regression of control unit covariates on control unit outcomes. The purpose of this imputation is to reduce post-match bias, in finite samples, due to remaining covariate imbalance. This process is like asking the question, "what would the outcomes of protected units have been in the absence of protection had their covariates influenced their outcomes in the same manner as the units that were not protected?"

Our methodology requires us to modify the LOESS procedure. In order to impute counterfactual outcome values for each treated unit, both protected and unprotected units must be used as inputs for the LOESS. Prior to the ith local estimation outlined in equations (38)-(41), a counterfactual value for each protected unit outcome in the span (s) is imputed according to

$$\widetilde{Y}_{i\in s} = Y_{i\in s:T=0} + \widehat{\mu}_0\left(X_{i\in s:T=1}\right) - \widehat{\mu}_0\left(X_{i\in s:T=0}\right), \tag{42}$$

where T is an indicator of treatment (0 and 1 indicating the unit is unprotected or protected, respectively) and $\widehat{\mu}_0(\cdot)$ represents the predicted values obtained from combining the coefficients from a control group regression, of outcome on covariates, with the respective treated or control covariates (see Tables 1 and 14 for a list of the covariates).[77] In addition to estimating a LOESS curve based on these counterfactual outcomes (dotted line in Figures 14 - 19), the counterfactual value $\widetilde{Y}_{i\in s}$ from (42) of the observation being evaluated (i^*) is stored in a vector for use in evaluating a LOESS for ATT.[78]

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

[78]Imputations within the LOESS were programmed in R v2.10.1. Code is available from authors upon request.

The LOESS curve for ATT is estimated using the difference between actual protected unit outcomes (Y_i) and their respective counterfactual outcomes $\left(\widetilde{Y}_i\right)$ from (42),

$$\left(Y_{i\in s} - \widetilde{Y}_{i\in s}\right) = \beta_0 + \beta_1 X_{i\in s} + \beta_2 X_{i\in s}^2 + \epsilon_i, \tag{43}$$

where the corresponding fits are estimated in a similar manner to (36). The standard error of the fit is used to form the confidence band (red/green shaded area) about the ATT LOESS curve (solid line in Figures 2 and 14 - 19).

The span for any LOESS estimator must be chosen so as to balance the bias/variance tradeoff. A relatively small span includes fewer data points and is considered to be more localized and therefore less biased. However, there will be greater variation, *ceteris paribus*, within a small span. Conversely, a relatively large span uses more data and produces smoother curves (less variation) that are considered to be more biased. For each of the LOESS estimators that we implement, we set the span (s) equal to 0.75. We choose this span for all analyses because: (1) after experimenting with many specifications we felt that it captured the important underlying variability with relatively little noise; and (2) we wanted to remain consistent across analyses.

B.4 PLM

B.4.1 Model

For all (moderating) covariates introduced in the Study Design Section we use a two-stage semiparametric partial differencing linear model (Yatchew 1997, 1998). The PLM is advantageous in that it allows us to control, linearly, for a vector of covariates that influence the outcome of interest and then map the outcome as a nonparametric function of the covariate of interest.

The data used in the PLM are the triples (Y_i, X_i, Z_i) where Y is the scalar outcome of interest, X is the scalar covariate for which the nonparametric function will be estimated and Z is a vector of covariates for which we wish to control in our estimation. Our first-stage equation is thus

$$Y_i = Z_i\beta + f(X_i) + \epsilon_i, \tag{44}$$

where β is a vector of coefficients and $f(\cdot)$ is an unknown real function. Our intention is to estimate $f(\cdot)$ net of the effects of Z. In order to achieve the final goal of removing the influence of Z on Y we must first remove the influence of X on Y. In the first stage we begin by ordering the data according to X such that $X_1 \leq ... \leq X_N$ where $i \in \{1, ..., N\}$. Yatchew (1997, 1998) shows that the influence of X on Y can be removed be taking the (first) difference (in (44)) according to X

$$\begin{aligned} Y_i - Y_{i-1} = (Z_i - Z_{i-1})\beta + (f(X_i) - f(X_{i-1})) + \\ +\epsilon_i - \epsilon_{i-1}, \quad i = 2, ..., N. \end{aligned} \tag{45}$$

Under the assumption that $\partial Y/\partial X$ is bounded by a constant, $(f(X_i) - f(X_{i-1}))$ goes to zero as N increases. Intuitively this assumption implies that, when the data are ordered according to X, the marginal influence of X on Y is zero, so that term can be dropped from the equation. OLS can then be run on (45) to return an estimate of $\widehat{\beta}_{diff}$. Yatchew (1997, 1998) shows that because $\widehat{\beta}_{diff}$ converges sufficiently quickly to β, $Z_i\widehat{\beta}_{diff}$ can be subtracted from both sides of (44) to obtain

$$Y_i - Z_i\widehat{\beta}_{diff} = Z_i\left(\beta - \widehat{\beta}_{diff}\right) + f(X_i) + \epsilon_i \tag{46}$$

$$\cong f(X_i) + \epsilon_i. \tag{47}$$

Denoting $Y_i - Z_i\widehat{\beta}_{diff} = Y_i - \widehat{Y}_{i,diff} = \widetilde{Y}$, the combination of the LHS of (46) and RHS of (46) is equivalent to

$$\widetilde{Y}_i = f(X_i) + \epsilon_i. \tag{48}$$

We are now able to estimate $f(\cdot)$, which is the nonparametric relationship between X and Y, net of the effects of Z. We do so for treatment, control and ATT estimates using the same LOESS estimator described above. Lokshin (2006) suggests using LOESS in the second stage and wrote a Stata ado file which performs the estimate. We wrote a similar function for R. The code is available from the authors upon request.

Yatchew (1997, 1998) noted that although $\widehat{\beta}_{diff}$ is an unbiased estimate of β, due to the differencing, $\widehat{\beta}_{diff}$ is relatively inefficient. However, he provides analytical higher order differencing weights that can be applied to a high order difference generalization of (45) to greatly improve the efficiency of estimates. We incorporate these weights into our estimation using the 10th order difference (the highest order for which weights are provided).[79] See (Yatchew 1997) for detailed description of the efficiency issues and a table of the analytical weights.

B.4.2 Empirical Specifications

For each of our PLM analyses we include in Z covariates that we believe affect the outcome of interest. This means that we control for the covariates used in each matching specification and the complementary moderating covariates. There are some notable exceptions, however, in which we exclude or add covariates as controls. For each of the analyses in which distance to major city is the moderating covariate of interest, we exclude distance to road from the controls due to high correlation (multicollinearity). For each of the Thailand socioeconomic analyses, we add

[79] The PLM estimates were programmed in R v. 2.11.1. The code is available from the authors upon request.

province level fixed effects to the vector of covariates. For a detailed account of the controls used in each specification see Table 18. Complete first stage results are available from the authors upon request.

B.4.3 Use of PLM and LOESS

We use LOESS to estimate the relationship between baseline poverty and the outcomes of interest in Costa Rica (Figure 2(a)) because we are interested in what actually happened to the poor over time rather than simply the effect of being poor. To identify the potential for protected areas to act as a mechanism for poverty traps, we do not want to partial out any of the variables that are correlated with being poor. We simply want to observe how areas with differing levels of baseline poverty fared over time.

We view the other covariates (slope, distance to city and percent agricultural workers) as moderating variables through which protection affects outcomes. For this reason we are interested in identifying the specific effect of these covariates, net of other influences, on our outcomes. Thus we use PLM. In addition, the use of PLM to isolate the specific effects of variables allows us to overlay these effects on the suitability maps with fewer concerns of confounding effects

B.5 Suitability Mapping

B.5.1 Motivation

The illustrative suitability maps presented in Chapter 2 characterize the suitability of end-period forested land for protection, based on past observed relationships between covariates and the environmental and socioeconomic outcomes. We characterize suitability along these two outcome dimensions because, while the targeting of protected areas is likely to be based on expected environmental outcomes, the opportunity costs of protection are socioeconomic in nature.

Therefore, it would be beneficial to a planner to understand the expected joint outcomes of the establishment of protected areas.

We choose to formulate our suitability maps based on slope and distance to major city for two reasons. First, these are measurements that are globally available and have been used in past studies of protected areas. Second, these covariates capture the notion of deforestation pressure (see main text) and are therefore likely to be considered in the establishment of protected areas.

B.5.2 Formulation

To map expected suitability for protection in Costa Rica and Thailand we begin by rasterizing the end-period forest cover shapefiles so that each raster cell is 3 ha in size. We then create a distance to city and slope raster based on these end-period forest cover rasters for each analysis.[80] The values of the rasters' cells are populated with measurements of distance to major city and slope, respectively.

The results from the PLM heterogeneity analyses act as the basis for our designation of expected suitability. The PLM results are appropriate for the creation of these maps because they map the continuous nonparametric *effect* of the covariates on the outcome of interest, net the effect of other influencing covariates. To allow for aggregation of suitability across covariates, we rescale the estimated covariate effects on avoided deforestation and poverty to fall within a range of 1 to 10.[81] For example, the maximum estimated effect of slope on avoided deforestation in Costa Rica is 0.139 at a slope of 14%, so it is rescaled to 10. Conversely, the minimum estimated effect is 0.00087 at a slope of 50%, so it is rescaled to 0. Similarly, all estimated effect between the `min` and `max` are rescaled and rounded. The rescaled values are then assigned to the distance to city and slope rasters for

[80] This leaves us with four initial rasters for each country: a distance to major city and slope raster for the deforestation analysis, and a distance to major city and slope raster for the poverty analysis.

[81] Mathematica has a Rescale command which we rewrote for `R`.

each analysis.[82] For example, all of the cells (of the slope raster) with slope values of 14 in the Costa Rica deforestation analysis are assigned a suitability score of 10. Comparable value assignments are made for each covariate in each analysis.

As a result of these assignments each parcel (in each country) has two rescaled environmental suitability scores and two rescaled socioeconomic suitability scores (one based on distance to city and the other based on slope). We use these values to calculate the average suitability (separately for environmental and socioeconomic outcomes) scores for each land parcel. Figure 12 and 13 show the aggregated environmental and socioeconomic suitability on separate maps for Costa Rica and Thailand, respectively. The final compound suitability maps (Figures 3 and 4 in Chapter 2) are created by overlaying the aggregate environmental and socioeconomic suitability maps.

On the final suitability maps, we highlight two types of land parcels: those with expected 'win-win' outcomes (yellow), and those with expected poverty exacerbation (black). A parcel is designated as 'win-win' if its average environmental and socioeconomic suitability scores are jointly greater than or equal to 6 (this corresponds to the top five deciles). Conversely, if the underlying covariate value of a parcel is associated with negative socioeconomic impacts then the parcels is designated as unsuitable for protection due to potential poverty exacerbation from protection. For instance, due to the relationship between agricultural suitability and slope, flat parcels in Costa Rica and Thailand are designated as unsuitable for protection.

[82]In the rescaling of the socioeconomic effects of the covariates, only positive expected outcomes are rescaled between 0 and 10. Any covariate value that is associated with socioeconomic effects deemed unsuitable for protection (see below).

B.5.3 Note on Thailand Results

In the final Thailand suitability map there are distinct concentric circles of predicted 'win-win' outcomes. It can be seen from the underlying suitability maps (Figure 13) and PLM results (Figure 2(d&e)) that these expected outcomes are driven by the nonparametric relationship between the outcomes of interest and distance to a major city. Figure 2(e) indicates that the greatest poverty reduction is expected between approximately 50km and 90km. Expected avoided deforestation is also positive along this range. The range 55-75km, where both expected outcomes are relatively high, is where a majority of the 'win-win' areas lie.

While distance to major city drives the concentric circles observed in Figures 4and 13, it is but a one facet in the determination of the joint suitability. In order for a parcel to be designated as 'win-win' there must be congruence in expected outcomes across distance to city and slope. Much of the land that lies within the 50-75km range is also relatively steeply sloped. Close examination of Figures 4 and 13 show that this is not the case throughout. In fact, there are many parcels within this range that are not designated as expected 'win-win' due to the underlying low slope.

B.6 Ancillary Analyses

B.6.1 Quantile Regression

In the Results Section of Chapter 2, we use the LOESS estimates to assert that the establishment of protected areas has not acted as a mechanism for poverty traps in Costa Rica. Our assertion stems from the fact that, in the mapping of the LOESS, there is a general trend of greater poverty alleviation in areas with higher baseline poverty. To corroborate these results from the nonparametric LOESS estimator, we use a parametric quantile regression (see Koenker and Hallock (2001) for a nice

overview). Quantile regressions estimate covariate effects at defined quantiles of the outcome. In our case, we use deciles of the poverty index in 2000. We are interested in the response to protection according to baseline poverty. To interpret the results of a quantile regression as a treatment effect on the distribution of outcomes, we must invoke a rank preservation assumption. This assumption implies that the poverty rank among census tracts remains stable over time. Given that the correlation coefficient between baseline and outcome poverty index is nearly 0.7, this assumption seems plausible.

We run a quantile regression (using deciles) of 2000 poverty index on an intercept and indicator of protection using the same matched set as is used in the LOESS analysis (described above). We do not include any additional controls in the regression because: (1) the LOESS estimator is (essentially) a univariate regression method, and thus our intention is to use similar specifications to that analysis; and (2) the quantile regression is run using the preprocessed matched set which is designed to be balanced across key confounding covariates. Figure 20 presents the results of the quantile regression in which the solid line represents the point estimates at each decile with the corresponding pointwise 95% confidence band in green. The point estimates can be interpreted as the effect on poverty of "moving" from unprotected to protected at each level of poverty. The results display a similar trend to that seen in the LOESS results (Figure 2(a) of the main text): namely that protection has had greater poverty alleviating effects on the poorer census tracts.

