

Estimating the Impacts of Bolivia's Protected Areas on Poverty

GUSTAVO CANAVIRE-BACARREZA

Universidad EAFIT, Medellin, Colombia

and

MERLIN M. HANAUER *

Sonoma State University, Rohnert Park, USA

Summary. — Bolivia has experienced a rapid proliferation of protected areas over the past several decades. Using rich biophysical and socioeconomic data, we estimate the impacts of Bolivia's protected areas on poverty between 1992 and 2001. We find no evidence that the establishment of protected areas exacerbated poverty. In fact, we provide robust estimates indicating that communities affected by protected areas exhibited differentially greater levels of poverty reduction as compared to similar unaffected communities. Our results add to the growing body of evidence that environmental conservation does not necessarily come at the expense of local communities. © 2012 Elsevier Ltd. All rights reserved.

Key words — poverty, protected areas, impacts, program evaluation, South America, Bolivia

1. 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, 2011). The sheer scale of protected areas' global coverage 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, Campbell, Miles, & Humphries, 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). Such facts raise concerns from poverty advocates that achieving environmental goals may come at the expense of the populations impacted by protection (Adams *et al.*, 2004; Coad *et al.*, 2008).

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 globally, the costs are borne only by those proximal to the areas (Adams *et al.*, 2004; Coad *et al.*, 2008; Wilkie *et al.*, 2006). 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, Amend, Painter, & Reid, 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 & Hanauer, 2011; Ferraro, Hanauer, & Sims, 2011; Andam, Ferraro, Sims, Healy, & Holland, 2010; Sims, 2010; Robalino

& Villalobos-Fiatt, 2011). 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.¹

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.² Yet despite having a wealth of natural resources, Bolivia is one of the poorest countries in Latin America, with poverty levels upwards of 60% (Canavire-Bacarreza & Mariscal, 2010). 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.

We combine these rich biophysical and socioeconomic data with quasi-experimental methods to ask, "what would poverty outcomes in Bolivian communities affected by protected areas have been had protected areas not been established?" We find

* Funding for this study comes from the Latin American and Caribbean Environmental Economics Program (LACEEP), Small Grants Program for Research in Environmental Economics. We would like to thank Andrew Balthrop, Allen Blackman, Mario Boudoin, Salvatore Di Falco, Paul J. Ferraro, Alfonso Malky, Juan Robalino, three anonymous reviewers and the participants at the LACEEP seminars in Costa Rica for helpful comments and suggestions. We extend deep gratitude to Juan Carlos Ledezma for providing and discussing data. Final revision accepted: June 12, 2012.

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 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 lead to an overestimation of the poverty alleviation associated with the establishment of protected areas. The contribution of our paper is thus twofold. First, the observed differences in impact estimates and the characteristics of protected areas, as compared to previous studies, show that not all aspects of such studies will be generalizable across countries. Therefore, in order to better understand the socioeconomic impacts of protected areas, similar studies to ours should be done on a country-by-country basis. Second, and most importantly, our study adds to the global knowledge base on the socioeconomic impact of protected areas and corroborates previous evidence that conservation policy is not necessarily at odds with development goals.

(a) Background

(i) Related literature

There have been only a handful of studies on the socioeconomic impacts of protected areas that properly control for their nonrandom establishment. Further, there is no formal literature regarding the impacts of Bolivia's protected area network on socioeconomic outcomes. 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 and 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 (2011) 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, nonagricultural 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, Hunt, & Plantinga, 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, 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 Bureau of Land Management 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.

(ii) 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 nontrivial amount of Bolivia's area was designated for protection between the late 30's and early 90's.³ 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 which lead to a general disregard of what rules were established (Sernap, 2007). Furthermore, there was a lack requisite enforcement within these areas, a phenomenon commonly referred to as, "paper parks" (e.g., Bruner, Gullison, Rice, & Da Fonseca, 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. Given the perceived failure of the older parks, many new parks were established, and old parks restructured to include local actors (e.g., populations living within park boundaries) in the management process.

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, esthetic, historical, economic, and social interest, in order to conserve and protect natural and cultural heritage of the country."⁴

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 impact of protected areas can be measured.⁵

2. DATA AND METHODS

(a) 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. We use Geographic Information Systems (GIS, ArcGIS 10) to compile and analyze all spatial data.

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 the municipal level, such as structural poverty measures, education, employment, housing, and health. Information regarding Bolivia's protected areas and their boundaries was obtained from Servicio 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).

(i) 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 327 municipalities with an average area of 325,083 ha (about the size of Rhode Island, USA; range: 28,928–1,298,121 ha). The maps in Figures 1 and 2 show that the municipalities in the mountainous and altiplano regions (southwest) tend to be smaller than those in the lowlands (east northeast; see Figure 3 for the topography of Bolivia).

(ii) Protected areas and treatment assignment

Bolivia has 23 protected areas that cover a total area of 17,131,507.48 ha or ~16% 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 establish a threshold (percentage), above which municipalities are considered protected. Our initial threshold is established at 10%.⁶ 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

