

11 Human choices and policies' impacts on ecosystem services

Improving evaluations of payment and park effects on conservation and carbon

Alexander Pfaff and Juan Robalino

Introduction

Conservation policies are receiving increased attention in light of the potential, under new climate change policies, for rewarding avoided deforestation. Deforestation could be avoided in many ways, from alterations of development policies through many forms of conservation, for example, the creation of protected areas. In particular, the potential of payments for environmental services has received significant attention of late (see Chomitz et al., 1999; Echavarria et al., 2004; Frank and Muller, 2003; Miranda et al., 2003; Rojas and Aylward, 2003; Rosales, 2003; Smith, 1995; Szentandrasi et al., 1995; Tikka, 2003) and it is a focus of this volume and chapter.

Neither payments nor more common conservation policies such as forest reserves are typically subjected to rigorous impact evaluation. Ferraro and Pattanayak (2006) call for the empirical evaluation of conservation actions, while noting that forest impact is complicated to measure because 'avoided deforestation' involves a constructed counterfactual – an estimate of the deforestation that would have occurred had the forest not received policy protection. Andam et al. (2008) is the first evaluation of protected area impact, as far as we are aware, to construct such an estimate in light of the non-random distribution of protection that can easily arise from the choices that determine policy location and programme enrolment. Generally, for evaluation, it is crucial to understand the effects of both agencies' and landowners' choices.

Correct evaluation of policy impact is crucial for avoided deforestation. Consider the impacts on the environment. With incorrect evaluation, credits for avoided deforestation can be given for policy that does not avoid deforestation. If so, a buying entity's rise in emissions after buying the credit will actually be increasing total greenhouse gas emissions. And, on the other side of the errors possible due to imperfect evaluation, if we do not credit all those who have, in fact, avoided deforestation, the incentive to generate further avoidance falls. In the likely setting of learning by doing, i.e. global trial and error, if the impacts of all reasonable early efforts to avoid deforestation cannot be distinguished then the costs of deforestation that is actually avoided will be much higher than expected; also, recall, emissions could increase.

This chapter conveys why human choices complicate correct evaluations of impacts. Unobservable land choices, choices affecting policy location and interactions among choices complicate both ex post impact evaluation and ex ante policy planning. Based on application of proper methods to Costa Rica, we then suggest how these hurdles can best be addressed.

We provide examples of: how a best practice deforestation baseline rightly conveys the constraints on the impact the pioneering Costa Rican eco-payments programme could have; why it may be critical to have different baselines for different locations to correctly infer the impacts of Costa Rican protected areas; and how choices by conservation agencies and landowners can determine the bias within heretofore typical approaches to impact evaluation.

Finally, focusing on ecosystem services payments in particular, we discuss the effect of scale of the policy. International payments for carbon based on national baselines may face a lower 'baseline hurdle' than that facing a domestic agency making payments to landowners, since baseline errors at site level cancel out if baselines are right on average. However, when nations act to earn global payments, the domestic issues reappear. Having seen how common these issues are, we suggest their implications for conservation planning in the final section.

Choice-based hurdles for impact evaluation

Unobservable choices: the need to guess at the baseline without any policy

Any number of public or environmental economics textbooks convey the general idea that conservation policies could help to achieve the socially efficient amount of conservation. Should the benefits of conservation accrue to society as a whole, and not just to those who make costly choices to provide habitat, then private choices may provide too little habitat.

There are many ways in which public actions might increase habitat. Altering planned development policy is one, such as moving a planned road to a more benign location (perhaps with global funding of incremental costs). If we were sure of what the less benign road would have caused in terms of degradation, we could confidently estimate the impact of this policy-altering project by comparing the observed actual degradation to what would have occurred.

Yet the baseline estimate of the less benign impacts would likely be imperfect. In that same light, consider the taxation of deforestation. It may well reduce clearing but how would we know? We observe how much deforestation occurs with the tax. But how much clearing would have occurred without a tax? We will never observe that by making the calculation of the impact of tax a challenge. A comparison with similar but taxless locations could help.

The same ideas apply to protected areas. If we are sure they would have been cleared without protection, then we know how much clearing protection avoids. Yet if protection is implemented we will never, in fact, be sure what would have happened without protection. Thus, again, for doing impact evaluation we are

limited to comparing the outcomes in places with implemented policies to those in places without. This can lead to errors in evaluation.

For ecosystem services payments the same reasoning applies. Payments go to land that starts and stays in forest. If we knew it would be cleared without payments, then we could know payment impact. But we do not know what would have occurred without the policy. Thus, most programmes will pay anybody who enrolls with forested land and maintains forest. However, that opens the door to paying people who were not going to clear forest anyway.

Such payments that do not make a difference arise due to missing information. We do not know the baseline, ‘no policy’ choices for places that are paid. If we did, we could write contracts to pay only when land use changed relative to a baseline, for example, land that was going to be cleared otherwise stayed in forest. That would guarantee that the policy has impact. Our discussion of Figure 11.1 shows how to get that information completely wrong.