B.6.2 Agricultural Workers

In Chapter 2, we use slope as a proxy for agricultural suitability. Slope has been used in a similar manner in previous studies (Ferraro and Hanauer 2011) as well as a proxy for other deforestation pressures (e.g., logging access; (19)). To support the conjecture that the slope analysis is indeed highlighting the impact of opportunity

costs from agriculture, we run a PLM analysis to study the heterogeneity of protection's impact conditional on baseline percentage of the workforce employed in agriculture in Costa Rica (where we have data on this measure). An opportunity cost argument would predict that avoided deforestation would be higher in areas with a high percentage of the workforce in agriculture and poverty impacts would be lower in these same areas. We observe this relationship in Figure 15 (bottom panel).

B.6.3 Standard Errors

All of our analyses are preceded by matching to improve balance across protected and unprotected units. Because the matching is performed with replacement there are repeated control observations in the final matched samples. The concern with repeat control observations is that precision of the standard error estimates in post-match analyses (e.g., regression) may be overstated. In response to this concern, we first note that our results are driven by the relationships presented in the ATT estimates, rather than the precision of these estimates. For example, we are more interested in the overall relationship between avoided deforestation and slope than knowing whether or not avoided deforestation was significantly different from zero at 45 percent slope.

Second, we note that the standard errors of the fit presented in the main text are not likely to be understated. The final estimate in each of our analyses (both LOESS and PLM) is designed to be interpreted in a manner similar to a post-matching, bias-adjusted difference in means. This design allows us to compare our results to the studies from which we draw. Thus we are performing the final stage LOESS using the independent variable of interest and the individual ATT, which is simply the difference between actual outcome and imputed counterfactual outcome for each protected unit (see LOESS section above). Therefore the degrees of freedom in the estimation of the standard error of the fit is based only on the number of

observations in the protected sample, rather the entire sample of protected and unprotected units (as would be the case in a typical regression context). The fact that unprotected units do not add to the degrees of freedom serves to mitigate the effect of repeated observations, which lie only in the unprotected units.

Third, to offer the reader more confidence that the standard errors used in Figure 2 are not substantially understated, we calculate standard errors via bootstrapping. The 95% pointwise confidence band is determined by the 2.5 and 97.5 percentile bootstrapped outcome at each point of interest along the range of the independent variable. In each analysis, the final stage LOESS estimate is bootstrapped 1000 times.[83] The bootstrapped standard errors are overlaid on the standard errors of the fit in Figures 17 - 19 in which it can be seen that the two standard error estimates coincide closely. One of the key insights that can be taken from Figures 17 - 19 is that our main results are robust to alternative methods of estimating the standard errors.

[83]The bootstrapping function was written in `R` v. 2.11.1. Code is available from authors upon request.

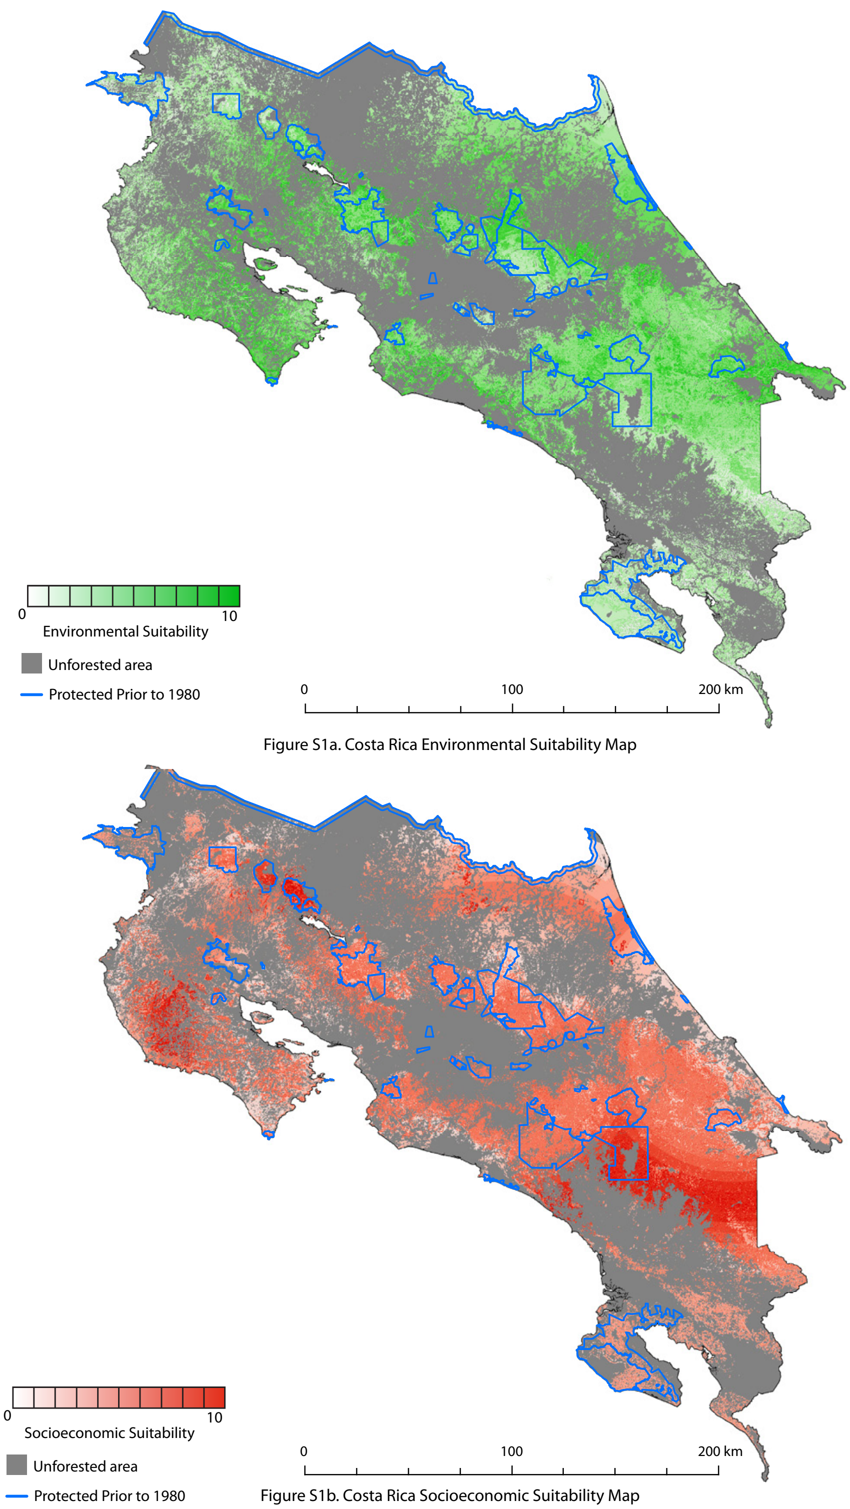

Figure 12: Costa Rica protected area suitability maps by environmental and socioeconomic suitability.

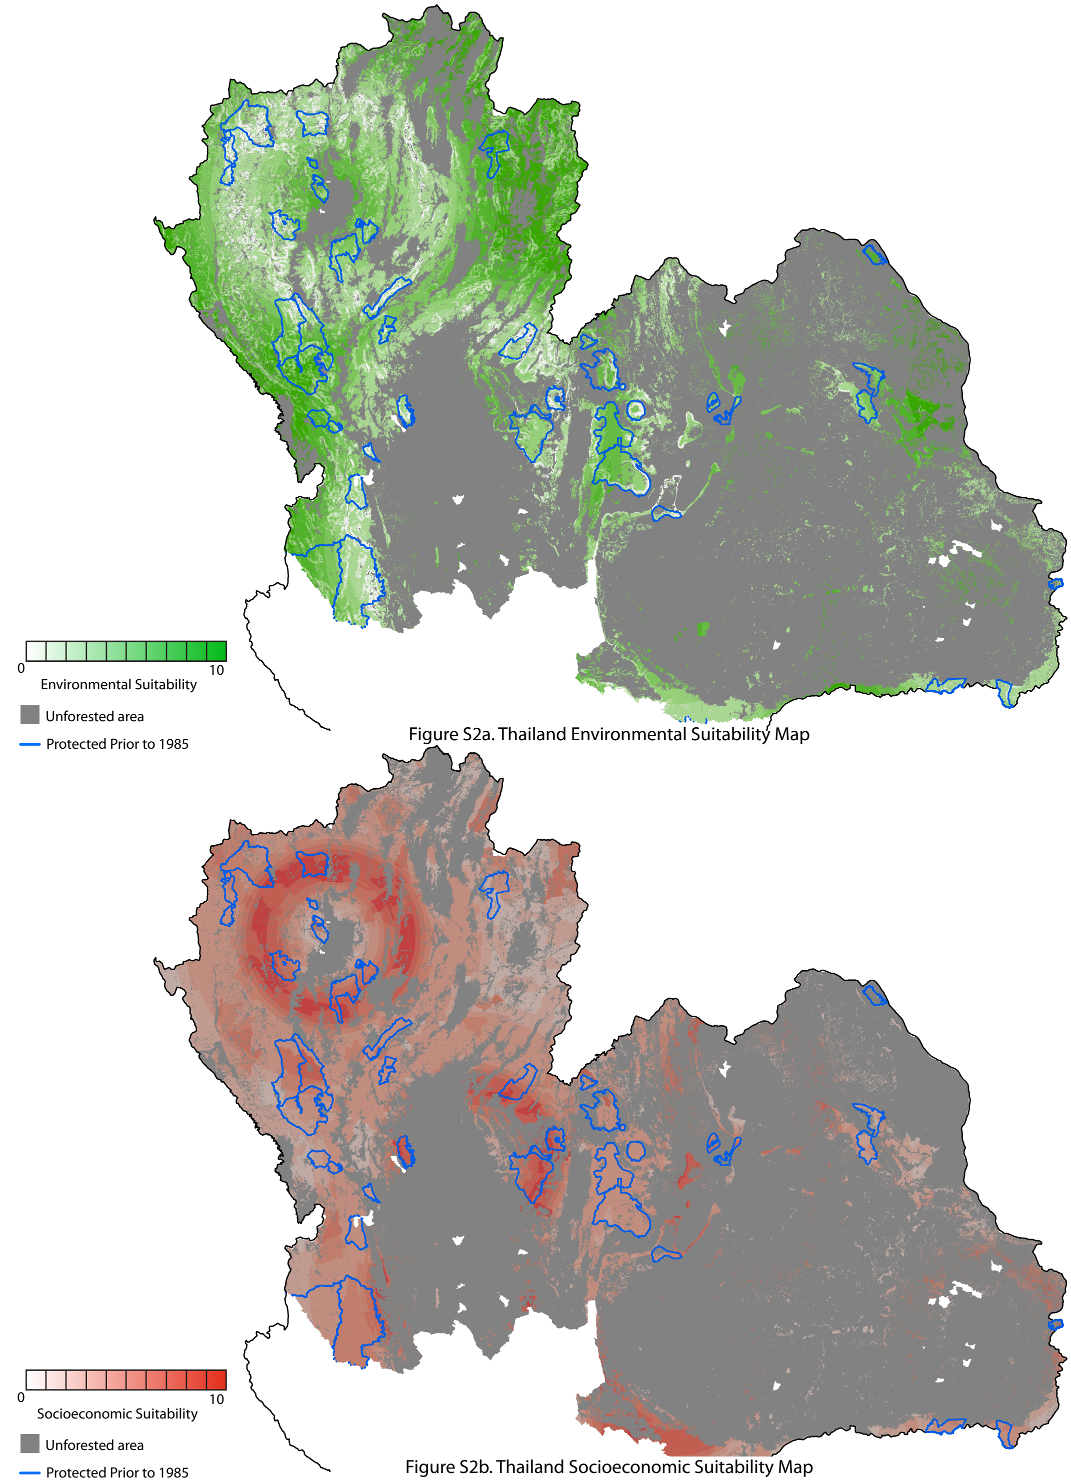

Figure 13: Thailand protected area suitability maps by environmental and socioeconomic suitability.