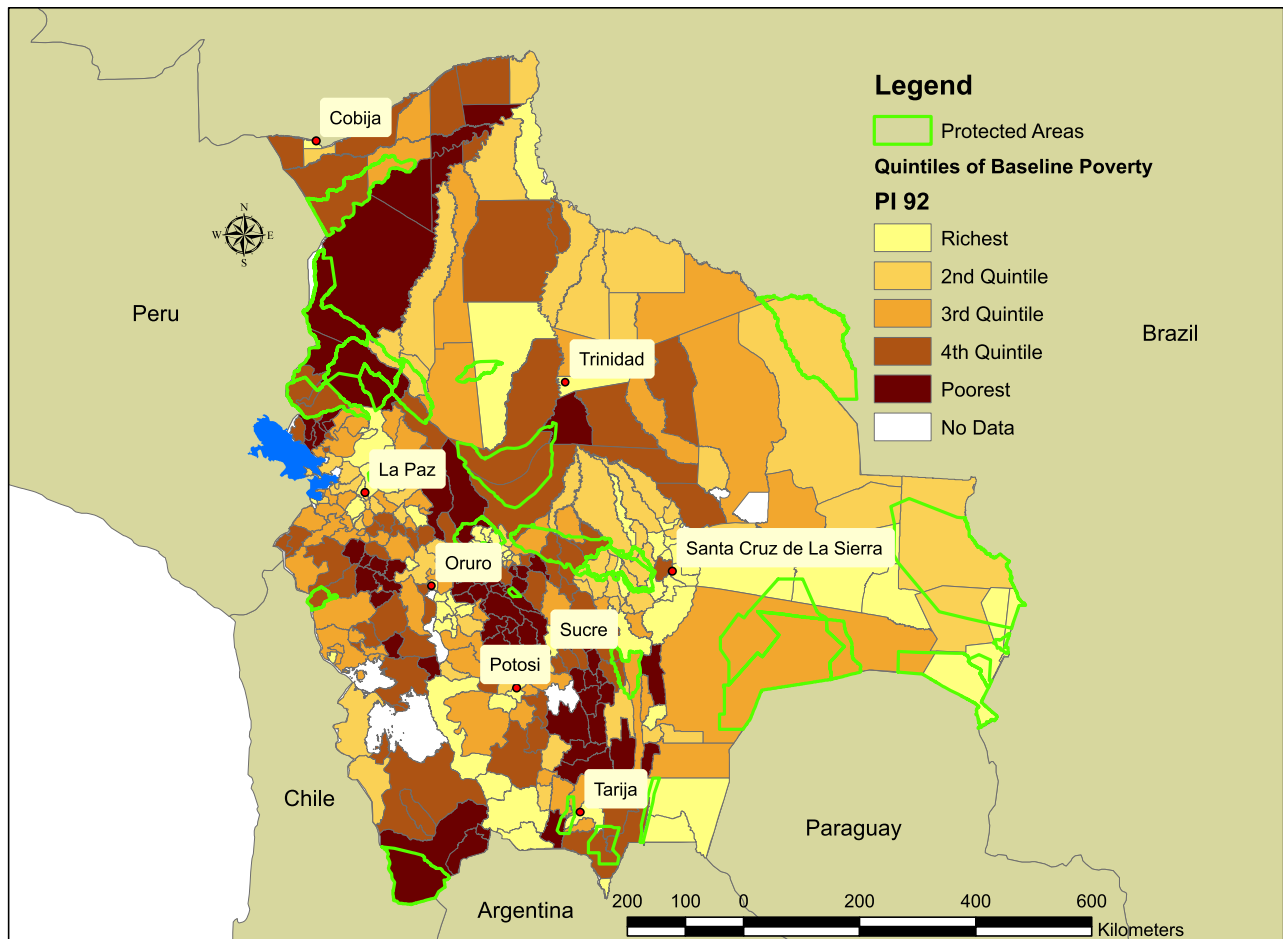


Figure 1. Map of protected areas, major cities, and quintiles of poverty in 1992 according to poverty index.

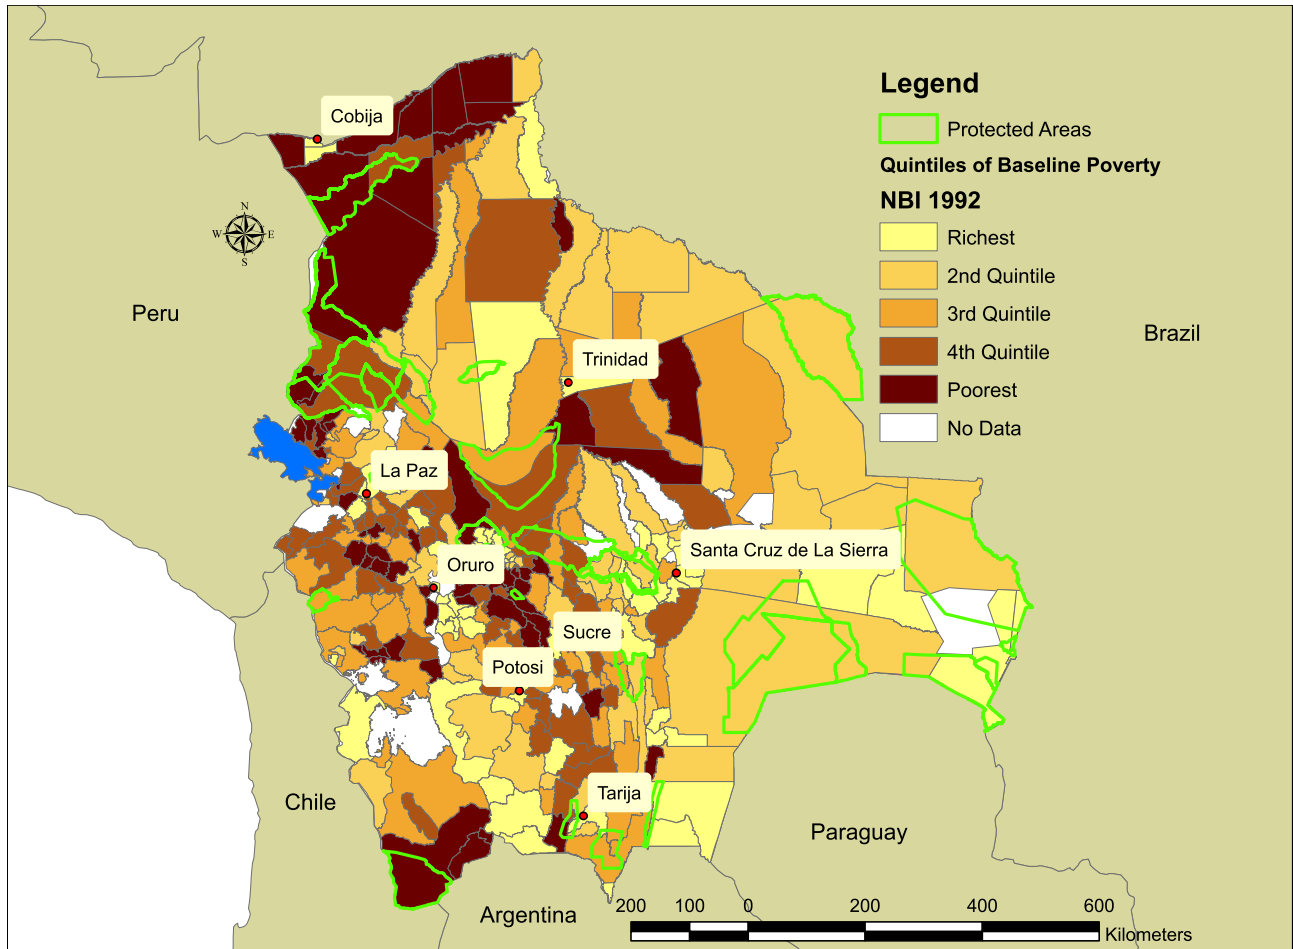


Figure 2. Map of protected areas, major cities, and quintiles of poverty in 1992 according to NBI.




Figure 3. 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 1 and 2. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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% to 100% (mean = 43.9%, median = 39.14%). We are left with 256 municipalities that are considered unprotected.

(iii) 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 & 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 (and El Alto), Sucre, Cochabamba, Cobija, Trinidad, Oruro, Potosi, Tarija, and Santa Cruz. Table 1 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; Joppa & Pfaff, 2009; Sims, 2010).