Figure 11.1, which depicts landowners’ choice of land use in order to maximize returns, provides a useful framework for communicating constraints both on payments’ impacts and on impact estimation. Land is ordered, along the horizontal axis, according to relative profitability of clearing (clearing profit minus forest profit), with clearing profit more favoured to the right. When relative clearing profit is positive, the land will be deforested; with no payments, the forest remains within $[0, x^N]$ while the rest of the forest, i.e. to the right of x^N , will be cleared.

If environmental services payments of P per forested unit compete with non-forest land uses, landowners sign up for payments in $[0, x^P]$. A crucial point is that not all who sign up would modify their behaviour as a result of the payment. Those within the interval $[0, x^N]$ do not clear their forest even without payments. In contrast, the parcels in the interval $[x^N, x^P]$ would be deforested in the absence of payments but would be kept in forest if P is being paid.




Figure 11.1 Baselines are what people would have chosen without the policy

Source: Authors

Thus, the impact of a programme of such payments impact depends critically on the fraction of total enrolled land coming from the interval $[x^N, x^P]$, a fraction we denote by α . If α equals 1, i.e. only land from $[x^N, x^P]$ is enrolled, then all payments prevent deforestation. But if α equals 0, i.e. only land from $[0, x^N]$ is enrolled, then the payments have no impact.

We emphasize that if all of those would benefit (within $[0, x^P]$) in fact apply,¹ which parcels are enrolled affects not only impact but also the accuracy of simple impact estimates. If $\alpha = 1$, i.e. targeting is 'good' in that all of $[x^N, x^P]$ but none of $[0, x^N]$ is enrolled, then forest locations outside the programme are $[0, x^N]$. Those are not similar to the enrolled parcels. None will be cleared, although all of the enrolled would have been. If we estimate what would have happened on enrolled parcels using the non-enrolled parcels, we get the baseline completely wrong. We would estimate zero impact from payments when, in fact, each payment mattered.

Relevant human choices: agencies, landowners and bias

In the preceding example, not only was a baseline based on non-enrolled parcels wrong, it was wrong in a particular way and that bias resulted from the agency's excellent targeting. If the agency is well informed and able to exclude applicants, it can enrol only the applicants who would have cleared without a payment. As noted, that is good for true payment impact.

However, when as analysts we are estimating impact after the fact, we do not observe the 'no policy' choices in locations with payments. We are constrained to compare the lack of clearing in enrolled places with the clearing in non-enrolled places. If the agency chose in a very sensible but clearly non-random way who to enrol, such a comparison yields an impact estimate that is biased. With excellent agency targeting, the estimate will be too low (as in the example, it could be zero). With poor agency targeting, the estimate will instead be too high. Landowner choices can bias estimates too. If the agency uses first-come-first-served basis, i.e. does not target, which landowners might get into the programme? It could be that those who are sure they will lose money in agriculture would be most likely to apply. In the extreme, if only parcels that would not have been cleared anyway are enrolled, then the actual impact of payments is zero. Any non-zero impact estimate is in fact an overestimate of the true impact.

Generally if the parcels treated by a policy (whether those receiving payments or those in protected areas or those near new roads) is not a random subset of all parcels, then policy impact estimates may be biased. As we have seen, this bias could be in either direction. What determines the bias is the net impact, on non-random enrolment, of all the actors' choices. Put another way, we are constrained to compare policy locations only with no policy locations but choices by agencies and landowners can imply that these are not similar groups. The method for better impact estimation we now demonstrate tries to create similar subsets.

Spatial and temporal choice interactions: various spillovers*Spatial interactions*

Human choice also implies that policy can have impacts on locations without policies. One implication is that in comparing policy with no policy locations to estimate impacts, we need to be careful not to include the no policy locations onto which policy impacts spillover. Even if avoiding that estimation pitfall, however, we may want to consider such ‘spillovers’. For instance, in cost–benefit analysis of a policy, likely all of its impacts should be included.

For example, when a new protected area is created, those who were using or planning to clear the forest for private reasons may relocate themselves and/or their planned activities to nearby forest areas that would not otherwise have been cleared. Thus no policy locations, i.e. parcels outside of the protected area, are cleared more. The total net reduction of clearing due to the protected area is lower than one would think if looking only within the boundaries of the protected area. This idea is referred to as ‘leakage’ in various parts of this volume (see Murray, Chapter 9 in this volume).

Further, recall that we will estimate the impact within the protected area’s boundaries by comparing deforestation there with outcomes outside. If the latter includes locations near the protected area that are cleared because of the spatial spillover then these will bias the estimate of protection’s impact upwards. Protection appears to have a greater deforestation-retarding influence within its own boundaries because it raises clearing nearby.² This could happen for payments, for example, a landowner can accept payments for a forested parcel and instead clear others.