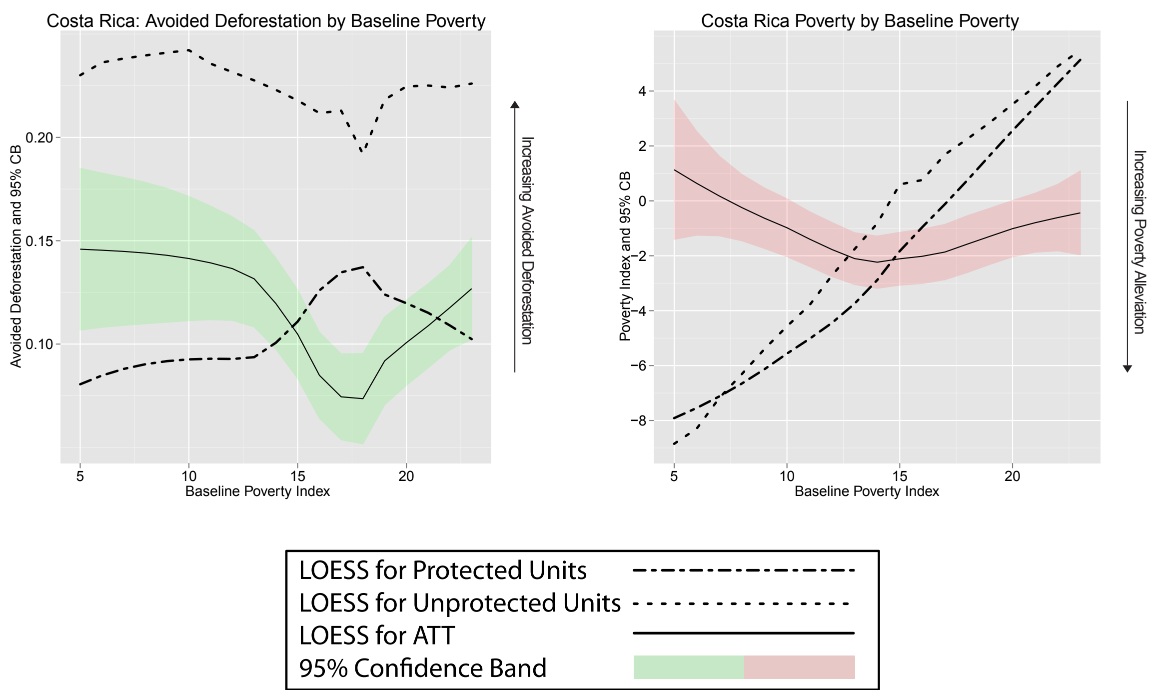

Figure 14: Costa Rica: Full LOESS results.

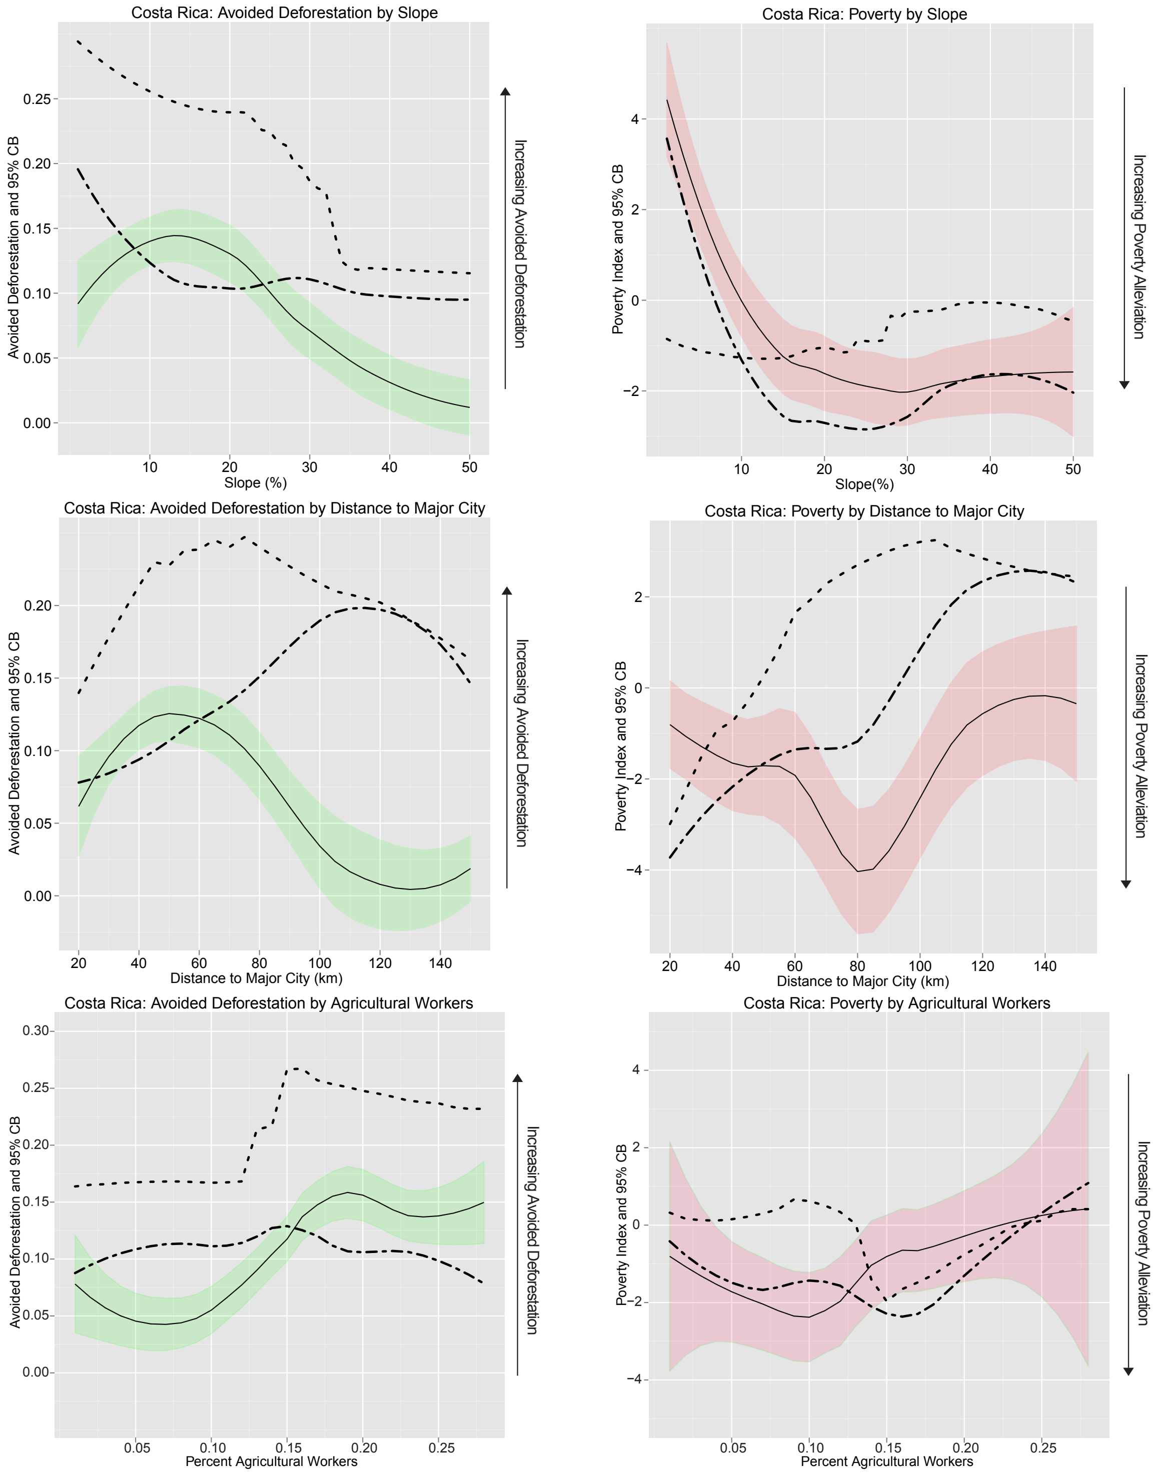

Figure 15: Costa Rica: full heterogeneous response to protection results.

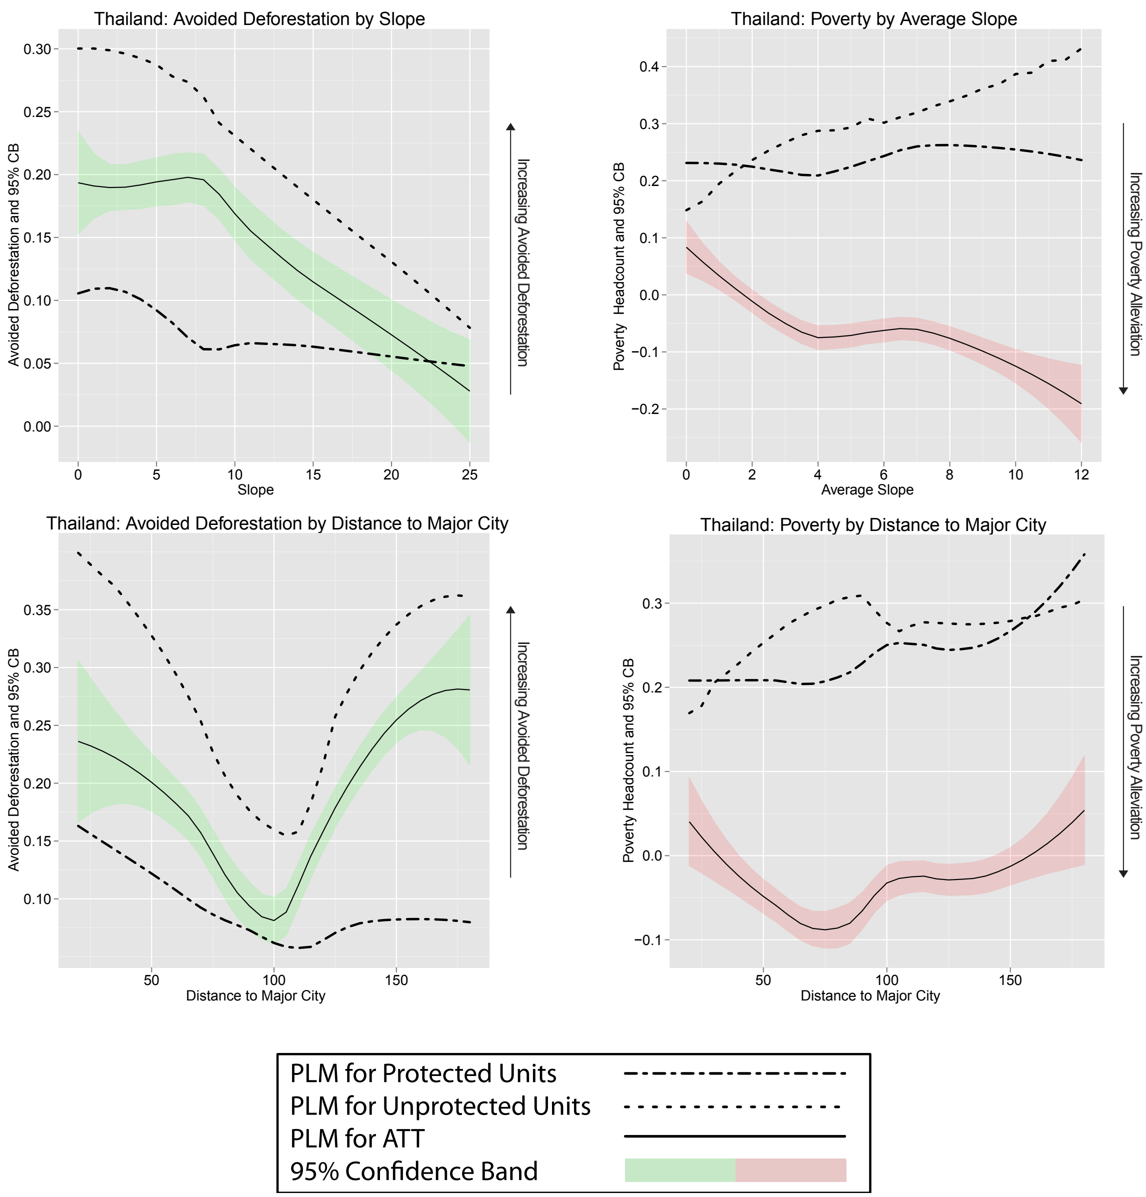

Figure 16: Thailand: full heterogeneous response to protection results.

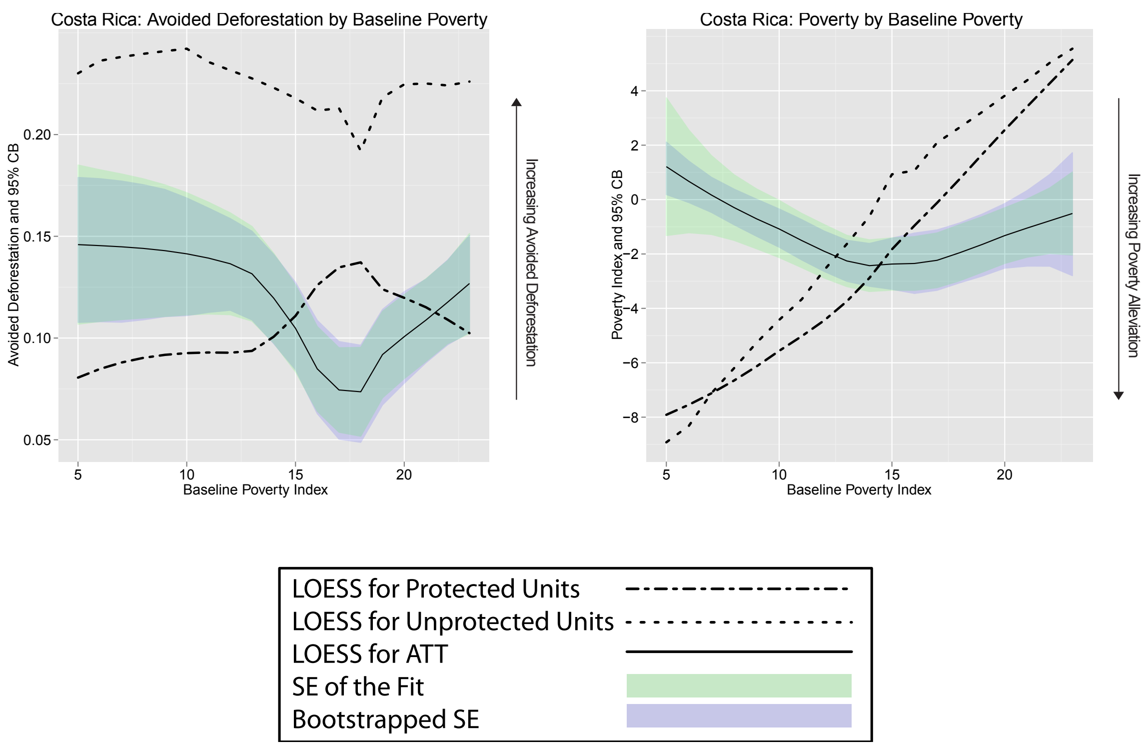

Figure 17: Costa Rica: Comparison of bootstrapped standard errors to standard errors of the fit.