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; Andam, Ferraro, Pfaff, Sanchez-Azofeifa, & Robalino, 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 in this case) and the distance from that parcel to the nearest road (in 1992). Therefore, higher measurements of roadless volume indicate fewer road networks within a municipality. Table 1 shows that roadless volume is greater within protected municipalities, which is consistent with previous studies (e.g., Andam *et al.*, 2010).⁷

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,

Table 1. Summary statistics for poverty outcomes and covariates of interest

Variable	Description	Status	Mean	Median	Std	Min	Max
Poverty index 1992	Asset-based poverty index for 1992	Unprotected	0.74	1.20	2.08	-5.89	3.97
Poverty index 2001**	Asset-based poverty index for 2001	Protected	0.26	1.06	2.41	-5.94	3.56
		Unprotected	-0.45	-0.05	2.01	-6.59	2.94
NBI 1992	% population with unmet basic needs 1992	Protected	-1.33	-1.13	2.31	-6.88	2.35
		Unprotected	91.05	95.33	11.52	44.21	100.00
NBI 2001***	% population with unmet basic needs 2001	Protected	87.52	93.28	14.07	45.76	99.67
		Unprotected	85.85	86.02	16.78	19.07	100.00
Area (m)***	Municipality area, square kilometers	Protected	76.17	76.7	22.06	23.83	99.84
		Unprotected	2,178	931	3,652	12	35,110
Percent Forest Cover 1991***	% of municipality under forest cover 1991	Protected	6,937	2,214	12,701	217	71,358
		Unprotected	0.19	0.00	0.31	0.00	0.99
Average Distance to a City (km)**	Avg. dist. to city from each 1ha parcel within each municipality	Protected	0.48	0.58	0.34	0.00	0.93
		Unprotected	108	94	72	6.8	384
Average Elevation (m)***	Average elevation of each 1ha parcel within each municipality	Protected	143	104	131	8.4	589
		Unprotected	2,713	3,265	1470	143	4,478
Average Slope (pct)**	Average slope of each 1ha parcel within each municipality	Protected	1,825	1,695	1357	149	4,589
		Unprotected	18.97	16.58	14.82	0.91	59.35
Roadless Volume (km)*	Sum of the product of area and dist. to road, each 1ha parcel	Protected	23.89	27.38	14.31	1.29	55.74
		Unprotected	8.349e+07	9.801E+6	2.604E+8	1.342E+4	2.53E+9
		Protected	2.978e+08	2.544E+7	8.477E+8	6.059E+5	5.61E+9

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.

* indicate that means between unprotected and protected municipalities are significantly different according to Welch Two Sample *t*-test at 10% levels.

** indicate that means between unprotected and protected municipalities are significantly different according to Welch Two Sample *t*-test at 5% levels.

*** indicate that means between unprotected and protected municipalities are significantly different according to Welch Two Sample *t*-test at 1% levels.

both globally (Joppa & Pfaff, 2009) and at the country-level (Andam *et al.*, 2010; Sims, 2010), have found that protected areas tend to be placed on land that is relatively steep and at high elevation. Therefore, it is important to control for the average slope and elevation of municipalities. We calculate the average slope and elevation for each municipality based on the respective values of each 1 ha parcel therein. Bolivia presents somewhat of a unique case. Bolivia is characterized by a dichotomous landscape in that the country is defined (generally) by the highlands and lowlands. Table 1 shows that the average slope in protected municipalities is greater than 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 1 shows that protected municipalities contained significantly more forested areas at baseline.

(iv) Poverty

In order to estimate the impact of protected areas on poverty, an outcome that objectively measures levels of socioeconomic welfare is necessary. Income or other data that are externally commensurable are not available during our study period. Therefore, we are tasked with developing metrics that adequately capture socioeconomic outcomes. Due to the characteristics of the available data and the nature of our analysis, we choose to use structural measures of poverty. Specifically, we employ two measures of poverty in our analyses: (1) an asset based poverty index (PI), and, (2) a measure of unsatisfied basic needs (Necesidades Basicas Insatisfechas or 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 outcomes 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.

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.

Poverty index (PI). 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).⁸ 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.

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 & 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

Table 2. Eigenvectors from principal component analysis

Eigenvectors, pooled variable	EigenV
Adult men in population ^a	-0.02836
Households without bathroom ^a	0.34984
Households that use fuelwood for cooking ^a	0.39719
Households with dirt floors ^a	0.39399
Low-quality houses ^a	0.33074
Households without electricity ^a	0.46972
Illiterate population ^a	0.17906
Population employed with salary ^a	-0.06499
Average persons per bedroom	0.00384
Households without access to public water ^a	0.39091
Households without sewer or septic ^a	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.

^a Indicates that variable is measured as a percentage.

changes in poverty. Table 2 lists the variables used in the construction of the poverty index, and the eigenvectors associated with each asset. These eigenvectors (from the first component) account for over 60% of the variation in the asset variables and provide the factor scores F_j for asset $j \in \{1, 2, \dots, J\}$ which indicate the weight and direction of the influence each asset a_j exerts on the PI. Factor scores are combined with observed asset levels to formulate the PI for municipality $i \in \{1, 2, \dots, N\}$,

$$PI_i = \sum_j F_j \left[\frac{a_{ij} - \bar{a}_j}{s_j} \right], \quad (1)$$

where a_{ij} is the observed level of asset j in municipality i , \bar{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 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 2 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. According to the manner in which the poverty index was constructed, poverty is decreasing in the negative orthant.

We provide further evidence of the internal validity of our PI in Table 3, 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 find that this is the case. For instance, the correlation between NBI and PI in 2001 is over 0.88, which can be observed spatially in the maps in Figures 1 and 2.

NBI. The NBI measures the percentage of the population within a municipality with unsatisfied basic human needs. NBI captures poverty by comparing municipality-level measures of goods and services (associated with well-being) that households possess to a norm (or ideal). The Bolivian NBI was estimated by the INE in coordination with UDAPE. It

Table 3. Mean asset values by deciles of 1992 poverty index

Variable	Richer			Deciles of PI 1992				Poorer		
	1	2	3	4	5	6	7	8	9	10
Adult men in population	57.77	55.96	55.78	54.74	54.45	56.5	55.1	54.56	54.3	57.9
Without bathroom	43.96	66.88	61.13	68.52	77.72	86.13	82.14	85.15	92.49	99.26
Use fuelwood for cooking	25.59	46.61	70.06	69.85	65.79	74.67	82.98	89.47	91.4	99.2
With dirt floors	28.77	49.74	63.72	72.77	76.89	79.72	80.52	81.61	92.19	91.7
Low-quality houses	15.23	23.07	37.02	41.09	49.08	47.7	57.57	58.63	68.46	82.57
Without electricity	26.35	45.1	63.42	73.31	82.13	86.89	94.17	94.65	97.19	99.87
Illiterate population	12.5	19.32	23.11	27.04	24.7	33.12	30.24	37.22	47.19	67.91
Population employed with salary	95.8	91.28	90.17	88.66	89.6	88.09	85.62	84.93	81.45	61.62
Average persons per bedroom	3.28	3.26	3.48	3.52	3.42	3.34	3.53	3.57	3.77	4.09
No access to public water	27.13	46.12	59.33	69.38	75.83	75.63	77.57	83.98	90.18	97.56
Without sewer or septic	65.05	87.19	90.97	93.19	96.16	96.98	97.34	98.5	99.03	99.87
Average years of education	6.77	5.1	4.66	4.2	4.42	3.68	3.77	3.36	2.72	1.45

comprises a set of factors such as housing, basic services, education, and health.⁹ 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 aforementioned 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 household.¹⁰ Accordingly, a higher measure of the NBI indicates greater poverty within the associated municipality.