Such a bias can also go in the other direction. A policy could lower deforestation rates in no policy locations. It is claimed, for instance, that payments to some parcels may create awareness of the value of nature and/or the fact that one can receive payments. Alternatively, private or public choices to conserve forest could lead neighbours to band together to do the same. Pointing in that same direction, in terms of bias, a choice to clear and to produce could lead neighbours to produce (see, for instance, Robalino and Pfaff, 2008 for related empirical analyses).

Temporal interactions

Human response to a policy may also occur over time for any given location in space. Above we note that policy in location A can affect contemporaneous behaviours at location B or many such locations. Here we note that the policy could also affect behaviours in location A in the periods before or after the policy is implemented. This, too, is a form of policy spillover.

An example is pre-emptive clearing of species habitat based on expected public action. For instance, when the US Environmental Protection Agency (EPA) moves a species upwards in the hierarchy of protection under the Endangered Species Act, owners of land where that species resides may well revise upwards

their expectations of restrictions on their land-use choices. In response, they may reduce that species' habitat on their land immediately in order to extract their private value from that land before potential public action takes place (see related discussion and empirical analysis in, e.g., Ferraro et al., 2007; List et al., 2006). Then a policy action that was aiming to conserve habitat has instead hastened its degradation.

Private land-use choice in expectation of policy is also observed in forest frontiers. If a new road plan is announced, or even simply leaked, private land speculators acquire land in its path to realize gains. An example that in a way blends these other examples concerns a canal announced in the state of Ceará in northeast Brazil that was intended to create new options for allocating scarce freshwater. Not only did private landowners purchase land in the path, ahead of the canal's arrival, but their investments also increased the total demand for water.³

Not surprisingly, policy can also change behaviour after the policy is implemented. For instance, a road's arrival will often lead to follow-on investments, including new roads.⁴ Thus, for instance, if a conservation policy succeeds in relocating a new road around a forest, its impact in reducing deforestation may be much greater over time than is observed early on.

Choice-adjusted impact estimation

Matching analysis

What is 'similar'?

As noted, we are unable to observe the 'without-policy' outcomes for policy locations. Thus, we estimate policy impact by comparing the outcomes in untreated, i.e. without policy, locations with outcomes in locations that are treated (are protected or receive payments). Matching is an effort to compare treated locations with the most 'similar' untreated parcels (see, for example, Rosenbaum and Rubin, 1983; Rubin, 1980). However, what is 'similarity'?

To highlight the importance of this definition, we consider two matching estimators. The first is a nearest neighbour propensity score-matching estimator (Hill et al., 2003; Rosenbaum and Rubin, 1983), using a fixed number of matched observations per treated observation. Such estimators define similarity based on the estimated probabilities of being treated, which are generated by a first-stage regression that attempts to explain which locations are treated using observed characteristics of locations (such as slope, soil quality, distance to markets).

Another approach is nearest neighbour covariate matching using an inverse weighting matrix to account for the difference in the scale of the covariates (Abadie and Imbens, 2006a). Again one might employ a fixed number of best matches per treated observation. Covariate matching estimators define similarity without a first-stage regression but rather using the simple distances, in the space of the matching covariates, between the treated and matched. Thus, for any treated location, say, a parcel in a protected area, we search for the four parcels outside

protected areas whose observed characteristics (such as slope, soil quality, distance to markets) are most similar to the observed characteristics of the treated location in question.

These two approaches define similarity quite differently and, as a result, will compare the treated locations to different subsets of the untreated locations, although in both cases that subset of the untreated is made up of those similar to the treated. The computation of standard errors, to indicate when differences are statistically significant, also differs across these and other approaches to matching (see Abadie and Imbens, 2006b).⁵

Whichever of these approaches is chosen, the commonality is the basic method. Thus, the chosen number of most similar untreated points for each treated location is identified and then the outcomes for those untreated points are compared to the outcomes for the treated. For instance, if average deforestation in protected areas is zero, then the average deforestation in the matched untreated locations can be used as the estimate of the impact of protection. If the treatment is, in fact, non-randomly allocated over space, this matched impact estimate often will differ from the typical impact estimate based on all of the untreated locations. With the matched subset of untreated locations in hand, another option is to run a regression using only those untreated plus the treated locations including data on the observed characteristics.

Match quality

The point of matching is to use a subset of the untreated, i.e. no policy, parcels which are more like the treated parcels than is the entire set of untreated locations. The extent to which using a subset improves similarity needs to be checked and demonstrated. The typical first approach for doing so is a balancing test, in which the means of the characteristics of the matched subset of untreated parcels are tested for significant difference from the means for the treated parcels. One may check whether the full set of untreated parcels is different from the treated and then check whether the matched untreated locations more closely approximate the treated. Even if closer, however, the goal is for the matched points to be indistinguishable.