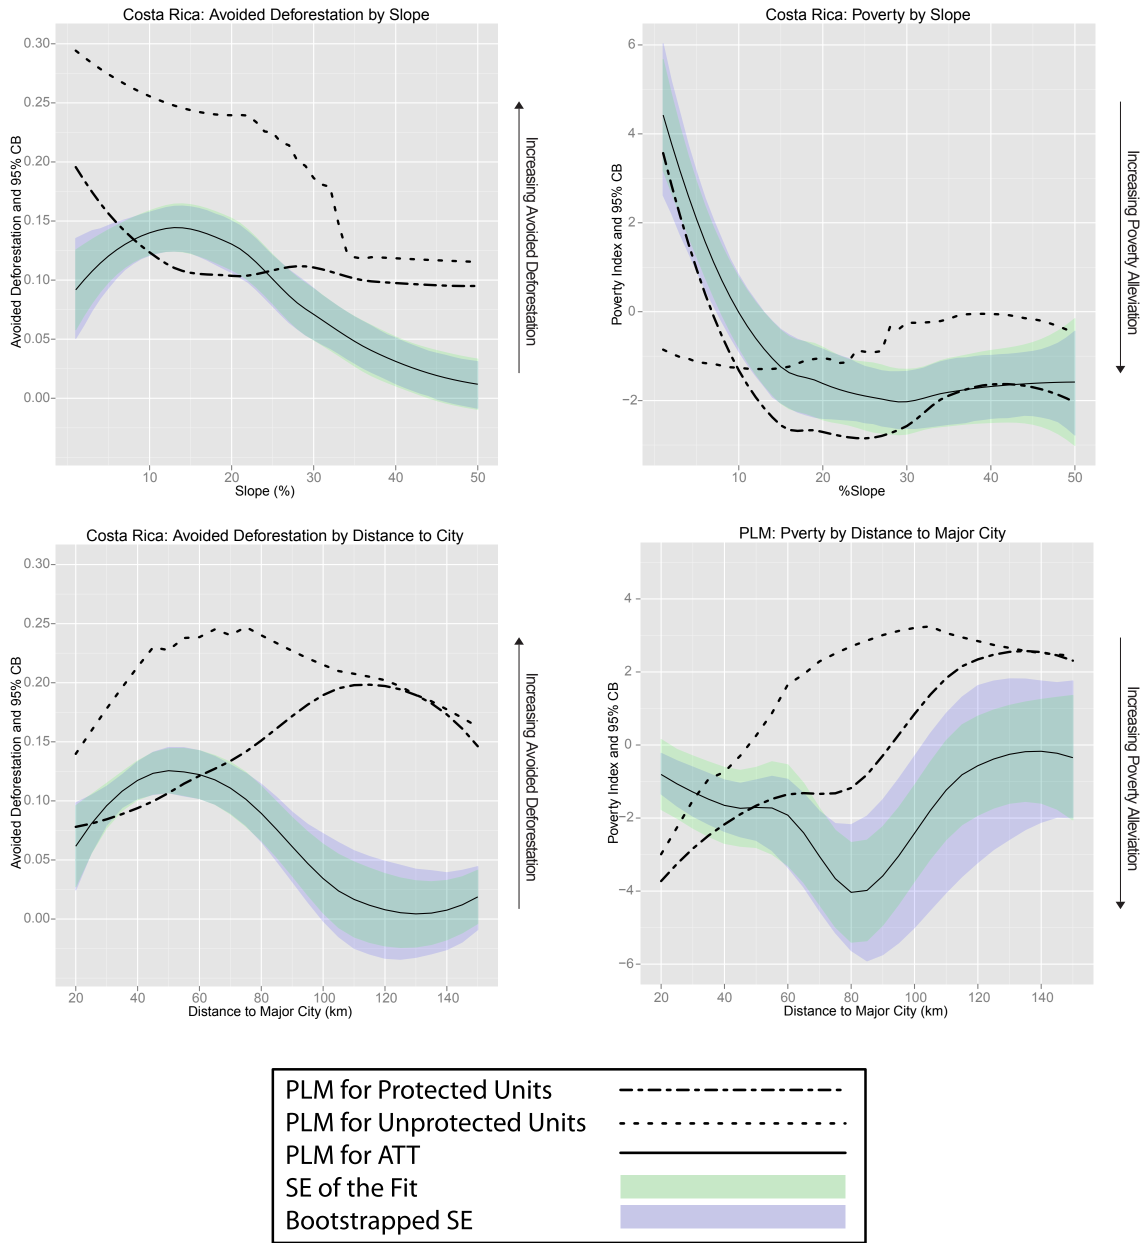

Figure 18: Costa Rica: Comparison of bootstrapped standard errors to standard errors of the fit.

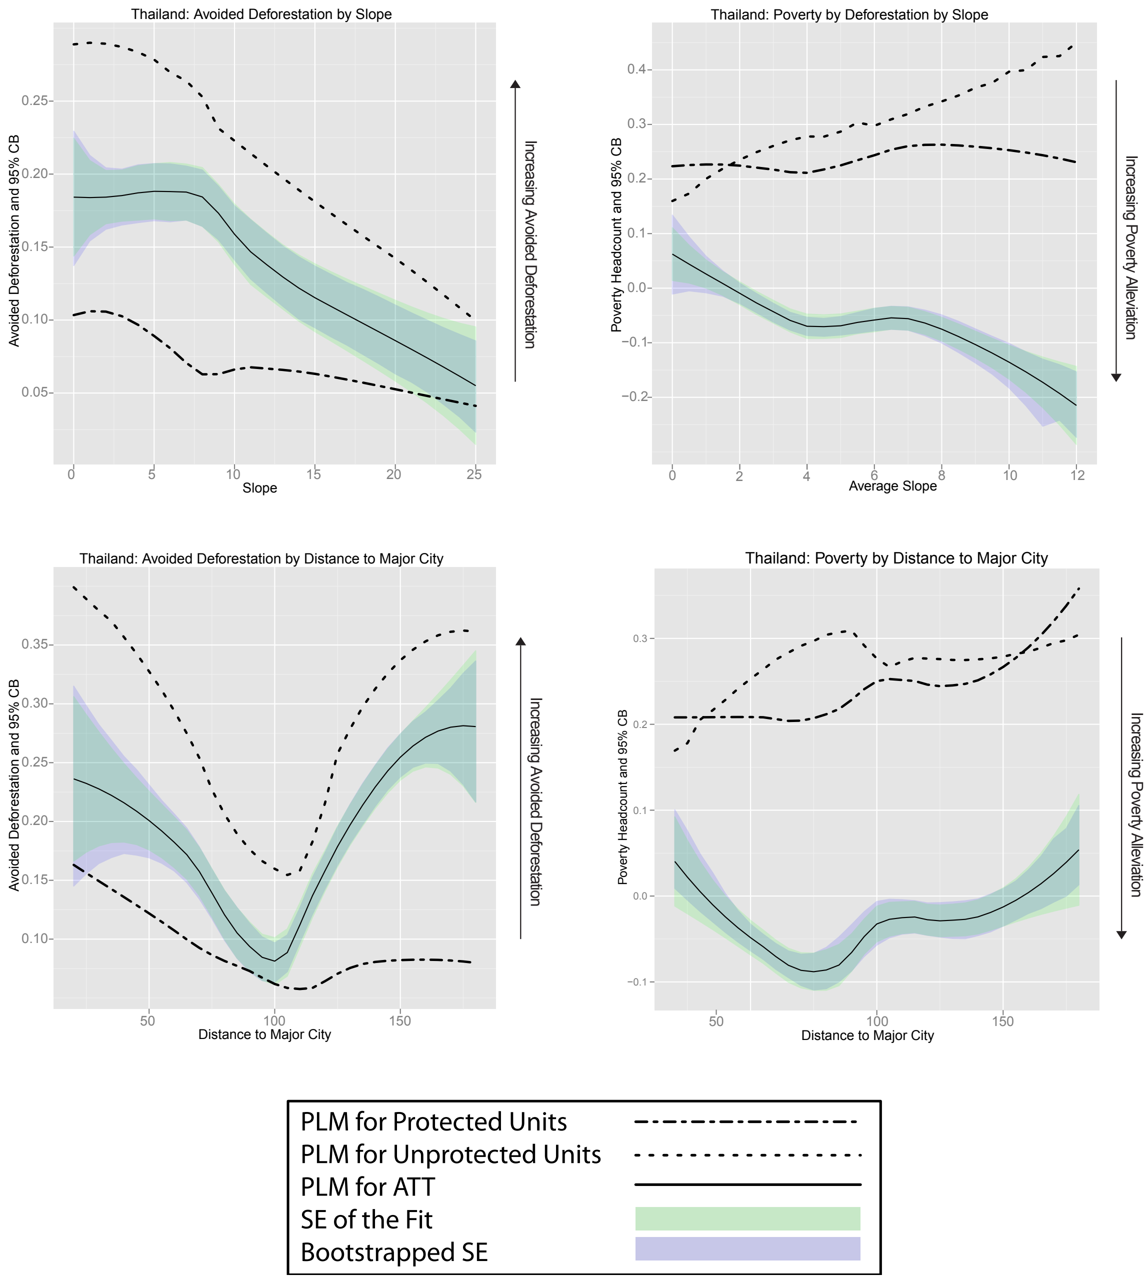

Figure 19: Thailand: Comparison of bootstrapped standard errors to standard errors of the fit.

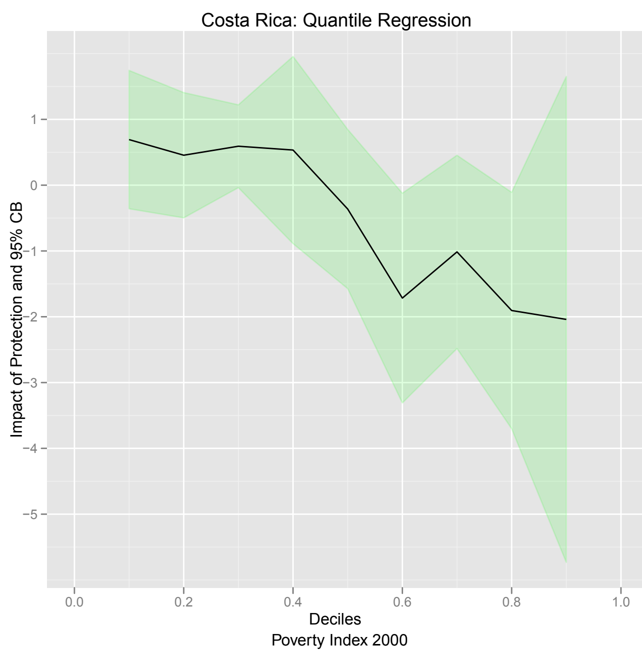

Figure 20: Costa Rica: Quantile regression estimating impact of protection according to deciles of 2000 poverty index.

Table 14: Thailand- Summary statistics and description of covariates used as controls to form counterfactual samples.

Variable	Description	Mean	Median	Standard Deviation	Range
Deforestation Covariates					
Slope	Slope of parcel (degrees)	5.905	5	5.48	0-43
Distance to Major rivers	Distance (km) to major river (flow accumulation greater than 5000)	30.598	27.62	19.552	0.004-109.3
Elevation	Elevation (m) of parcel	555.535	497	316.942	0-2183
Distance to Forest Edge	Distance (km) to the edge of the forest in 1973	2.747	1.884	2.775	0.0001-19.58
Distance to Road	Distance (km) to nearest road in 1962	21.08	116.4	17.682	0.00076-93.8
Distance to Major City	Distance (km) to nearest major city (pop greater than 100,000)	113.573	113.5	42.621	7.26-254.3
Socioeconomic Covariates					
Average Slope	Average slope of subdistrict (degrees)	1.018	0.0504	2.042	0-14.33
Maximum Slope	Maximum slope of subdistrict (degrees)	4.05	0.9882	6.99	0-46.99
Distance to Major River	Distance (km) to major river (flow accumulation greater than 5000)	21.61	0	16.61	0.01-97.82
Forest Cover 1973	Percent of subdistrict covered by forest, 1973	0.194	0.00423	0.315	0-1
Distance to Major City	Distance (km) to nearest major city (pop greater than 100,000)	85.59	81.03	44.51	10.05-222.6
Distance to Major Road	Distance (km) to major road in 1962	5.26	7.615	6.22	0.002-76.16
Distance to Any Road	Distance (km) to minor road in 1962	10.42	3.448	0.002	88.08
Distance to Thai Border	Distance (km) to Thailand border	91.62	91.33	52.36	0.062-218.9
Near Watershed	Within 1 km of major watershed boundary	0.461	0	0.499	0-1
Distance to Rail Line	Distance (km) to rail line	55.05	42.95	45.76	0.015-222.1
Dist. to Mineral Deposit	Distance (km) to nearest mineral deposit	119.46	102.7	84.73	1.371-376.4
Temperature	Average temperature (C) for subdistrict	25.37	25.89	1.448	18.07-27.85
Rainfall	Average monthly rainfall (mm)	1064	1021	225.3	375.8-2308

Table 15: Costa Rica - Covariate balance for baseline avoided deforestation analysis.