(v) Baseline covariate distributions

Previous studies at the global and country level have found significant differences in the biophysical (Andam *et al.*, 2010; Joppa & Pfaff, 2009; Pfaff, Robalino, Sanchez-Azofeifa, Andam, & Ferraro, 2009; Sims, 2010) and socioeconomic (Andam *et al.*, 2010) characteristics underlying protected and unprotected areas. The differences underscore the nonrandom 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 & 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 1 shows 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 municipalities is lower than in unprotected municipalities (see Figure 3). 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 conventional wisdom. Although some of characteristics exhibited by Bolivia's protected areas correspond to trends seen in previous studies, while others do not, it is important to note that all covariates (other than baseline poverty) differ significantly between protected and unprotected municipalities.

(b) Methods

The underlying differences in covariate values between protected and unprotected municipalities underscores 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. Nonrandom allocation leads to the observed imbalance across these key covariates that jointly determine selection into protection and socioeconomic outcomes (see Table 1), 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.

(i) Matching

We measure the impact of protected areas on poverty in surrounding municipalities using matching to quantify 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 requires 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.¹¹

The key to matching as an identification strategy is the balancing of covariate distributions across treatment arms (protected and unprotected), thus mimicking the identification strategy of a randomized experiment. Covariate balance is achieved in expectation through 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,

$$E[Y(1)|T = (1)] = E[Y(1)|T = (0)] = E[Y(1)], \quad (2)$$

$$E[Y(0)|T = (1)] = E[Y(0)|T = (0)] = E[Y(0)]. \quad (3)$$

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 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),

$$E[Y(0)|T = 1, X] = E[(Y(0)|T = 0, X)] = E[Y(0)|X]. \quad (4)$$

Eqn. (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 the ATT on matched data sets. By ensuring that units are comparable across treatment and control groups, we make the CIA, which is necessary for causal inference, more defensible (Angrist & Pischke, 2009).

(ii) 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 interest.¹² 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 Section 3(b) 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 Dehejia & Wahba, 2002; Imbens & Wooldridge, 2009) during matching, use a post-match regression bias-adjustment (Abadie, Drukker, Herr, & Imbens, 2004; Abadie & Imbens, 2006; Imbens & Wooldridge, 2009), 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, thus isolating the only remaining variation between treatment arms to be the establishment of protected areas and allowing the unbiased estimation of ATT.

3. RESULTS

(a) Primary results

Table 4 present the results from our primary and ancillary robustness analyses (Tables 5 and 6 provide balancing results from our primary matching specifications). In this subsection we focus on the results found in the penultimate row of Table 4, Genetic Matching, which presents the ATT estimates from our primary genetic matching algorithm.

For both our PI and NBI poverty indicators we find no evidence to suggest the establishment of protected areas in Bolivia had deleterious effects on poverty. To the contrary, all of our point estimates indicate 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 estimate from the primary specification is 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.

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 4 highlights the underlying difference between the naïve (first row of Table 4) 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 (86.29) 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). The associated counterfactual outcome estimates are more similar to the treated sample (-0.805 for the PI and 81.16 for NBI), resulting in a more modest estimate of the poverty reduction associated with the establishment of protected areas for both poverty metrics.¹³

Our results are concordant with previous studies from Costa Rica (Andam *et al.*, 2010) and Thailand (Andam *et al.*, 2010;

Table 4. *Results from primary and ancillary analyses. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each treatment column is calculated as $Y(T = 1) - Y(T = 0)$*

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Naïve difference in means	-1.33 [56]	-0.451 [252]	-0.838*** {0.014}	76.18 [53]	86.29 [242]	-10.11*** {0.005}
Regression on raw data	-1.33 [56]	-0.828 [268]	-0.502*** (0.098)	76.18 [53]	81.7 [258]	-5.52*** (1.17)
Regression dropping marginal	-1.33 [56]	-0.795 [252]	-0.535*** (0.099)	76.18 [53]	81.79 [242]	-5.62*** (1.2)
Post-match frequency weighted regression	-1.33 [56]	-0.836 [41]	-0.494*** (0.106)	76.18 [53]	78.81 [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}.

*** Treatment effect estimates are different from zero at 1% levels.

Table 5. *Balance results for primary GenMatch specification—PI*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%Improve mean diff.
Poverty index 1992	Unmatched	0.267	0.744	-0.477	0.141	0.436	
	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 major city (km)	Unmatched	142.8	108	34.79*	0.173	35.95	
	Matched	142.8	136.6	6.15	0.029	32.94	82.32
Average elevation	Unmatched	1825	2713	-888***	0.307	884	
	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+8	8.35E+7	2.14E+8*	0.194	1.96E+8	
Volume 1992 (km)	Matched	2.98E+8	2.53E+8	4.53E+7	0.030	1.27E+8	78.89

None of the post-match differences are significant.

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

* Means are significantly different according to Welch Two Sample *t*-test at 10% levels.

*** Means are significantly different according to Welch Two Sample *t*-test at 1% levels.

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.

(b) 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.

(i) Internal robustness of matching specification

In any observational study it must be acknowledged that the ability to eliminate bias associated with nonrandom selection is limited by one's understanding of the underlying selection process (Meyer, 1995), and 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 are uncorrelated with observable characteristics for which we control but 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

Table 6. Balance results for primary GenMatch specification—NBI

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%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 major city	Unmatched	136.4	108.6	27.79	0.140	29.58	
	Matched	136.4	132.6	3.73	0.017	25.16	86.58
Average elevation	Unmatched	1888	2751	−863.1***	0.299	862.8	
	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+8	8.53E+7	2.20E+8*	0.195	1.99E+8	
Volume 1992	Matched	3.06E+8	2.58E+8	4.80E+7	0.033	1.06E+8	78.22

** Means are significantly different according to Welch Two Sample *t*-test at 5% levels. None of the post-match differences are significant.

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

* Means are significantly different according to Welch Two Sample *t*-test at 10% levels.

*** Means are significantly different according to Welch Two Sample *t*-test at 1% levels.

Table 7. Rosenbaum upper bound on *P*-value at given levels of Γ for primary matching analysis—PI

Γ	Upper bound <i>P</i> -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

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). Table 7 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 (see Appendix for full details on Rosenbaum bounds).

(ii) 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 Table 4.

Genetic matching with calipers. In our primary econometric specification we use one-to-one nearest neighbor matching based on the genetic matching algorithm. Although we achieve a high level of balance across treatment arms (see Tables 5 and 6), 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 differs by more than one standard deviation across covariate values.