Yet even if a balancing test is passed after all the matching choices described earlier, matching controls solely for the impacts of the group differences in all the *observable* factors. On recognizing this, many analysts ask whether, and if so, why, typical regression analysis for all the available observations is inferior to matching analyses in this policy setting. For comparison, consider the matching regression with observed covariates but using just treated locations plus the matched untreated observations instead of typical use of all observations – those are very similar regressions differing in terms of which untreated points are included.

To the extent that the observed characteristics of locations (slope, soil, others) are not in fact similar in the treated and untreated groups, any regression using all of the observations applies the information at hand to control for the impact of the differences in characteristics. However, when group characteristics differ, the

burden on the specification is considerable.⁶ If the specification of the regression is not correct, the control for characteristics is wrong.

Matching is intended precisely to reduce that burden by comparing 'apples to apples' or using untreated and treated locations with similar characteristics. Yet this problem will plague matching, too, when match quality is not good. That can motivate the use of a subset of the treated observations, too, in particular dropping all of the treated points with poor matches.

Crump et al. (2006) address a lack of covariate overlap between untreated and treated, noting that many common estimators become sensitive to the choice of specification (much as Cochran noted and following related prior work including Heckman et al., 1997, 1998). Crump et al. (2006) characterize optimal subsamples for which treatment effects can be estimated most precisely. In the following, we do not use optimal subsamples but do emphasize this by examining robustness to dropping high probability of treatment treated observations for which we do not find good untreated matches (Cochran and Rubin, 1973).

Application I: Costa Rican ecopayments

Unobserved choices: low deforestation baseline

We do not observe what would have happened without payments in parcels receiving payments. However, as in Sánchez et al. (2007), we can estimate this by looking at locations not receiving payments. From Figure 11.1, this is not sensible if either all land that would not be cleared without payments or all land that would be cleared without payments is included in the programme. However, the Costa Rican PSA programme was quite small so both are unlikely.

Anecdotally (de Camino-Velozo et al., 2000; Hartshorn et al., 2005), starting in 1997 Law 7575 on forestry greatly slowed clearing. Consistent with this observation, simple measures of forest find that during the first 3 years of the PSA payment programme (which started in 1997) the annual deforestation rate outside the programme areas was 0.21 per cent, i.e. one-fifth of 1 per cent. Thus, very little of the land enrolled in PSA was actually being protected from deforestation. That may reflect significant impacts of all previous conservation policies (1997 law, protection, ecotourism and more). If other policies more or less put a halt to forest clearing, then there may not have been very much deforestation for payments to prevent.

Landowner choices: lower deforestation impact

We believe that the 1997–2000 payments were not strongly targeted by the relevant agency (see also Engel et al., Chapter 12 in this volume). Thus, agency choice may not be a major issue during this time period. However, landowner choices seem likely to matter, since the PSA programme was voluntary. That is, parcels were volunteered by their owners. While this is not necessarily the dominant story, there is good reason to believe that owners may have frequently volunteered for the

forest-maintenance payments those parcels which were the poorest for agriculture, perhaps so poor that they would never have been cleared anyway.

Pfaff et al. (2007a) directly examine this possibility, a general issue for such payments and relevant for countries without Costa Rica's low deforestation baseline. First, a pixel-level regression generates the predicted probabilities of 1997–2000 deforestation. Then we examine whether enrolment leans towards areas of higher or lower deforestation. Without question the bias, if any, is towards lower pressure. That is, the land that is less likely than average to be cleared without payments is over-represented in enrolment for payments. That suggests that even the low deforestation national baseline will overestimate the PSA policy's impact.

Pfaff et al. (2007b) go beyond this bounding argument by applying matching analysis (as described previously). We compare the treated (or payment) parcels with untreated (or no payment) parcels, which, per their observable characteristics (such as slope of the land, distance from urban markets and from roads) are most similar to the treated. This will avoid or ameliorate errors in policy-versus-no-policy comparisons such as discussed earlier. In this fashion, for example, using the propensity score-matching approach, we estimate the 1997–2000 Costa Rican ecopayment impact at 0.08 per cent or less than one-tenth of 1 per cent rate of deforestation per year prevented, i.e. impact well below the average of 0.21 per cent for all untreated areas.

Robalino et al. (2007) then examine payments after 2000 given that, by all accounts, the programme increased targeting after 2000. Yet our understanding is that targeting was not of land more likely to have been cleared without payments but instead land whose conservation would most matter ecologically or, more generally, in terms of ecosystem services provision. The baseline for this period is essentially unchanged from the very low 1997–2000 clearing, although net clearing was even lower during 2000–2005 and, in fact, was negative. However, there was more deforestation (even if it was outweighed on net by reforestation). This implies more potential for deforestation prevention by payments. Preliminary results suggest in addition that post-2000 payments were not as biased towards low pressure areas. Between the shifts in baseline trends and in the allocation of payments over space, we find that 0.5 per cent, of deforestation per year, was prevented. This is still a relatively low rate but it conveys that changes in policy design and trends can affect impact.