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
High Land Use	Unmatched	0.008	0.205	-0.197	0.307	0.197	
Capacity	Matched	0.008	0.008	0.000	0.000	0.000	100.0%
Medium-High	Unmatched	0.029	0.198	-0.170	0.259	0.170	
Land Use Capacity	Matched	0.029	0.029	0.000	0.000	0.000	100.0%
Median-Low	Unmatched	0.080	0.507	-0.427	0.563	0.427	
Land Use Capacity	Matched	0.080	0.080	0.000	0.000	0.000	100.0%
Distance to	Unmatched	2.857	2.045	0.812	0.162	0.886	
Forest Edge	Matched	2.857	2.713	0.143	0.031	0.148	82.3%
Distance to	Unmatched	17.354	15.336	2.017	0.078	2.099	
Road	Matched	17.354	16.709	0.645	0.026	0.975	68.0%
Distance to	Unmatched	76.980	80.515	-3.535	0.037	15.894	
Major City	Matched	76.980	77.912	-0.933	0.008	2.295	73.6%

Table 16: Thailand- Baseline avoided deforestation analysis.

	Difference in Means	Mahalanobis Matching†
Avoided Deforestation	-0.2595***	-0.14738***
$(Y_{T=1} - Y_{T=0})$	{0.0062}	(0.0175)
N Protected	2,808	2,808
N Available Controls	NA	13,609

*** Indicates significance at the 1% level
† ATT is post-match difference in means using regression bias-adjustment to control for bias in finite samples
(Abadie-Imbens heteroskedasticity robust standard errors)
{Standard errors}

Table 17: Thailand - Covariate balance for baseline avoided deforestation analysis.

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Distance to	Unmatched	109.67	114.06	-4.395	0.040	8.498	
Major City	Matched	109.67	110.42	-0.756	0.008	4.958	82.8%
Distance	Unmatched	31.16	17.16	13.999	0.401	14.001	
to Road	Matched	31.16	29.42	1.741	0.039	2.054	87.6%
Distance to	Unmatched	3.62	2.26	1.359	0.242	1.358	
Forest Edge	Matched	3.62	3.30	0.316	0.051	0.317	76.7%
Slope	Unmatched	7.96	4.99	2.970	0.254	2.972	
	Matched	7.96	7.81	0.152	0.012	0.437	94.9%
Distance to	Unmatched	35.93	27.24	8.691	0.217	8.689	
Major River	Matched	35.93	34.89	1.043	0.024	2.464	88.0%
Elevation	Unmatched	697.13	486.61	210.519	0.307	210.448	
	Matched	697.13	635.13	62.004	0.093	62.061	70.6%

Table 18: Exclusions and Inclusions, with respect to the baseline set of controls, in PLM analyses.

Covariate	Exclusions	Inclusions	Justification
		Costa Rica	
Slope	Land Use Cap.†‡	NA	LUC is a function of slope
D. MCity	Dist.e to Road†‡	NA	Colinearity with distance to city
% AgWorkers	NA	NA	NA
		Thailand	
Slope	NA	Province Fixed Effects‡	Control for baseline poverty
D. MCity	Dist. to Road†‡	–	Colinearity with distance to city
	Dist. to Railroad‡	–	Colinearity with distance to city
	–	Province Fixed Effects‡	Control for baseline poverty

Baseline set of controls for each analysis include all matching covariates (Table 1) and other mediating covariates

† Indicates inclussion/exclussion from the deforestation analysis

‡ Indicates inclussion/exclussion from the socioeconomic analysis

Appendix C

Appendix to Chapter IV

C.1 Threshold Analyses

In our primary specifications we designate a municipality as protected if at least 10% of its area is occupied by a protected area. A 10% protection threshold is in line with the goals set forth at the 4^{th} World Congress on National Parks and Protected areas (Andam et al. 2010)[84] and previous studies (Andam et al. 2010, Ferraro and Hanauer 2011, Ferraro et al. 2011). However, it could be argued that this threshold is somewhat arbitrary. We, therefore, test the robustness of our results to changes in this threshold assignment.

Table 19 provides comparisons of ATT for the primary genetic matching specification at the 5%, 10% (primary specification in main analysis), 20%, 30% and 50% protection thresholds for both PI and NBI (Tables 20-23 provide full results for all protection thresholds). Table 19 shows that as the protection threshold increases, i.e., as we increase the protected area land coverage required for a municipality to be considered protected, the number of protected municipalities drops (as expected). In the final 50% threshold specification, only 18 (17) treated units remain in the PI (NBI) analyses.

In the PI analyses we find that the ATT remains relatively stable, and statistically significant, across the range of protection thresholds. The mean

[84] As mentioned in the main text, one of the goals set forth by the 4^{th} World Congress on National Parks and Protected areas was to protect 10% of the earth's ecosystems by the year 2000.

outcomes for treated and control groups jumps up (in absolute terms) at the 50% threshold providing indication that there is somewhat lower poverty within municipalities with greater area protected (note that their matched counterparts have relatively low average poverty levels as well). We see a similar phenomenon in the NBI analyses in Table 19. The lowest average poverty outcomes (for protected and matched unprotected units) are observed at the highest levels of protection. The ATT according to NBI is increasing monotonically (in absolute terms) with percent protection, however none of the estimates are significant at the 5% level (the 50% threshold specification is significant at the 10% level).

The results in Tables 19 - 23 provide evidence that our primary results are not driven by the choice of threshold. Rather, our results are robust and consistent across protection threshold specifications.

C.2 Placebo Analysis

In our main and ancillary analyses we show that the estimated poverty alleviation associated with the establishment of protected areas is robust to a number of econometric specifications and ancillary analyses. However, there is always the concern that the difference in outcomes between protected and unprotected municipalities stems from our inability to select a control group that closely enough resembles the protected group.[85] To address this potential confounding issue we perform a placebo analysis.

The goal is to see if our covariates of interest perform well in the construction of a counterfactual for municipalities that are observably similar (on average) to protected municipalities, but were never protected. In other words, to see if the gap in poverty outcomes between protected and unprotected groups was due to something other than protection. If protection was the only remaining source of

[85] This concern is unlikely true given the high degree of balance across treatment arms in Tables 11 and 12.

variation across treatment arms (to which the treatment effect can be attributed) in our main analyses, then we should observe *no* difference in expected outcomes between the placebo group and its matched controls (i.e., our covariates are creating a quality counterfactual).

We proceed by selecting a placebo group that is observably similar to the original protected group. We then run the same genetic matching specification, assigning the placebo group as the "treated" group, as in the primary matching analysis. The placebo group in this analysis comprises the 56 matched controls from the primary genetic matching analysis in the main text.[86] This group is observably similar to the original protected group (on average; see Table 11) with the exception that the placebo group was not affected by protected areas. Therefore, if our covariates are capturing underlying poverty trajectories well, then by selecting unprotected municipalities that are observably similar to our placebo group, we should observe *no* difference in outcomes because there is no longer protection as a source of variation between the two groups.

The results in Table 24, for both the full and unique placebo groups, indicate that there is no placebo effect. In other words, our covariates of interest appear to be predicting poverty trajectories well (see Tables 25 and 26 for balance results). These results buttress our claims that the treatment effects present in our main analyses are due to the establishment of protected areas rather than an inability to estimate quality counterfactuals.

C.3 Spillover Analysis

The central result from our main analyses is that municipalities that were affected by protected areas had differentially greater poverty reduction than comparable

[86]Recall that matching was performed with replacement so the placebo group has 15 repeat observations. We choose this for our primary placebo group because it most closely resembles our original protected group in expectation . We perform an additional analysis assigning only unique matched controls, from the original analysis, to the placebo group (see Table 25).

municipalities that were unaffected by protected areas. An often voiced concern is that protected areas, rather than having a positive impact of proximal populations, caused those most negatively impacted to emigrate from the impacted communities. If such emigration was undertaken by relatively poor populations (a supposition supported by Robalino 2007) it would have two effects. First, the departure of a relatively poor population would result in a decrease in the measured average poverty level within protected municipalities. Second, immigration of these relatively poor populations would have negative impacts on measured average poverty in surrounding municipalities. The former affect is a concern because it implies that there were no truly positive mechanisms through which protected areas affected poverty (e.g., tourism, infrastructure development, ecosystem services, etc.). Instead, the former implies that protected areas didn't make surrounding populations better-off, it just compelled those that they made the worst-off, to emigrate. The latter effect is one that we attempt to test. To do so we first assume that if the poor are negatively affected by protected areas, they will migrate to the nearest unaffected communities.

Our analysis to test local migration effects is thus framed as a spillover analysis. Using GIS we select all the municipalities that neighbor (congruent to) municipalities with at least 10% of their area occupied by a protected area (see Figure 21). If protected areas caused poor populations to migrate to surrounding communities, then we would expect an increase in poverty between 1992 and 2001 in these neighboring municipalities, compared to observably similar (unprotected) municipalities. To test this hypothesis we treat the 99 neighboring municipalities[87] as "treated" units and match them to observably similar unprotected municipalities (according to our covariates of interest). Under the null hypothesis of no spillover, there should be no treatment effect in the resulting matched sample. In other

[87] There are 116 municipalities that are congruent to a protected municipality. 17 are dropped from the sample because they are considered marginally protected.

words, there should be no difference in poverty outcomes between congruent (to protected municipalities) and matched unprotected municipalities.

We run our primary genetic matching and regression specifications on the spillover data, the results of which can be found in Table 27 (see Tables 28 and A11 for balance results). We find no evidence of negative spillover effects from protected municipalities into congruent communities. For both specifications in which we designate the PI as the outcome, the estimated impacts are quite small and statistically insignificant. In the specifications using NBI, we find estimates of poverty alleviation in congruent municipalities (compared to similar unprotected municipalities). These results are statistically significant (insignificant) in the regression (genetic matching) specification.

The results from our spillover analysis indicate that municipalities congruent to protected municipalities fared no worse (and by most indications, better) than similar unprotected municipalities. We, therefore, propose that the positive poverty impacts associated with the establishment of protected areas are unlikely due to the emigration of poor populations to surrounding communities. While this proposition may hold for our regional spillover analysis, it is difficult to test for broader general equilibrium migration effects.

One potential scenario is that the emigrants move to urban areas. If this were the case then we would expect to see less poverty reduction in urban areas as compared to protected municipalities, *ceteris paribus*. We attempt to capture this potential migration effect by limiting our control sample to municipalities that lie within 50km of a major city.[88] Our resulting sample comprises the original 56 protected municipalities and 53 unprotected municipalities. The regression results from this sample can seen found in Table 30. We find that, compared to relatively

[88]Unprotected municipalities remain in the sample if their average euclidean distance from each 1ha parcel is within 50km from a major city. We choose 50km to balance the tradeoff between capturing urban areas and sample size.

urban areas, protected areas still had differentially greater level of poverty reduction between 1992 and 2001. Though coarse, this provides evidence that our results are not consistent with poor populations being driven to urban areas.

One final piece of evidence that indicates our results are unlikely driven by complex migration patterns comes from the preceding placebo analysis. Aside from localized (congruent) or urban migration, it is not illogical to presume that adversely affected poor populations might migrate to municipalities that are observably similar to the protected municipalities from which they originate. Under this scenario we might reasonably assume that poor populations would migrate to the unprotected municipalities found in the matched control group from the primary genetic matching specification. However, these are the municipalities that compose our "treated" placebo group, for which we found no difference in poverty outcomes (compared to similar unprotected municipalities) in the placebo analysis. If our more complex migration scenario were occurring, we would expect to find a negative (poverty exacerbation) treatment effect in the placebo analysis.

Unfortunately it is not possible to fully capture all the potential general equilibrium poverty effects of protected areas. However, given the limited mobility of poor populations and migration scenarios explored, we believe that our analyses provide strong evidence that the positive impacts of protected areas are not driven by the emigration of poor populations.

C.4 Areas Formally Protected in the 1990s

Fundamental to our identification of the impacts of protected areas in our study period was Law 1333 and the associated restructuring and enforcement or protected areas subsequent to 1992. Despite the evidence of that protected areas were so-called "paper parks" (e.g., (Bruner et al. 2001)), if the 10 protected areas that

were established prior to our study period were in fact effective, this may be biasing our results.

There are a number of plausible impacts to our results that stem from including protected areas established (and maintained effectively) prior to our study period. However, our primary concern is that such inclusion would significantly increase the probability of inferring poverty alleviation associated with the establishment of protected areas. To address this potential bias, we drop the 10 protected areas that were originally established prior to our study period. We are left with 32 (30) protected municipalities when the PI (NBI) is used as the outcome of interest. We run our primary specifications on this sample and find the results to be strikingly similar to those from the original sample (see Table 31 for results and Tables 32 and 34 for balance results).

C.5 Rosenbaum Bounds

The ATT estimates from the primary genetic matching specifications represent unbiased estimates of the impact of protected areas on poverty under the assumption that we have sufficiently controlled for all covariates that jointly determine the spatial establishment of protection and poverty (conditional independence assumption). However, if there exists an unobserved covariate or group of covariates, that is highly correlated with protection and poverty, and uncorrelated with the covariates for which we do control, then we may be concerned that this confounder might be biasing our results. The fundamental concern is that the poverty alleviation observed in protected municipalities is due to systematic differences between protected and unprotected municipalities, other than protection.