Table 4 shows that the results for the PI are relatively robust to the introduction of calipers (bottom row), 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 4 indicates that this change is due to opposite movements in average poverty across the protected and unprotected samples. Average NBI in protected municipalities increased from 76.18 to 79.04 while average NBI in unprotected municipalities increases slightly from 81.16 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 4 (see Appendix Table 22 for full regression results). 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.¹⁴

Both the regression-based analyses for which the PI is the outcome of interest return estimates that are strikingly similar to one another, and to the primary matching specification. The precision of the estimates is slightly higher than for 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.

(c) *Results summary*

The central finding from our results is that it does not appear as though Bolivia's protected areas that were established between 1992 and 2000 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.

(i) *Effect size of impacts*

To measure poverty, and changes therein, we use two metrics that are neither directly comparable to one another nor to external measures. It is, therefore, difficult to interpret whether or not the estimated socioeconomic impacts that we report are consistent across metrics, or economically interesting in terms of magnitude. To address these issues, we calculate the effect size (Cohen, 1988) for each of our primary impact estimates.

Calculation of the effect size involves dividing the respective ATT estimates by the standard deviation of the matched unprotected units from the respective analyses. The effect size is, therefore, a standardized measure of by how much the distribution of the treatment group differs from that of the control group. The effect size estimates from our primary PI and NBI analyses are both -0.25 . This can be interpreted as follows: protected municipalities had, on average, 0.25 standard deviations lower PI and NBI than the associated matched controls.¹⁵ According to Cohen (1988) this is a relatively small effect size; nontrivial, however.

Another benefit of calculating the effect size is that it allows us to draw comparison between our results and those of similar studies. Our estimates fall in between those from Andam *et al.* (2010) who estimate poverty alleviating effect sizes of -0.2 and -0.43 for Costa Rica and Thailand, respectively.

(ii) *Why are the matching-based and regression-based estimates so similar?*

The goal 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 imposing a functional relationship between treatment, covariates, and outcome. Regression 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 scheme¹⁶ 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 are so similar across matching- and regression-based estimates can be seen in Figure A2. 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 A2). Therefore, it makes sense that the matching- and regression-based results are comparable because, even with the full sample, the regression results are not plagued by "out of sample" predictions.

4. 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 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, our results do not indicate that surrounding communities 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 experienced differentially greater poverty reduction than those unaffected by protected areas. Indeed, we employ two separate measures of poverty in our analyses and find that the point estimates of poverty reduction are robust across our econometric specifications.

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 differ subtly, but significantly. In those studies the authors found that controlling for key observable covariates led 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 socioeconomic impacts of protected areas. 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-by-country analyses of the socioeconomic impacts of protected areas.

Protected areas in Bolivia exhibit many of the characteristics observed globally (e.g. located relatively distant from major cities, roads, and on steeper slopes), however, some of the key drivers of poverty differ in important ways from previous global and country-level observations. We believe such differences 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 of protected areas requires that the scientific body of knowledge be built on a country-by-country basis.

Further, we note that, as with any observational study, unbiased impact estimates rely on satisfaction of the conditional independence assumption (CIA). While there is no formal test of the CIA, we attempt to show the robustness of our results by employing a variety of econometric specifications and multiple measures of poverty. Perhaps the best means to improve the estimates for Bolivia would be through the procurement and use of more highly disaggregated data. Such data would allow for more precise identification of the communities truly affected by protected areas and greater statistical precision.

In addition, our study brings to the fore important questions regarding *why* protected areas have been found to be associ-

ated with reductions in poverty. In discussions with local experts several common themes arose. First, the establishment of joint management between communities and the government allowed for the inclusion of communities in environmental education and implementation of projects and programs. Second, nature-based tourism has affected the structural employment composition in communities surrounding protected areas, e.g., hiring of locals as park guards. Third, after the restructuring of the park system in 1992 attention was paid to access. This led to the establishment of more, and improved, infrastructure surrounding protected areas. All of these observations provide insight into the potential channels through which protected areas have positively affected poverty in Bolivia. Unfortunately, empirical tests of such observations require data with a greater level socioeconomic disaggregation than are currently available.

Despite these challenges, further studies in Bolivia and elsewhere should strive to identify and quantify the mechanisms through which protected areas affect poverty (e.g., Ferraro & Hanauer, 2012; Robalino & Villalobos-Fiatt, 2011). 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.

NOTES

1. 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.

2. 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 (2007). Moreover, according to a UNESCO report, Bolivia has the largest water reserves in Latin America and ranks 6th in the world in terms of tropical moist forest resources (the third in the continent after Brazil and Mexico).

3. 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.

4. 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 fundamental to the goals of the system.

5. We use all categories of nationally established protected areas in our analyses: National Park, National Monument, Sylvester life Sanctuary, Silvester Reserve, and Integrated Management.

6. We use the 10% threshold in accordance with previous studies (Andam *et al.*, 2010; Ferraro & Hanauer, 2011; Ferraro *et al.*, 2011). A 10% threshold was chosen because one of the goals set forth by 4th

World Congress on National Parks and Protected Areas was to protect 10% of the world's 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 for full threshold results).

7. 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.

8. A similar asset based poverty index was developed for Costa Rica (Andam *et al.*, 2010; Cavatassi, Davis, & Lipper, 2004) and used by the Mexican government in the analyses of the PROGRESA program (cited by Cavatassi *et al.* (2004).

9. The health component of the Bolivian NBI is not fully comparable between censuses as the questions have changed. However the change has been accepted and approved by the Bolivian government. The full methodology can be found at <http://www.ine.gob.bo/pdf/Metodologias2004/NBI.doc>.

10. 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 requires 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 is 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 (three of which are considered protected) for which NBI was not calculated in 1992.

11. 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 than ATE.

12. 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.

13. The reduced sample due to missing 1992 NBI observations does not appear to be driving the results. We designate the PI as the outcome of interest using the reduced sample and results do not change significantly.

14. 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 matched unprotected 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, Imai, King, and Stuart (2007) because the second stage regression helps to eliminate any residual differences across treatment arms that remain after matching.

15. Another way to temperate this effect size measurement would be that approximately 60% of the protected municipalities had lower PI and NBI measures than the matched controls.

16. 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.