Application II: Costa Rican protected areas

Unobserved choices: temporal shifts affect the baseline

For the same country, of course, the facts about the 1997–2000 baseline will still apply. However, protected areas raise the issue of shifts over time in forest baselines because within Costa Rica, most of the protected areas were created well before the ecopayments had begun.

Andam et al. (2008) considers the impacts of protection from the 1960s onward and find that significant amounts of deforestation were prevented (although much

less than usually assumed – more later). That is not surprising neither does it contradict the ecopayments result because the deforestation rate in Costa Rica used to be much higher than it was after 1997 and thus any effective protection (i.e., zero deforestation in protected areas, which is true in Costa Rica if not in all other countries) certainly could be expected to have had an impact.

Comparing this with Pfaff et al. (2007c) emphasizes the importance of understanding the default or baseline rate of land clearing that an effective policy prevents. Pfaff et al. (2007c) examine forest impacts for the period 1986–1997. They, too, find some significant impact but its magnitude is lower than in Andam et al. (2008) as deforestation had already started to slow. More importantly, Pfaff et al. (2007c) find quite significant variation in the baseline over space within this time period. Thus, the baseline deforestation rate and, as a result, the conservation impact of protection is shown to vary with the distance to San José, the distance to a road and slope.

Agency targeting choices: lower deforestation impact

Whereas ecopayment policy gave a role to landowners' choices, the location of areas designated for protection did not. Agencies literally locate reserves. That need not bias the set of locations but one can imagine motivations such that it would. An agency might target locations that were going to be cleared, to maximize one kind of impact. Or an agency might want protection to last for centuries and to avoid political pressure about taking private lands (an issue which arose in Costa Rica) and might then look for low value, uncontested lands.

While differing in average impact found, due to the shifts in baselines over time, both Andam et al. (2008) and Pfaff et al. (2007c) find strong statistical evidence that protection is not on a random set of locations but rather on locations less likely to be cleared than the average parcel. Thus, estimating impact by looking at all non-protected locations will overestimate impacts of protection. The Andam et al. (2008) and Pfaff et al. (2007c) estimates of the clearing prevented by protection are under half of what one finds if ignoring the non-random location of protection by comparing deforestation under protection to that on all the unprotected land.

Spatial choice interactions

Both baseline (unobserved no policy choices) and selection (landowner and agency choices) issues that affect additionality, as discussed earlier, must be considered in estimating correctly the impact of a policy in its intended location and time period. That is not true of spatial and/or temporal interactions or leakage in space or time that result from the policy.

That such leakage is outside of the initial space-and-time focus explains in part why often it is omitted from cost–benefit analyses though it represents relevant costs and benefits. For protected areas, in considering the total net policy impact one natural question is whether completely preventing land clearing within a protected area influenced the land uses outside.

In examining whether land around a protected area was deforested more or less than in the case of no protection, baseline issues apply. We do not observe deforestation next to protection in the case of no protection. We can only compare next-to-protection with not-near-protection parcels. Yet next-to-protection parcels may not be a random set. Still they might be compared with a subset of not-next-to-protection lands. Robalino et al. (2008) examines this empirically for Costa Rica, finding preliminary evidence of some deforestation spillovers, i.e. that rates of deforestation in some areas next to protection are higher than if there was no protection.

Does scale reduce choice-based hurdles?

International ‘REDD’ payments

Consider the payments for avoided deforestation being considered as part of global climate change strategy. In principle, those payments could be made, as in current domestic programmes, directly to individual landowners. However, since the transactions costs of doing so could be prohibitive, instead a country could serve as the ‘owner’. Thus a country could contract with other countries to receive payments if it captured more carbon, or if it emitted less carbon, than in an agreed baseline. Payments could then be distributed domestically.

Conceptually, baseline issues still apply. If a policy exists, then the world will never observe what land use would have occurred in the country without the payments incentives. Thus one asserts at a national level how much deforestation would have occurred without a policy. This has its challenges. Assertions will be too high for some countries but lower for others. However, working at the national level allows the many errors one could make at the parcel scale to cancel out. Being right on average across parcels works for any given country, in the sense of credits for reduced carbon emissions being valid (not ‘hot air’). While hurdles remain conceptually, the magnitude of the baseline problem may be reduced at larger scale.

Further, spatial interaction may be greatly reduced. While Sohngen and Brown (2004) show that linkages can occur through global markets – if timber is not cut and not put on the market from one country, the global price response could induce more cutting in other sites – local leakage would not be an evaluation issue at the country scale. If a conservation policy in a given country does cause local leakage, then that is taken into account in calculating global payments to the country, as it reduces the amount of additional (i.e. above baseline) forest.