One of the desirable properties of matching is that under CIA we can invoke many of the methods of inference used in a randomized experiment. Under pure randomization each selected unit has an equal probability of being assigned to the

treatment or control group. Therefore, under the null hypothesis of no treatment effect there is a $\Pr = 0.5$ that any unit within a pair chosen across treatment arms has a greater outcome than the other unit within the pair (outcomes are "exchangeable" within pairs). In other words, under randomization, if treatment has no effect, we should observe treated units within pairs exhibiting greater outcomes approximately 50% of the time, and control units exhibiting greater outcomes approximately 50% of the time.[89] If matching satisfies CIA then similar logic, and inference, can be applied to matched pairs.

Suppose that matching perfectly accounts for all covariates that affect outcome and selection. Similar to randomization, under the null hypothesis of no treatment effect, we should observe treated units exhibiting greater outcome values within matched pairs approximately 50%, and *vice versa.* This type of inference is valid within matched pairs because, conditional on covariates X, the probability of treatment within these pairs in equal and, therefore, outcomes within matched pairs are considered exchangeable.[90] Now suppose that there is some unobserved covariate, u, that is uncorrelated with X, but correlated with outcomes Y and treatment T. There are two ways of looking at the impact of u on inference. (1) u affects the probability of treatment such that exchangeability is no longer satisfied, therefore, invalidating permutation-based inference using a null of no treatment effect. (2) The differences in u, which are systematically related to T, are driving the observed differences in Y, otherwise attributed to T.

Rosenbaum (2002) proposes measures by which we can test the sensitivity of our matching results to the presence of u. Rosenbaum bounds allow us to measure how strong a confounder, u, would need to be to the invalidate our statistical findings.[91]

[89]Under randomization, there are a number of permutation-based inference tests by which to estimate exact p-values based on this logic (e.g., Rosenbaum 2002).

[90]Another way of expressing this is that, conditional of X, there remains no other source of variation between treated and control groups that affects Y, other than T.

[91]It should be noted prior to exposition that any measurement of sensitivity to unobserved bias, or varying degrees therein, does not imply the presences of unobserved bias.

In Rosenbaum's model the probability of assignment to treatment π_j for unit j can be expressed in terms of odds as $\frac{\pi_j}{(1-\pi_j)}$. Under randomization $\pi_j = \pi_k$ for $j \neq k$. Similarly, in an observational setting and in absence of u, $\pi_j = \pi_k$ when $x_j = x_k$. In other words, conditional of similar values of X within matched pairs, the probability of treatment is equal for treated and control units. The departure from randomization (or the influence of u), can be expressed by Γ in the odds ratio between matched pairs

$$\frac{1}{\Gamma} \leq \frac{\pi_j\,(1-\pi_k)}{\pi_k\,(1-\pi_j)} \leq \Gamma, \text{ for all } j,k \text{ with } x_j = x_k,$$

where $\Gamma = 1$ under randomization. Conversely, if $x_j = x_k$ but $\pi_j \neq \pi_k$ this implies the presence of u, the degree of which is captured by $\Gamma \neq 1$. For instance, if $\Gamma = 2$ (but $x_j = x_k$) this implies that the presence of u is causing the odds of treatment between j and k to differ by a factor of 2. In Rosenbaum's sensitivity test we ask how large Γ would need to be (i.e., how strong a confounder u would need to be) in order to alter matching-based inference.

To frame Γ explicitly in terms of the unobserved bias, u, Rosenbaum (2002) shows that the log odds ratio for j is equivalent to

$$\log\left(\frac{\pi_j}{(1-\pi_j)}\right) = k(x_j) + \gamma u_j, \text{ with } 0 \leq u_j \leq 1,$$

which states that the odds of treatment are an unknown function of x plus an unknown parametrization of u.[92] The odds ratio can therefore be rewritten as

$$\frac{\pi_j\,(1-\pi_k)}{\pi_k\,(1-\pi_j)} = \exp\{\gamma(u_j - u_k)\},$$

[92]See Rosenbaum (2002) for a discussion of the restriction on u.

where $k(\cdot)$ cancels when $x_j = x_k$. By stating the odds ratio in terms of u it can be seen that in the absence of u, i.e., when u does not influence π (or when $u_j = u_k$), $e^{\gamma} = \Gamma = 1$. Conversely, as the influence of u increases (or as u_j and u_k diverge) Γ and the absolute difference in probability of treatment between treated and untreated units increases.

In observational studies we cannot observe the presence of u or its potential influence as measured by $\Gamma = e^{\gamma}$. In Rosenbaum's sensitivity test we impose increasing levels of Γ to measure at what influence of unobserved bias our inference would be invalidated (shown to be insignificant). If inference is altered by a level of Γ close to 1 this implies that a study is sensitive to unobserved bias. However, we reiterate that estimated sensitivity to unobserved bias in no way implies the presence of unobserved bias.

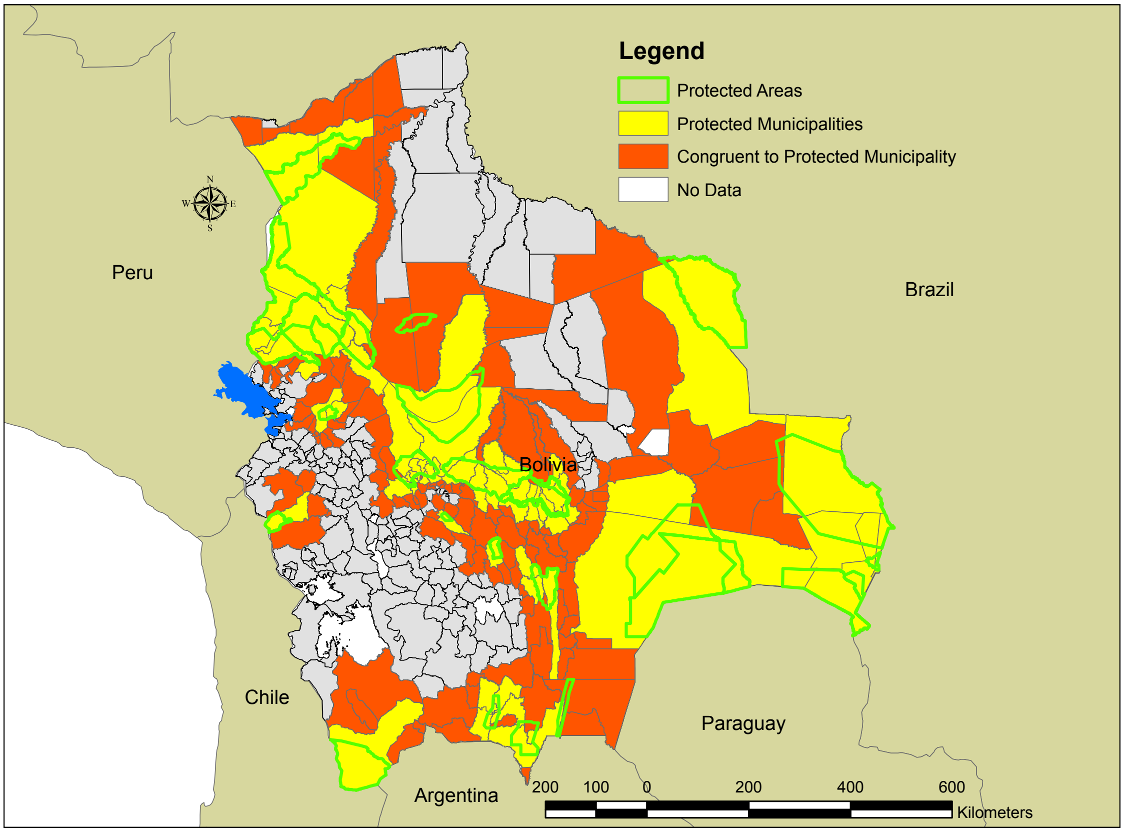

Figure 21: Map of municipalities congruent (orange) to a municipality with at least 10% area occupied by a protected area (yellow).

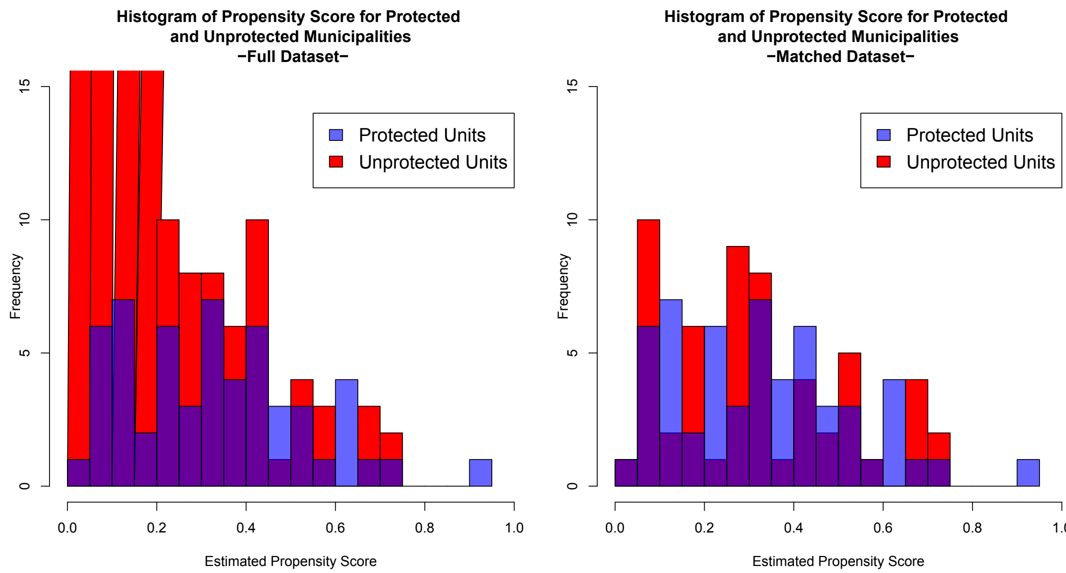

Figure 22: Histogram of propensity score distributions according to PI for full and matched datasets. Red bars indicate frequency of unprotected units and blue bars indicate frequency of protected units (purple represent areas where bars overlap).

Table 19: ATT Estimates from Primary GenMatch Specification by Protection Threshold

	Poverty Index			NBI		
Threshold	Y(T=1)	Y(T=0)	Treatment	Y(T=1)	Y(T=0)	Treatment
5%	-1.23	-0.761	-0.465	77.12	81.78	-4.66
	[63]	[63]	(0.136)	[60]	[60]	(3.394)
10%	-1.33	-0.805	-0.525	76.18	84.16	-4.99
(Primary Specification)	[56]	[56]	(0.142)	[53]	[53]	(3.67)
20%	-1.214	-0.725	-0.489	76.21	82.3	-6.075
	[42]	[42]	(0.149)	[39]	[39]	(4.75)
30%	-1.162	-0.731	-0.431	76.64	82.92	-6.282
	[38]	[38]	(0.154)	[36]	[36]	(5.1)
50%	-1.67	-1.223	-0.454	70.14	76.49	-6.35
	[18]	[18]	(0.272)	[17]	[17]	(3.52)

(Abadie-Imbens Heteroskedasticity Robust Standard Errors)
[Observation]

Table 20: Results from Primary and Ancillary Analyses, Treatment Threshold=5%

	Poverty Index			NBI		
Method	Y(T=1)	Y(T=0)	Treatment	Y(T=1)	Y(T=0)	Treatment
Nave Difference	-1.23	-0.4946	-0.731	77.12	85.89	-8.77
in Means	[63]	[261]	{0.0215}	[60]	[251]	{0.004}
Regression Dropping	NA	NA	-0.42	NA	NA	-2.04
Marginal	[63]	[46]	(0.103)	[60]	[42]	(1.55)
Post-Match Frequency	-1.23	-0.761	-0.465	77.12	81.78	-4.66
Weighted Regression	[63]	[63]	(0.136)	[60]	[60]	(3.394)
Genetic Matching	-0.986	-0.516	-0.471	79.72	82.73	-3.01
	[56]	[56]	(0.14)	[54]	[54]	(2.09)
Genetic Matching,	-1.07	-0.511	-0.56	79.04	81.51	-2.47
Calipers=1sd	[49]	[49]	(0.147)	[47]	[47]	(1.55)

[Number of observations]
(Standard errors)
{P-value}

Table 21: Results from Primary and Ancillary Analyses, Treatment Threshold=20%

	Poverty Index			NBI		
Method	Y(T=1)	Y(T=0)	Treatment	Y(T=1)	Y(T=0)	Treatment
Nave Difference	-1.21	-0.551	-0.663	76.2	85.3	-9.13
in Means	[42]	[282]	{0.018}	[39]	[272]	{0.013}
Regression Dropping	NA	NA	-0.527	NA	NA	-7.23
Marginal	[42]	[252]	(0.113)	[39]	[242]	(1.39)
Post-Match Frequency	NA	NA	-0.47	NA	NA	-2.83
Weighted Regression	[42]	[31]	(0.116)	[39]	[29]	(1.95)
Genetic Matching	-0.996	-0.475	-0.52	78.95	82.48	-3.53
	[36]	[36]	(0.155)	[35]	[35]	(3.1)
Genetic Matching,	-1.07	-0.511	-0.56	79.04	81.51	-2.47
Calipers=1sd	[49]	[49]	(0.147)	[47]	[47]	(1.55)