REFERENCES

- Abadie, A., Drukker, D., Herr, J., & Imbens, G. (2004). Implementing matching estimators for average treatment effects in stata. *Stata Journal*, 4, 290–311.
- Abadie, A., & Imbens, G. (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1), 235–267.
- Adams, W., Aveling, R., Brockington, D., Dickson, B., Elliott, J., Hutton, J., et al. (2004). Biodiversity conservation and the eradication of poverty. *Science*, 306(5699), 1146–1149.
- Andam, K., Ferraro, P., Sims, K., Healy, A., & Holland, M. (2010). Protected areas reduced poverty in Costa Rica and Thailand. *Proceedings of the National Academy of Sciences*, 107(22), 9996.
- Andam, K. S., Ferraro, P. J., Pfaff, A., Sanchez-Azofeifa, G., & Robalino, J. A. (2008). Measuring the effectiveness of protected area networks in reducing deforestation. *Proceedings of the National Academy of Sciences*, 105(42), 16089–16094.
- Angrist, J., & Pischke, J. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Bruner, A., Gullison, R., Rice, R., & Da Fonseca, G. (2001). Effectiveness of parks in protecting tropical biodiversity. *Science*, 291(5501), 125.
- Canavire-Bacarreza, G., & Mariscal, M. (2010). Macroeconomic policy, external shocks and social protection system effects on poverty, inequality, and social vulnerability in Bolivia: A CGE analysis. UDAPE Published Report.
- Cavatassi, R., Davis, B., & Lipper, L. (2004). Estimating poverty over time and space: Construction of a time-variant poverty index for Costa Rica. Technical Report. ESA Working Paper.
- Coad, L., Campbell, A., Miles, L., & Humphries, K. (2008). *The costs and benefits of protected areas for local livelihoods: A review of the current literature*. UNEP.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Lawrence Erlbaum.
- Dehejia, R., & Wahba, S. (2002). Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and Statistics*, 84(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(1), 109–136.
- Ferraro, P. J. (2008). Protected areas and human well-being in the tropics: A strategic dialogue. Conference Paper.
- Ferraro, P. J., & Hanauer, M. M. (2011). Protecting ecosystems and alleviating poverty with parks and reserves: 'Win-win' or tradeoffs?. *Environmental and Resource Economics*, 48(2), 269.
- Ferraro, P. J., & Hanauer, M. M. (2012). Causal mechanisms of protected area impacts. Working Paper.
- Ferraro, P. J., Hanauer, M. M., & Sims, K. R. (2011). Conditions associated with protected area success in conservation and poverty reduction. *Proceedings of the National Academy of Sciences*, 108(34), 13913–13918.
- Filmer, D., & Pritchett, L. (2001). Estimating wealth effects without expenditure data-or tears: An application to educational enrollments in states of India. *Demography*, 38(1), 115–132.
- Fleck, L. C., Amend, M., Painter, L., & Reid, J. (2006). Regional economic benefits from condensation: The case of Madidi. Conservation Strategy Fund Working Paper.
- Ho, D., Imai, K., King, G., & Stuart, E. (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(396), 945–960.
- Imbens, G., & Wooldridge, J. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1), 5–86.
- Joppa, L., & Pfaff, A. (2009). High and far: Biases in the location of protected areas. *PLOS One*, 4(12).
- Lewis, D., Hunt, G., & Plantinga, A. (2002). Public conservation land and employment growth in the northern forest region. *Land Economics*, 78(2), 245.
- Lewis, D., Hunt, G., & Plantinga, A. (2003). Does public lands policy affect local wage growth?. *Growth and Change*, 34(1), 64–86.
- MEA (2005). *Ecosystems and human well-being. Millennium ecosystem assessment*. Island Press.
- Meyer, B. (1995). Natural and quasi-experiments in economics. *Journal of Business and Economic Statistics*, 13(2), 151–161.
- Pfaff, A., Robalino, J., Sanchez-Azofeifa, G., Andam, K., & Ferraro, P. (2009). Park location affects forest protection: Land characteristics cause differences in park impacts across Costa Rica. *The B.E. Journal of Economic Analysis and Policy*, 9(2), 5.
- Robalino, J. (2007). Land conservation policies and income distribution: Who bears the burden of our environmental efforts?. *Environment and Development Economics*, 12(04), 521–533.
- Robalino, J., & Villalobos-Fiatt, L. (2011). Conservation policies and labor markets: Unraveling the effects of national parks on local wages in Costa Rica. IDEAS Discussion Papers.
- Rosenbaum, P. (2002). *Observational studies*. Springer.
- Sachs, J., Baillie, J., Sutherland, W., Armsworth, P., Ash, N., Beddington, J., et al. (2009). Biodiversity conservation and the millennium development goals. *Science*, 325(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(2), 1–51.
- Sernap, B. (2007). *Informe sobre el sistema de areas protegidas. SERNAP-Servicio Nacional de Areas Protegidas La Paz*. SERNAP.
- Sims, K. R. (2010). Conservation and development: Evidence from Thai protected areas. *Journal of Environmental Economics and Management*, 60(2), 94–114.

- Watts, R., Compton, R., McCammon, J., Rich, C., Wright, S., Owens, T., *et al.* (2007). Roadless space of the conterminous United States. *Science*, 316(5825), 736.
- WDPA. (2011). World database on protected areas: 2010 annual release.
- Wilkie, D., Morelli, G., Demmer, J., Starkey, M., Telfer, P., & Steil, M. (2006). Parks and people: Assessing the human welfare effects of establishing protected areas for biodiversity conservation. *Conservation Biology*, 20(1), 247.
- Yáñez, E. (2006). Existe un vínculo entre las áreas protegidas y la pobreza? evidencia de tres áreas protegidas de Bolivia. The Nature Conservancy Mimeo.

APPENDIX A

A.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 4th World Congress on National Parks and Protected areas (Andam *et al.*, 2010) and previous studies (Andam *et al.*, 2010; Ferraro & 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 8 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 9–12 provide full results for all protection thresholds). Table 8 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 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 8. 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 8–12 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.

Table 8. *ATT Estimates from primary GenMatch specification by protection threshold*

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

(Abadie–Imbens Heteroskedasticity Robust Standard Errors).
[Observation].

Table 9. *Results from primary and ancillary analyses, treatment threshold = 5%*

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Naïve difference in means	-1.23 [63]	-0.4946 [261]	-0.731 {0.0215}	77.12 [60]	85.89 [251]	-8.77 {0.004}
Regression dropping marginal	NA [63]	NA [46]	-0.42 (0.103)	NA [60]	NA [42]	-2.04 (1.55)
Post-match frequency weighted regression	-1.23 [63]	-0.761 [63]	-0.465 (0.136)	77.12 [60]	81.78 [60]	-4.66 (3.394)
Genetic matching	-0.986 [56]	-0.516 [56]	-0.471 (0.14)	79.72 [54]	82.73 [54]	-3.01 (2.09)
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}.

Table 10. Results from primary and ancillary analyses, treatment threshold = 20%. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each Treatment column is calculated as $Y(T = 1) - Y(T = 0)$

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Naïve difference in means	-1.21 [42]	-0.551 [282]	-0.663 {0.018}	76.2 [39]	85.3 [272]	-9.13 {0.013}
Regression dropping marginal	NA [42]	NA [252]	-0.527 (0.113)	NA [39]	NA [242]	-7.23 (1.39)
Post-match frequency weighted regression	NA [42]	NA [31]	-0.47 (0.116)	NA [39]	NA [29]	-2.83 (1.95)
Genetic matching	-0.996 [36]	-0.475 [36]	-0.52 (0.155)	78.95 [35]	82.48 [35]	-3.53 (3.1)
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}.