Domestic problems redux

Still selection issues can arise in an international payments scheme, analogous to landowner choices in the Costa Rican payments programme. Following the logic of Montero (1999) and its application to deforestation in Kerr (2007), we must realize that the decision to participate in a global payments policy may depend on the error in a country’s assigned baseline. If it was assumed that more forest would be

present without payments than the country believes is correct, it can simply refuse to join as the cost of earning payments by having a forest level above the baseline would be simply too costly. If the assumed forest level is erroneously low, a country will rapidly join. This process means the world overpays for impact. Even if baselines were right on average for all the countries, they would be biased towards assuming too little forest for those that will join. Thus as in our earlier payments example, some countries will pay some others for unchanged behaviours (or 'hot air').

At least as important is a domestic perspective on earning global payments. As noted, international payments are earned if national outcomes (forest area, carbon) are above agreed baselines. Assuming the baseline is correct, i.e. that the baseline is indeed what would occur without policy, the national agency that will receive global payments need to do something to translate that global incentive into domestic behaviours. Additional domestic ecopayments and protection are likely candidates for doing so. That brings us back to the inefficiencies in implementation of such policies shown earlier, due to a lack of baseline knowledge at parcel level. That can greatly increase the domestic transfers required to earn the global payments.

Discussion

Microeconomic theory supports the idea that conservation policy may improve social efficiency should conservation's benefits accrue to society, i.e. not just to those who provide it at a cost. Thus a basic theory of private choices that affect others suggests that policies like public forest reserves or ecosystem services payments could have impacts that are beneficial.

However, whether such conservation policies actually have those impacts in reality is another question entirely and is not usually asked. To answer it, we do not observe what would have happened had the policy not been implemented. Further, using observations of locations without policy to guess at this may provide the wrong impact estimates because the human choices that drive the location of policies often imply their locations are non-random.

These issues create challenges for ex post estimates of policy impact. However, as we demonstrated in two applications to Costa Rica (ecopayments and protected areas), there are sensible methods for 'human choice adjusting' the empirical estimation of policy impacts.

These issues with ex post impact estimation also complicate ex ante policy planning: policymakers must guess the future of each candidate parcel without policy being applied; policy location may be driven by influential actors with various goals; policies in one place may affect outcomes in other. The only difference for planning is that one is choosing where to locate a policy, which adds prediction (in place of hindsight) to all of our earlier challenges.

What can be done? A focus on forest parcels under greater clearing pressure could in principle raise impact. First, if an ecopayments programme remains over-subscribed, in selecting parcels to admit an agency can focus on those more likely to be deforested in the absence of payments. That would essentially be a form of

targeting of high clearing threat. There may be countervailing arguments like land price but threat should at least be considered (see Engel et al. and Alix-Garcia et al., Chapters 12 and 13 in this volume, respectively). Putting it another way, to provide a given level of eco-services more costly land could, in fact, lower policy cost.

As noted in the discussion of the Costa Rican payments programme after 2000, targeting may be on dimensions of environmental service, i.e. not explicitly concerning threat. Those may not be correlated with threat. More threat targeting could raise gains for a given budget.

Programme design could also permit higher payments to areas under higher threat due to higher returns. One payment five times as high as current payments could yield a positive impact on one parcel even if current payments to five parcels yield nothing. Such adjustments would require an agency to have information about rents. Such information will be imperfect but it may be sufficient (again, see Engel et al and Alix-Garcia et al., Chapters 12 and 13 in this volume, respectively).

Such adjustments could change not only efficiency but also the distributional effects. For a fixed budget, having high payments means concentrating payments on fewer people, for example, providing nothing to four of any five people who would have enrolled in the current regime but five times as much as currently paid for the fifth person. Such a shift could, in principle, help to alleviate poverty if that one person is relatively poor. Targeting higher rent land, however, may mean that the person receiving payments has relatively high profits and is more likely to have relatively high wealth. Combining the targeting of the higher rents with fewer but higher payments can shift a programme away from any redistributive policy objective.

Acknowledgements

For financial support we thank the Tinker Foundation, Inc.; NSF's MMIA (National Science Foundation, Methods and Models for Integrated Assessment); NCEAS (NSF's National Center for Ecological Analysis and Synthesis); SSHRC (Social Science and Humanities Research Council of Canada); CERC/Earth Institute at Columbia University, FAO (U.N. Food and Agriculture Organization); IAI (Inter-American Institute for Global Change Research); LACEEP (Latin American and Caribbean Environmental Economics Program); and EfD (Environment for Development initiative).

Notes

- 1 Note that, in reality, not all of the landowners who would like to enrol (because they are in $[0, x^p]$) will in fact apply for the programme. Some landowners may not know about the programme or may face high application costs. Further, not all those who apply are guaranteed to be enrolled, if the programme lacks sufficient funds.
- 2 See Robalino (2007), which also discusses spillovers to the wages of those who work on land in the area.
- 3 Personal communication, Renzo Taddei, from field interviews with both landowners and agency officials.