[Number of observations]
(Standard errors)
{P-value}

Table 22: Results from Primary and Ancillary Analyses, Treatment Threshold=30%

Method	Poverty Index Y(T=1)	Poverty Index Y(T=0)	Poverty Index Treatment	NBI Y(T=1)	NBI Y(T=0)	NBI Treatment
Nave Difference	-1.16	-0.567	-0.595	76.6	85.2	-8.551
in Means	[42]	[282]	{0.108}	[36]	[275]	{0.029}
Regression Dropping	NA	NA	-0.523	NA	NA	-7.42
Marginal	[38]	[252]	(0.119)	[36]	[242]	(1.44)
Post-Match Frequency	NA	NA	-0.45	NA	NA	-2.91
Weighted Regression	[38]	[30]	(0.124)	[36]	[28]	(2.08)
Genetic Matching	-1.162	-0.731	-0.431	76.64	82.92	-6.282
	[38]	[38]	(0.154)	[36]	[36]	(5.1)
Genetic Matching,	-0.908	-0.506	-0.402	79.71	83.59	-3.872
Calipers=1sd	[32]	[32]	(0.157)	[31]	[31]	(3.1)

[Number of observations]
(Standard errors)
{P-value}

Table 23: Results from Primary and Ancillary Analyses, Treatment Threshold=50%

Method	Poverty Index Y(T=1)	Poverty Index Y(T=0)	Poverty Index Treatment	NBI Y(T=1)	NBI Y(T=0)	NBI Treatment
Nave Difference	-1.67	-0.577	-1.1	70.14	85.01	-14.87
in Means	[18]	[306]	{0.09}	[17]	[294]	{0.04}
Regression Dropping	NA	NA	-0.684	NA	NA	-11.3
Marginal	[18]	[252]	(0.162)	[17]	[242]	(1.95)
Post-Match Frequency	NA	NA	-0.627	NA	NA	-8.52
Weighted Regression	[18]	[15]	(0.218)	[17]	[15]	(3.3)
Genetic Matching	-1.67	-1.223	-0.454	70.14	76.49	-6.35
	[18]	[18]	(0.272)	[17]	[17]	(3.52)
Genetic Matching,	-1.23	-0.451	-0.782	75.9	82.22	-6.33
Calipers=1sd	[15]	[15]	(0.245)	[14]	[14]	(6.95)

[Number of observations]
(Standard errors)
{P-value}

Table 24: Results from Placebo Matching Test

Method	Poverty Index Y(T=1)	Y(T=0)	Treatment
Genetic Matching	-0.797	-0.84	0.044
Full Matched Sample	[56]	[56]	(0.17)
Genetic Matching	-0.915	-0.7189	-0.196
Dropped Repeat	[41]	[41]	(0.156)

[Number of observations]
(Standard errors)

Table 25: Balance Results for Placebo Matching Analysis- Full Sample

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	0.319	0.834	-0.515	0.157	0.502	
1992	Matched	0.319	0.264	0.054	0.016	0.216	89.48%
% Forest 1991	Unmatched	0.466	0.148	0.317	0.531	0.315	
	Matched	0.466	0.458	0.008	0.012	0.035	97.51%
Distance to	Unmatched	136600	102500	34100	0.205	33260	
Major City	Matched	136600	137500	-868.1	0.004	13100	97.45%
Average	Unmatched	1794	2860	-1066	0.375	1052	
Elevation	Matched	1794	1936	-142	0.051	204.9	86.68%
Average Slope	Unmatched	23.89	18.06	5.822	0.191	5.989	
	Matched	23.89	24.17	-0.287	0.008	4.508	95.07%
Roadless	Unmatched	2.526E+14	6.101E+13	1.916E+14	0.259	1.860E+14	
Volume 1992	Matched	2.526E+14	1.269E+14	1.257E+14	0.162	1.547E+14	34.38%

Table 26: Balance Results for Placebo Matching Analysis- Unique Sample

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	0.282	0.834	-0.551	0.170	0.537	
1992	Matched	0.282	0.283	-0.001	0.000	0.188	99.89%
% Forest 1991	Unmatched	0.427	0.148	0.279	0.458	0.274	
	Matched	0.427	0.426	0.002	0.002	0.024	99.41%
Distance to	Unmatched	136100	102500	33560	0.201	32230	
Major City	Matched	136100	132900	3230	0.017	10430	90.38%
Average	Unmatched	1956	2860	-904.1	0.313	887.6	
Elevation	Matched	1956	2000	-44.070	0.015	147.3	95.13%
Average Slope	Unmatched	23.62	18.06	5.556	0.181	5.591	
	Matched	23.62	23.03	0.594	0.017	3.361	89.30%
Roadless	Unmatched	1.992E+14	6.101E+13	1.382E+14	0.212	1.265E+14	
Volume 1992	Matched	1.992E+14	1.415E+14	5.764E+13	0.080	1.097E+14	58.28%

Table 27: Spillover Analyses, Municipalities Congruent to Protected Municipalities Considered Treated

Method	Poverty Index Y(T=1)	Poverty Index Y(T=0)	Poverty Index Treatment	NBI Y(T=1)	NBI Y(T=0)	NBI Treatment
Regression Dropping Marginal	NA [99]	NA [153]	-0.14 (0.09)	NA [95]	NA [147]	-3.12 (0.98)
Genetic Matching	-0.629 [99]	-0.726 [99]	0.097 (0.185)	81.7 [95]	83.88 [95]	-2.187 (2.29)

[Number of observations]
(Standard errors)

Table 28: Balance Results for Congruent Spillover Analysis- PI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	0.678	0.787	-0.109	0.034	0.323	
1992	Matched	0.678	0.616	0.062	0.019	0.331	43.27%
% Forest 1991	Unmatched	0.319	0.113	0.206	0.355	0.208	
	Matched	0.319	0.311	0.008	0.012	0.020	96.30%
Distance to	Unmatched	109000	107400	1569	0.009	10650	
Major City	Matched	109000	110500	-1567	0.009	9140	0.06%
Average	Unmatched	2124	3095	-970.3	0.342	959.5	
Elevation	Matched	2124	2109	15.45	0.005	195.9	98.41%
Average Slope	Unmatched	21.66	17.23	4.43	0.145	4.688	
	Matched	21.66	20.70	0.96	0.029	3.121	78.28%
Roadless	Unmatched	7.926E+13	8.624E+13	-6.976E+12	0.009	3.20E+13	
Volume 1992	Matched	7.926E+13	7.774E+13	1.524E+12	0.002	1.95E+13	78.15%

Table 29: Balance Results for Congruent Spillover Analysis- NBI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
NBI 1992	Unmatched	89.760	91.880	-2.114	0.049	1.957	
	Matched	89.760	90.560	-0.797	0.018	1.936	62.32%
% Forest 1991	Unmatched	0.309	0.108	0.201	0.346	0.203	
	Matched	0.309	0.304	0.005	0.007	0.023	97.56%
Distance to	Unmatched	109200	108200	1048	0.006	10110	
Major City	Matched	109200	107300	1894	0.011	7354	-80.79%
Average	Unmatched	2158	3135	-977	0.348	965.8	
Elevation	Matched	2158	2137	20.97	0.008	167.7	97.85%
Average Slope	Unmatched	22.08	17.75	4.332	0.141	4.579	
	Matched	22.08	22.72	-0.643	0.019	2.743	85.15%
Roadless	Unmatched	8.04E+13	8.84E+13	-8.04E+12	0.010	3.26E+13	
Volume 1992	Matched	8.04E+13	7.70E+13	3.43E+12	0.004	2.08E+13	57.36%

Table 30: Spillover Regression Results, Controls Within 50km of Major City

Covariate/Outcome	Standard PI	Standard NBI
(Intercept)	-2.16***	-48.9***
	(0.266)	(6.09)
Protected	-0.469**	–4.88*
	(0.161)	(2.58)
Baseline Poverty	0.918***	1.26***
	(0.025)	(0.0648)
Percent Forest 1991	0.643*	16.48**
	(0.349)	(5.70)
Distance to Major City	2.44E-06***	1.87E-06***
	(6.87E-07)	(1.17E-05)
Average Elevation	2.15E-04**	-0.0049***
	(8.18E-05)	(0.00135)
Average Slope	-9.99E-04	0.071
	(0.0052)	(0.082)
Roadless Volume	7.83E-17	4.01E-15*
	(1.11E-16)	(1.77E-15)
	R^2=0.944	R^2=0.836
	DF=101	DF=96
	F=242	F=69.7

Notes: Outcomes are indicated at column heads and represent 2001 poverty index and NBI.
***, **, * Indicate significance at the (0.01, 0.05 and 0.1 level, respectively.)
(Standard Errors)

Table 31: Results from Primary and Ancillary Analyses, Protected Areas Established in 1990s

Method	Poverty Index Y(T=1)	Poverty Index Y(T=0)	Poverty Index Treatment	NBI Y(T=1)	NBI Y(T=0)	NBI Treatment
Regression Dropping	NA	NA	-0.599	NA	NA	-4.23
Marginal	[32]	[252]	(0.128)	[30]	[242]	(1.48)
Post-Match Frequency	NA	NA	-0.54	NA	NA	-1.87
Weighted Regression	[32]	[24]	(0.151)	[30]	[24]	(2.09)
Genetic Matching	-0.79	-0.313	-0.485	79.95	85.85	-5.89
	[32]	[32]	(0.205)	[30]	[30]	(5.38)

[Number of observations]
(Standard errors)

Table 32: Balance Results for Analysis Using Protected Areas Established in 1990s- PI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
Poverty Index	Unmatched	0.552	0.744	-0.192	0.056	0.329	
1992	Matched	0.552	0.563	-0.010	0.003	0.199	94.54%
% Forest 1991	Unmatched	0.510	0.194	0.316	0.534	0.316	
	Matched	0.510	0.501	0.009	0.014	0.048	97.28%
Distance to	Unmatched	175400	108000	67370	0.314	65140	
Major City	Matched	175400	145900	29480	0.123	41660	56.24%
Average	Unmatched	1649	2713	-1065.000	0.373	1057	
Elevation	Matched	1649	1584	65.060	0.027	130.400	93.89%
Average Slope	Unmatched	24.970	18.970	6.006	0.193	6.080	
	Matched	24.970	24.890	0.087	0.003	2.305	98.56%
Roadless	Unmatched	3.825E+14	8.350E+13	2.990E+14	0.234	2.558E+14	
Volume 1992	Matched	3.825E+14	2.808E+14	1.017E+14	0.061	1.632E+14	65.98%

Table 33: Balance Results for Analysis Using Protected Areas Established in 1990s- NBI

Covariate	Status	Mean Prot.	Mean Unprot.	Diff. in Means	Norm. Diff.	Mean eQQ Diff.	%Improve Mean Diff.
NBI 1992	Unmatched	89.11	91.050	-1.934	0.043	1.821	
	Matched	89.11	89.130	-0.017	0.000	1.069	0.991%
% Forest 1991	Unmatched	0.488	0.187	0.301	0.510	0.298	
	Matched	0.488	0.483	0.006	0.009	0.036	0.982%
Distance to	Unmatched	165600	108600	57030	0.271	54590	
Major City	Matched	165600	139400	26230	0.108	34620	0.540%
Average	Unmatched	1731	2751	-1021	0.360	1013	
Elevation	Matched	1731	1746	-14.74	0.006	112.6	0.986%
Average Slope	Unmatched	25.83	19.45	6.376	0.204	6.4	
	Matched	25.83	27.43	-1.602	0.045	3.239	0.749%
Roadless	Unmatched	3.924E+14	8.528E+13	3.071E+14	0.234	2.600E+14	
Volume 1992	Matched	3.924E+14	3.027E+14	8.969E+13	0.052	1.691E+14	0.708%

Table 34: Regression Results from Primary Specifications

	Standard		Post-Match Weighted	
Covariate/Outcome	PI	NBI	PI	NBI
(Intercept)	-2.01***	-42.3***	-1.42***	-39.4
	(0.157)	(3.41)	(0.288)	(6.37)
Protected	-0.535***	-5.62***	-0.494***	-2.63
	(0.099)	(1.2)	(0.106)	(1.7)
Baseline Poverty	0.896***	1.26***	0.927***	1.31***
	(0.017)	(0.0361)	(0.0235)	(0.066)
Percent Forest 1991	0.574**	6.22***	-0.0298	-2.73
	(0.192)	(2.34)	(0.302)	(4.75)
Distance to Major City	2.77E-06***	-2.47E-05***	2.22E-06***	1.12E-05
	(4.88E-08)	(5.98E-06)	(6.33E-07)	(9.69E-06)
Average Elevation	1.95E-04***	-0.0035***	-4.51E-05	6.70E-04
	(4.39E-05)	(0.00054)	(9.09E-05)	(0.0014)
Average Slope	-0.003	0.027	0.006	0.071
	(0.0027)	(0.032)	(0.005)	(0.074)
Roadless Volume	2.67E-17	1.96E-15*	-7.78E-17	1.58E-15
	(9.26E-17)	(1.11E-15)	(8.29E-17	(1.41E-15)
	R^2=0.918	R^2=0.848	R^2=0.951	R^2=0.845
	DF=300	DF=287	DF=89	DF=83
	F=481	F=229	F=245	F=64.8

Notes: Outcomes are indicated at column heads and represent 2001 poverty index and NBI.
***, **, * Indicate significance at the 0.01, 0.05 and 0.1 level, respectively.
(Standard Errors)

References

Abadie, A., D. Drukker, J. Herr, and G. Imbens. 2004. Implementing matching estimators for average treatment effects in stata. *Stata Journal* 4: 290–311.