Table 11. Results from primary and ancillary analyses, treatment threshold = 30%. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each Treatment column is calculated as $Y(T = 1) - Y(T = 0)$

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Naïve difference in means	-1.16 [42]	-0.567 [282]	-0.595 {0.108}	76.6 [36]	85.2 [275]	-8.551 {0.029}
Regression dropping marginal	NA [38]	NA [252]	-0.523 (0.119)	NA [36]	NA [242]	-7.42 (1.44)
Post-match frequency weighted regression	NA [38]	NA [30]	-0.45 (0.124)	NA [36]	NA [28]	-2.91 (2.08)
Genetic matching	-1.162 [38]	-0.731 [38]	-0.431 (0.154)	76.64 [36]	82.92 [36]	-6.282 (5.1)
Genetic matching, calipers = 1sd	-0.908 [32]	-0.506 [32]	-0.402 (0.157)	79.71 [31]	83.59 [31]	-3.872 (3.1)

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

Table 12. Results from primary and ancillary analyses, treatment threshold = 50%. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each treatment column is calculated as $Y(T = 1) - Y(T = 0)$

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Naïve difference in means	-1.67 [18]	-0.577 [306]	-1.1 {0.09}	70.14 [17]	85.01 [294]	-14.87 {0.04}
Regression dropping marginal	NA [18]	NA [252]	-0.684 (0.162)	NA [17]	NA [242]	-11.3 (1.95)
Post-match frequency weighted regression	NA [18]	NA [15]	-0.627 (0.218)	NA [17]	NA [15]	-8.52 (3.3)
Genetic matching	-1.67 [18]	-1.223 [18]	-0.454 (0.272)	70.14 [17]	76.49 [17]	-6.35 (3.52)
Genetic matching, calipers = 1sd	-1.23 [15]	-0.451 [15]	-0.782 (0.245)	75.9 [14]	82.22 [14]	-6.33 (6.95)

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

A.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. 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 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. This group is observably similar to the original protected group (on average; see Table 5) 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.

Table 13. *Results from placebo matching test*

Method	Poverty index		
	$Y(T = 1)$	$Y(T = 0)$	Treatment
Genetic matching full matched sample	−0.797 [56]	−0.84 [56]	0.044 (0.17)
Genetic matching dropped repeat	−0.915 [41]	−0.7189 [41]	−0.196 (0.156)

[Number of observations].
(Standard errors).

Table 14. *Balance results for placebo matching analysis—full sample*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%Improve mean diff.
Poverty index 1992	Unmatched	0.319	0.834	−0.515	0.157	0.502	
	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 major city	Unmatched	136600	102500	34100	0.205	33260	
	Matched	136600	137500	−868.1	0.004	13100	97.45
Average elevation	Unmatched	1794	2860	−1066	0.375	1052	
	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 volume 1992	Unmatched	2.526E+14	6.101E+13	1.916E+14	0.259	1.860E+14	
	Matched	2.526E+14	1.269E+14	1.257E+14	0.162	1.547E+14	34.38

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

Table 15. *Balance results for placebo matching analysis—unique sample*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%Improve mean diff.
Poverty index 1992	Unmatched	0.282	0.834	−0.551	0.170	0.537	
	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 major city	Unmatched	136100	102500	33560	0.201	32230	
	Matched	136100	132900	3230	0.017	10430	90.38
Average elevation	Unmatched	1956	2860	−904.1	0.313	887.6	
	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

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

The results in Table 13, 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 14 and 15 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.

A.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 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, thus leading to positive impact estimates. 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 increase average poverty measured in surrounding municipalities. The former effect 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 the establishment of protected areas did not make surrounding populations better-off, it just compelled those that they made the worst-off, to emigrate. In conjunction, such a scenario would lead to erroneous estimates of poverty alleviation associated with the establishment of protected areas. To test for this potential false identification we first assume that if the poor are negatively affected by protected areas, they will migrate to the nearest unaffected municipalities.

Our analysis to test local migration effects is thus framed as a spillover analysis. Using GIS we select all the municipalities that neighbor (contiguous to) municipalities with at least 10% of their area occupied by a protected area (see Figure A1). If protected areas caused poor populations to migrate to surrounding municipalities, 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 that are not “marginally protected” 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 words, there should be no difference in poverty outcomes between contiguous (to protected municipalities) and matched unprotected municipalities.