- 4 See, for example, Pfaff, October 2006, LBA Brasilia conference presentation on evolving Amazon road networks.
- 5 For example, Abadie and Imbens (2006b) show that the common practice of bootstrapping standard errors is invalid with non-smooth, nearest neighbour estimators such as the propensity score-matching estimator with a fixed number of matches (versus with smoothly declining weights to less well-matched untreated locations).
- 6 Consider, for instance, the following claim from Cochran, in Rubin (1984):

Unless the regression equation holds in the region in which observations are lacking, covariance will not remove all the bias, and in practice may remove only a small part or it. Secondly, even if the regression is valid in the no man's land, the standard errors of the adjusted means become large, because the standard error formula in a covariance analysis takes account of the fact that extrapolation is being employed. Consequently the adjusted differences may become insignificant merely because the adjusted comparisons are of low precision. When the groups differ widely in x , these differences imply that the interpretation of an adjusted analysis is speculative rather than soundly based.

References

- Abadie, A. and Imbens, G.W. (2006a) 'Large sample properties of matching estimators for average treatment effects', *Econometrica*, 74(1): 235–267.
- Abadie, A. and Imbens, G.W. (2006b) *On the Failure of the Bootstrap for Matching Estimators*, NBER Technical Working Paper 325 [online]. Available at: <http://www.nber.org/papers/t0325.pdf> [accessed 11 August 2008].
- Alix-Garcia, J., de Janvry, A. and Sadoulet, E. (2009) 'The role of risk in targeting payments for environmental services', in C. Palmer and S. Engel (eds) *Avoided Deforestation: Prospects for Mitigating Climate Change*, Oxford: Routledge.
- Andam, K., Ferraro, P., Pfaff, A., Robalino, J. and Sánchez-Azofeifa, G.A. (2008) 'Measuring the effectiveness of protected area networks in reducing deforestation: a rigorous impact evaluation approach', *Proceedings of the National Academy of Sciences* (accepted).
- Chomitz, K., Brenes, E. and Constantino, L. (1999) 'Financing environmental services: the Costa Rican experience and its implications', *The Science of the Total Environment*, 240(1): 157–169.
- Cochran, W. and Rubin, D. (1973) 'Controlling bias in observational studies: a review', *Sankhya*, Series A 35: 417–446.
- Crump, R.K., Hotz, V.J., Imbens, G.W. and Mitnik, O.A. (2006) 'Moving the goalposts: addressing limited overlap in estimation of average treatment effects by changing the estimand', mimeo, Department of Economics, University of California Berkeley.
- de Camino-Velozo, R. and World Bank (2000) *Costa Rica: Forest Strategy and the Evolution of Land Use*, Washington, DC: World Bank.
- Echavarría, M., Vogel, M.J., Alban, M. and Meneses, F. (2004) *The Impacts of Payments for Watershed Services in Ecuador: Emerging Lessons from Pimampiro and Cuenca*, London: International Institute for Environment and Development (IIED).
- Engel, S., Wünscher, T. and Wunder, S. (2009) 'Increasing the efficiency of forest conservation: the case of payments for environmental services in Costa Rica', in C. Palmer and S. Engel (eds) *Avoided Deforestation: Prospects for Mitigating Climate Change*, Oxford: Routledge.
- Ferraro, P.J., McIntosh, C. and Ospina, M. (2007) 'The effectiveness of the US endangered