Abadie, A. and G. Imbens. 2006. Large sample properties of matching estimators for average treatment effects. *Econometrica* 74, no. 1: 235–267.

Adams, W., R. Aveling, D. Brockington, B. Dickson, J. Elliott, J. Hutton, D. Roe, B. Vira, and W. Wolmer. 2004. Biodiversity conservation and the eradication of poverty. *Science* 306, no. 5699: 1146–1149.

Andam, K., P. Ferraro, K. Sims, A. Healy, and M. Holland. 2010. Protected areas reduced poverty in costa rica and thailand. *Proceedings of the National Academy of Sciences* 107, no. 22: 9996.

Andam, K. S., P. J. Ferraro, A. Pfaff, G. Sanchez-Azofeifa, and J. A. Robalino. 2008. Measuring the effectiveness of protected area networks in reducing deforestation. *Proceedings of the National Academy of Sciences* 105, no. 42: 16089–16094.

Angrist, J. and J. Pischke. 2009. *Mostly harmless econometrics: an empiricist's companion*. Princeton Univ Pr.

Bandyopadhyay, S. and G. Tembo. 2010. Household consumption and natural resource management around national parks in zambia. *Journal of Natural Resources Policy Research* 2, no. 1: 39–55.

Baron, R. and D. Kenny. 1986. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology* 51, no. 6: 1173–1182.

Brockington, D., J. Igoe, and K. Schmidt Soltau. 2006. Conservation, human rights, and poverty reduction. *Conservation Biology* 20, no. 1: 250–252.

Bruner, A., R. Gullison, R. Rice, and G. Da Fonseca. 2001. Effectiveness of parks in protecting tropical biodiversity. *Science* 291, no. 5501: 125.

Canavire-Bacarezza, G. and M. M. Hanauer. 2011. Estimating the impacts of bolivia's protected areas on poverty. *Working Paper*.

Cavatassi, R., B. Davis, and L. Lipper. 2004. Estimating poverty over time and space: construction of a time-variant poverty index for costa rica. Technical report, ESA Working Paper.

Cleveland, W.. 1979. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*: 829–836.

Cleveland, W. and S. Devlin. 1988. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association* 83, no. 403: 596–610.

Coad, L., A. Campbell, L. Miles, and K. Humphries. 2008. The costs and benefits of protected areas for local livelihoods: a review of the current literature.

Costanza, R., R. d'Arge, R. d. Groot, S. Farber, M. Grasso, B. Hannon, K. Limburg, S. Naeem, R. v. O'Neill, J. Paruelo, R. G. Raskin, P. Sutton, and M. v. d. Belt. 1997. The value of the world's ecosystem services and natural capital. *Nature* 387, no. 15: 253–260.

Cropper, M., C. Griffiths, and M. Mani. 1999. Roads, population pressures, and deforestation in thailand, 1976-1989. *Land Economics* 75, no. 1: 58–73.

Cropper, M., J. Puri, and C. Griffiths. 2001. Predicting the location of deforestation: The role of roads and protected areas in north thailand. *Land Economics* 77, no. 2: 172.

Crump, R., V. Hotz, G. Imbens, and O. Mitnik. 2008. Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics* 90, no. 3: 389–405.

Crump, R., V. Hotz, G. Imbens, and O. Mitnik. 2009. Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96, no. 1: 187.

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

Deaton, A.. 2009. Instruments of development: Randomization in the tropics, and the search for the elusive keys to economic development.

Dehejia, R. and S. Wahba. 2002. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and Statistics* 84, no. 1: 151–161.

Duffy-Deno, K.. 1998. The effect of federal wilderness on county growth in the intermountain western united states. *Journal of Regional Science* 38, no. 1: 109–136.

Elbers, C., J. Lanjouw, and P. Lanjouw. 2003. Micro-level estimation of poverty and inequality. *Econometrica*: 355–364.

Evans, S.. 1999. *The green republic: A conservation history of Costa Rica.* University of Texas Press.

Ferraro, P. J.. 2008. Protected areas and human well-being.

Ferraro, P. J. and M. M. Hanauer. 2011. Protecting ecosystems and alleviating poverty with parks and reserves: 'win-win' or tradeoffs? *Environmental and Resource Economics.* 48, no. 2: 269.

Ferraro, P. J., M. M. Hanauer, and K. R. Sims. 2011. Conditions associated with protected area success in conservation and poverty reduction. *Proceedings of the National Academy of Sciences.* Forthcoming.

Filmer, D. and L. Pritchett. 2001. Estimating wealth effects without expenditure dataor tears: An application to educational enrollments in states of india*. *Demography* 38, no. 1: 115–132.

Fleck, L. C., M. Amend, L. Painter, and J. Reid. 2006. Regional economic benefits from condercation: The cas of madidi. *Conservation Strategy Fund*.

Flores, C. and A. Flores-Lagunes. 2011. Identification and estimation of causal mechanisms and net effects of a treatment under unconfoundedness.

Frangakis, C. and D. Rubin. 2002. Principal stratification in causal inference. *Biometrics*: 21–29.

Gindling, T. and K. Terrell. 2004. Legal minimum wages and the wages of formal and informal sector workers in costa rica.

Hanauer, M. M.. 2011. Causal mechanisms of protected area impacts. *Working Paper*.

Healy, A. and S. Jitsuchon. 2007. Finding the poor in thailand. *Journal of Asian Economics* 18, no. 5: 739–759.

Heckman, J.. 2010. Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of Economic Literature* 48, no. 2: 356–98.

Ho, D., K. Imai, G. King, and E. Stuart. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis* 15: 199–236.

Holland, P.. 1986. Statistics and causal inference. *Journal of the American Statistical Association* 81, no. 396: 945–960.

ICT. 2010, July. Instituto costarricense de turismo.

Imai, K., L. Keele, and T. Yamamoto. 2010. Identification, inference, and sensitivity analysis for causal mediation effects. *Statistical Science* 25, no. 1: 51–71.

Imbens, G.. 2004. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics* 86, no. 1: 4–29.

Imbens, G. and J. Wooldridge. 2009. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature* 47, no. 1: 5–86.

Joppa, L. and A. Pfaff. 2009. High and far: Biases in the location of protected areas. *PLOS One* 4, no. 12.

Joppa, L. and A. Pfaff. 2010. Reassessing the forest impacts of protection: The challenge of nonrandom location and a corrective measure. *Annals of the New York Academy of Sciences* 1185: 135–139.

Koenker, R. and K. Hallock. 2001. Quantile regression. *Journal of Economic Perspectives* 15, no. 4: 143–156.

Lewis, D., G. Hunt, and A. Plantinga. 2002. Public conservation land and employment growth in the northern forest region. *Land Economics* 78, no. 2: 245.

Lewis, D., G. Hunt, and A. Plantinga. 2003. Does public lands policy affect local wage growth? *Growth and Change* 34, no. 1: 64–86.

Lokshin, M.. 2006. Difference-based semiparametric estimation of partial linear regression models. *Stata Journal* 6, no. 3: 377.

Manski, C.. 2005. *Social choice with partial knowledge of treatment response*. New York: Springer.

Margules, C. and R. Pressey. 2000. Systematic conservation planning. *Nature* 405, no. 6783: 243–253.

MEA. 2005. *Ecosystems and human well-being*. Millennium Ecosystem Assessment. Island Press.

Mealli, F. and D. Rubin. 2003. Assumptions allowing the estimation of direct causal effects. *Journal of Econometrics* 112, no. 1: 79–87.

Menkhaus, S. and D. Lober. 1996. International ecotourism and the valuation of tropical rainforests in Costa Rica. *Journal of Environmental Management* 47, no. 1: 1–10.

Meyer, B.. 1995. Natural and quasi-experiments in economics. *Journal of Business and Economic Statistics* 13, no. 2: 151–161.

Morgan, S. and C. Winship. 2007. *Counterfactuals and causal inference: Methods and principles for social research*. Cambridge Univ Pr.

Naidoo, R., A. Balmford, P. Ferraro, S. Polasky, T. Ricketts, and M. Rouget. 2006. Integrating economic costs into conservation planning. *Trends in Ecology & Evolution* 21, no. 12: 681–687.

Nelson, A. and K. M. Chomitz. 2009. Protected area effectiveness in reducing tropical deforestation: A global analysis of the impact of protection status. *World Bank Independent Evaluation Group*.

Pearl, J.. 2009. *Causality: models, reasoning, and inference*. Cambridge Univ Pr.

Pfaff, A., J. Robalino, G. Sanchez-Azofeifa, K. Andam, and P. Ferraro. 2009. Park location affects forest protection: Land characteristics cause differences in park impacts across costa rica. *The B.E. Journal of Economic Analysis and Policy* 9, no. 2: 5.

Reibel, M.. 2007. Geographic information systems and spatial data processing in demography: A review. *Population Research and Policy Review* 26, no. 5: 601–618.

Robalino, J.. 2007. Land conservation policies and income distribution: who bears the burden of our environmental efforts? *Environment and Development Economics* 12, no. 04: 521–533.

Robalino, J. and L. Villalobos-Fiatt. 2010. Conservation policies and labor markets: Unraveling the effects of national parks on local wages in costa rica. *IDEAS Discussion Papers*.

Rosenbaum, P.. 1984. The consquences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society. Series A (General)* 147, no. 5: 656–666.

Rosenbaum, P.. 2002. *Observational studies*. Springer.

Rubin, B. and G. Hyman. 2000. The extent and economic impacts of soil erosion in costa rica. In C. Hall, C. Perez, and G. Leclerc (Eds.), *Quantifying Sustainable Development: The Future of Tropical Economies*, pp. 449. Academic Press.

Rubin, D.. 2004. Direct and indirect causal effects via potential outcomes. *Scandinavian Journal of Statistics* 31, no. 2: 161–170.

Sachs, J., J. Baillie, W. Sutherland, P. Armsworth, N. Ash, J. Beddington, T. Blackburn, B. Collen, B. Gardiner, and K. Gaston. 2009. Biodiversity conservation and the millennium development goals. *Science* 325, no. 5947: 1502.

Sekhon, J.. 2007. Multivariate and propensity score matching software with automated balance optimization: The matching package for r. *Journal of Statistical Software* 10, no. 2: 1–51.

Sims, K. R.. 2010. Conservation and development: Evidence from thai protected areas. *Journal of Environmental Economics and Management* 60, no. 2: 94–114.

Somanathan, E., R. Prabhakar, and B. Mehta. 2009. Decentralization for cost-effective conservation. *Proceedings of the National Academy of Sciences* 106, no. 11: 4143.

UDAPE. 2006. Economia y politica economica. *Report*.

Watts, R., R. Compton, J. McCammon, C. Rich, S. Wright, T. Owens, and D. Ouren. 2007. Roadless space of the conterminous united states. *Science* 316, no. 5825: 736.

WDPA. 2009. World database on protected areas: 2009 annual release.

Wilkie, D., G. Morelli, J. Demmer, M. Starkey, P. Telfer, and M. Steil. 2006. Parks and people: Assessing the human welfare effects of establishing protected areas for biodiversity conservation. *Conservation Biology* 20, no. 1: 247.

Yáñez, E.. 2006. Existe un vinculo entre las areas protigidas y la pobreza? evidencia de tres areas protegidas de bolivia. *The Nature Conservancy Mimeo*.

Yatchew, A.. 1997. An elementary estimator of the partial linear model. *Economics Letters* 57, no. 2: 135–143.

Yatchew, A.. 1998. Nonparametric regression techniques in economics. *Journal of Economic Literature* 36, no. 2: 669–721.

VITA

Merlin Mack Hanauer was born July 22, 1978 in Forks of Salmon, California, USA. He holds a Bachelor of Arts in Economics (Minor: Computer Information Systems) from Humboldt State University (Arcata, CA), and a Master of Arts in Economics from Georgia State University (Atlanta, GA).

Merlin matriculated to Georgia State's Doctoral program in 2006 to study Environmental Economics. He has worked as a graduate research assistant for Dr. Paul J. Ferraro and as a graduate teaching assistant to Dr. Bruce Seaman (Principles of Microeconomics, Summer 2009). He was the sole instructor of Principles of Microeconomics (Fall 2009) and The Global Economy (Spring 2010).

While at Georgia State Merlin has been the recipient of several internal and external awards. In 2010 he received the Theodor C. Boyden Excellence in Teaching Economics Award, and the European Association of Environmental and Resource Economists Travel Award. In 2011 Merlin received the Jack Blicksilver Scholarship in Economics, the Andrew Young School Excellence in Teaching Economics Award, and was the invited speaker at the Andrew Young School of Policy Studies Honors Day Dinner.

Merlin has also received internal and external funding for his research while at Georgia State. In 2010 he and a colleague received a grant from the Latin American Center for Environmental Economics Program (LACEEP) to study the impacts of protected areas on poverty in Bolivia. Also in 2010, Merlin was awarded Georgia State's Dissertation Improvement Grant.

Merlin's research has been published in *The Proceedings of the National Academy of Sciences* and *Environmental and Resource Economics*. He has presented his research at numerous universities, the Bellpaso International Summer School, and at conferences held by the Southern Economic Association and the Western Economic Association. He has also served as a referee for the *Energy Journal*.

Merlin received his Doctor of Philosophy degree in Economics from Georgia State University in May, 2011. He accepted a tenure-track faculty position as an Assistant Professor of Economics (Environmental) at Sonoma State University in Rohnert Park, California. To contact Merlin, please visit his homepage, `https://sites.google.com/site/merlinhanauer/`.