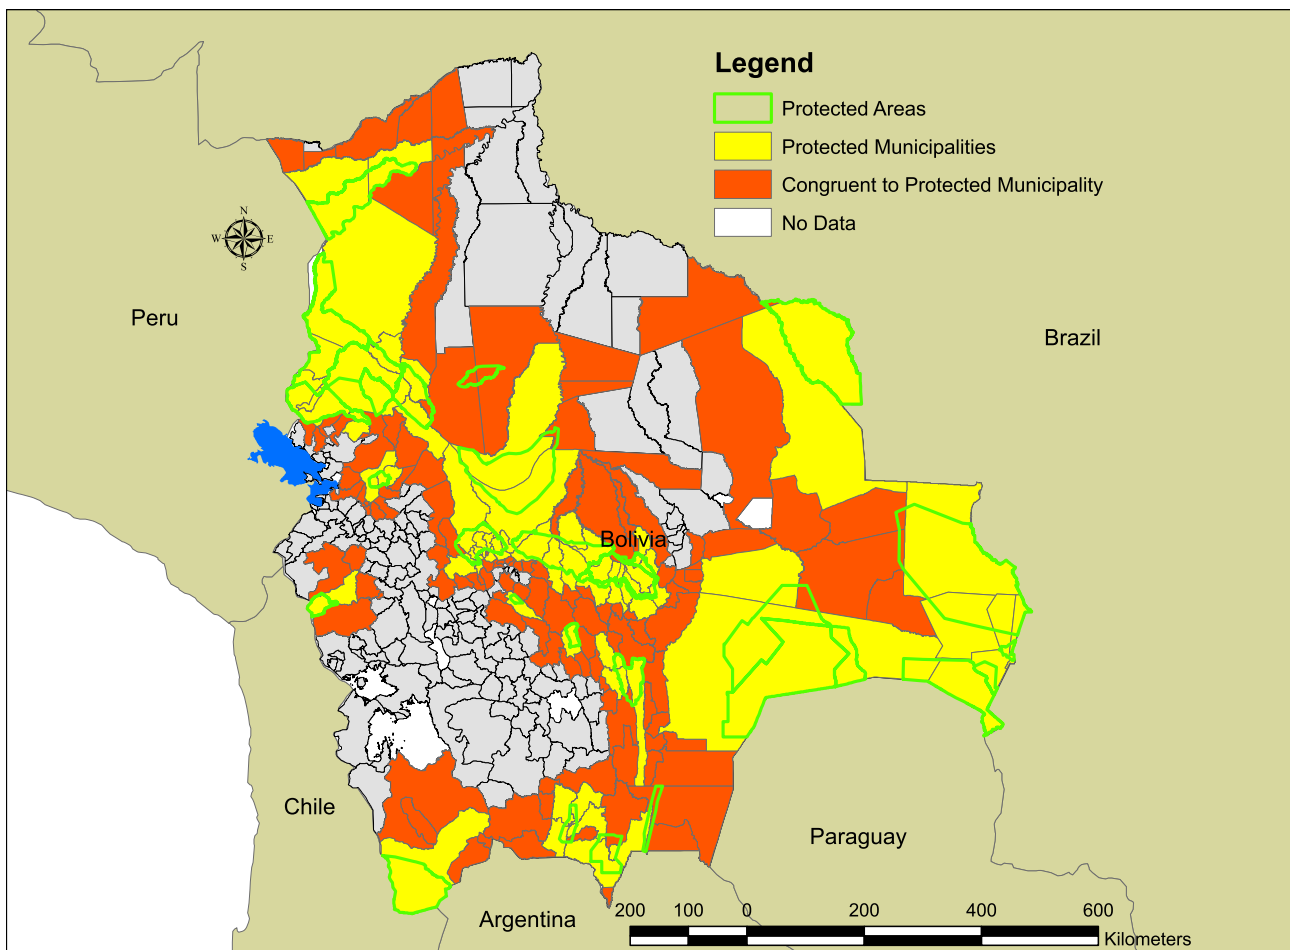


Figure A1. Map of municipalities contiguous (orange) to a municipality with at least 10% area occupied by a protected area (yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Figure A2. Histogram of propensity score distributions according to *PI* for full and matched data sets. Red bars indicate frequency of unprotected units and blue bars indicate frequency of protected units (purple represent areas where bars overlap). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 16. Spillover analyses, municipalities contiguous to protected municipalities considered treated. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each treatment column is calculated as $Y(T = 1) - Y(T = 0)$

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	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 17. Balance results for contiguous spillover analysis—*PI*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%Improve mean diff.
Poverty index 1992	Unmatched	0.678	0.787	-0.109	0.034	0.323	
	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 major city	Unmatched	109,000	107,400	1,569	0.009	10,650	
	Matched	109,000	110,500	-1,567	0.009	9,140	0.06
Average elevation	Unmatched	2,124	3,095	-970.3	0.342	959.5	
	Matched	2,124	2,109	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

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

We run our primary genetic matching and regression specifications on the spillover data, the results of which can be

found in Table 16 (see Tables 17 and 18 for balance results). We find no evidence of negative spillover effects from pro-

Table 18. *Balance results for contiguous spillover analysis—NBI*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%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 major city	Unmatched	109,200	108,200	1,048	0.006	10,110	
	Matched	109,200	107,300	1,894	0.011	7,354	-80.79
Average elevation	Unmatched	2,158	3,135	-977	0.348	965.8	
	Matched	2,158	2,137	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

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile–quantile plots.

tected municipalities into contiguous municipalities. 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 contiguous 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 contiguous 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 final piece of evidence that indicates our results are unlikely driven by complex migration patterns comes from the preceding placebo analysis. Aside from localized (contiguous) 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.

A.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 of 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 19 for results and Tables 20 and 21 for balance results).

A.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. If matching satisfies CIA then similar logic, and inference, can be applied to matched pairs.

Table 19. *Results from primary and ancillary analyses, protected areas established in 1990s. For each method $Y(T = 1)$ and $Y(T = 0)$ represent the average observed and counterfactual measures of poverty for protected municipalities, respectively. Each treatment column is calculated as $Y(T = 1) - Y(T = 0)$*

Method	Poverty index			NBI		
	$Y(T = 1)$	$Y(T = 0)$	Treatment	$Y(T = 1)$	$Y(T = 0)$	Treatment
Regression dropping marginal	NA	NA	-0.599	NA	NA	-4.23
	[32]	[252]	(0.128)	[30]	[242]	(1.48)
Post-match frequency weighted regression	NA	NA	-0.54	NA	NA	-1.87
	[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 20. *Balance results for analysis using protected areas established in 1990s—PI*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%Improve mean diff.
Poverty index 1992	Unmatched	0.552	0.744	-0.192	0.056	0.329	
	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 major city	Unmatched	175,400	108,000	67,370	0.314	65,140	
	Matched	175,400	145,900	29,480	0.123	41,660	56.24
Average elevation	Unmatched	1,649	2,713	-1,065.000	0.373	1,057	
	Matched	1,649	1,584	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

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile-quantile plots.

Table 21. *Balance results for analysis using protected areas established in 1990s—NBI*

Covariate	Status	Mean prot.	Mean unprot.	Diff. in means	Norm. diff. ^a	Mean eQQ diff. ^b	%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 major city	Unmatched	165,600	108,600	57,030	0.271	54,590	
	Matched	165,600	139,400	26,230	0.108	34,620	0.540
Average elevation	Unmatched	1,731	2,751	-1,021	0.360	1,013	
	Matched	1,731	1,746	-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

^a Normalized difference in means is the difference in means divided by the square root of the sum of the squared standard deviations of the treated and untreated covariate samples.

^b Mean eQQ difference is the mean of the raw differences in the empirical quantile-quantile plots.

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 is equal and, therefore, outcomes within matched pairs are considered exchangeable. Now suppose that there is some unobserved covariate, u , that is uncorrelated with X , but correlated with outcome 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, and; (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 invalidate our statistical findings. In Rosenbaum's model the probability of assignment to treatment π_j for unit j can be expressed in terms of odds as

Table 22. Regression results from primary specifications

Covariate/outcome	Standard		Post-match weighted	
	PI	NBI	PI	NBI
(Intercept)	-2.01*** (0.157)	-42.3*** (3.41)	-1.42*** (0.288)	-39.4 (6.37)
Protected	-0.535*** (0.099)	-5.62*** (1.2)	-0.494*** (0.106)	-2.63 (1.7)
Baseline poverty	0.896*** (0.017)	1.26*** (0.0361)	0.927*** (0.0235)	1.31*** (0.066)
Percent forest 1991	0.574** (0.192)	6.22*** (2.34)	-0.0298 (0.302)	-2.73 (4.75)
Distance to major city	2.77E-06*** (4.88E-08)	-2.47E-05*** (5.98E-06)	2.22E-06*** (6.33E-07)	1.12E-05 (9.69E-06)
Average elevation	1.95E-04*** (4.39E-05)	-0.0035*** (0.00054)	-4.51E-05 (9.09E-05)	6.70E-04 (0.0014)
Average slope	-0.003 (0.0027)	0.027 (0.032)	0.006 (0.005)	0.071 (0.074)
Roadless volume	2.67E-17 (9.26E-17)	1.96E-15* (1.11E-15)	-7.78E-17 (8.29E-17)	1.58E-15 (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. (Standard Errors).

*Significance at the 0.1 level.

**Significance at the 0.05 level.

***Significance at the 0.01 level.

$\frac{\pi_j}{(1-\pi_j)}$. 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

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

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

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

which states that the odds of treatment are an unknown function of x plus an unknown parametrization of u (see Rosenbaum (2002) for a discussion of the restriction on u). The odds ratio can therefore be rewritten as

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

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.