- species act: an econometric analysis using matching methods', *Journal of Environmental Economics and Management*, 54(3): 245–261.
- Ferraro, P.J. and Pattanayak, S.K. (2006) 'Money for nothing? A call for empirical evaluation of biodiversity conservation investments', *PLoS Biology*, 4(4): e105.
- Frank, G. and Muller, F. (2003) 'Voluntary approaches in protection of forests in Australia', *Environmental Science and Policy*, 6(3): 261–269.
- Hartshorn, G., Ferraro, P.J., Spergel, B. and Sills, E. (2005) *Evaluation of the World Bank–GEF Ecomarkets Project in Costa Rica*, Report of the Blue Ribbon Evaluation Panel [online]. Available at: http://www2.gsu.edu/~wwwceec/docs/doc%20updates/NCSU_Blue_Ribbon_Panel_Final.pdf [accessed 23 June 2008].
- Heckman, J.J., Ichimura, H. and Todd, P. (1997) 'Matching as an econometric evaluation estimator: evidence from evaluating a job training program', *Review of Economic Studies*, 64(4): 605–654.
- Heckman, J.J., Ichimura, H. and Todd, P. (1998) 'Matching as an econometric evaluation estimator', *Review of Economic Studies*, 65(2): 261–294.
- Hill, J., Waldfogel, J. and Brooks-Gunn, J. (2003) 'Sustained effects of high participation in an early intervention for low-birth-weight premature infants', *Developmental Psychology*, 39(4): 730–744.
- Kerr, S. (2007) *China Workshop Presentation on Effects of Selection with Errors*, mimeo.
- List, J., Margolis, M. and Osgood, D. (2006) *Is the Endangered Species Act Endangering Species?*, Working Paper 12777, Cambridge, MA: National Bureau of Economic Research (NBER).
- Miranda, M., Porras, I.T. and Moreno, M.L. (2003) *The Social Impacts of Payments for Environmental Services in Costa Rica: a Quantitative Field Survey and Analysis of the Virilla Watershed*, London: Environmental Economics Program, International Institute for Environment and Development (IIED).
- Montero, J.P. (1999) 'Voluntary compliance with market-based environmental policy: evidence from the U.S. acid rain program', *Journal of Political Economy*, 107(5): 998.
- Murray, B. (2009) 'Leakage from an avoided deforestation compensation policy: concepts, empirical evidence and corrective policy options', in C. Palmer and S. Engel (eds) *Avoided Deforestation: Prospects for Mitigating Climate Change*, Oxford: Routledge.
- Pfaff, A., Robalino, J. and Sánchez, A. (2007a) *Into the REDD?*, mimeo, Duke University.
- Pfaff, A., Robalino, J. and Sánchez, A. (2007b) *Payments for Environmental Services: Empirical Analysis for Costa Rica*, Duke Sanford Working Paper SAN08-05, Durham, NC: Terry Sanford Institute of Public Policy, Duke University.
- Pfaff, A., Robalino, J., Sánchez, A., Andam, K. and Ferraro, P. (2007c) *Location Affects Protection: Observable Characteristics Drive Park Impacts in Costa Rica*, Duke Sanford Working Paper SAN08-06, Durham, NC: Terry Sanford Institute of Public Policy, Duke University.
- Robalino, J. (2007) 'Land conservation policies and income distribution: who bears the burden of our environmental efforts?', *Environment and Development Economics*, 12(4): 521–533.
- Robalino, J. and Pfaff, A. (2008) *Contagious Development: Neighbor Interactions in Deforestation*, mimeo, Centro Agronómico Tropical de Investigación y Enseñanza (CATIE).
- Robalino, J., Pfaff, A., Sánchez, A., Alpizar, F., Leon, C. and Rodriguez, C.M. (2007). *Deforestation Impacts of Environmental Services Payments: Costa Rica's PSA Program 2000–2005*, EFD (Environment for Development) Discussion Paper 08–24, Gothenburg: University of Gothenburg and Washington, DC: Resources for the Future (RFF).

- Robalino, J., Pfaff, A., Sánchez, A. and Villalobos, L. (2008) *Protected Area Spillovers: Leakage from Protection in Costa Rica 1986–1997*, mimeo, Centro Agronómico Tropical de Investigación y Enseñanza (CATIE).
- Rojas, M. and Aylward, B. (2003) *What Are We Learning from Experiences with Markets for Environmental Services in Costa Rica? A Review and Critique of the Literature*, London: International Institute for Environment and Development (IIED).
- Rosales, R.M.P. (2003) *Developing Pro-Poor Markets for Environmental Services in the Philippines*, London: International Institute for Environment and Development (IIED).
- Rosenbaum, P.R. and Rubin, D.B. (1983), 'The central role of the propensity score in observational studies for causal effects', *Biometrika*, 70(1): 41–55.
- Rubin, D.B. (1980) 'Bias reduction using Mahalanobis-metric matching', *Biometrics*, 36(2): 293–298.
- Rubin, D.B. (1984) 'William G. Cochran's contributions to the design, analysis, and evaluation of observational studies', in D. Rubin (ed.) *Matched Sampling for Causal Effects*, Cambridge: Cambridge University Press.
- Sánchez-Azofeifa, A., Pfaff, A., Robalino, J. and Boomhower, J. (2007) 'Costa Rican payment for environmental services program: intention, implementation and impact', *Conservation Biology*, 21(5): 1165.
- Sohngen, B. and Brown, S. (2004) 'Measuring leakage from carbon projects in open economies: a stop timber harvesting project in Bolivia as a case study', *Canadian Journal of Forest Resources*, 34: 829–839.
- Smith, R.B.W. (1995) 'The conservation reserve program as a least-cost land retirement mechanism', *American Journal of Agricultural Economics*, 77(1): 93–105.
- Szentandrasi, S., Polasky, S., Berrens, R. and Leonard, J. (1995) 'Conserving biological diversity and the conservation reserve program', *Growth and Change*, 26(3): 383–404.
- Tikka, P. (2003) 'Conservation contracts in habitat protection in southern Finland', *Environmental Science and Policy*, 6(3): 271–278.