

The Inherent Trade-Off Between the Environmental and Anti-Poverty Goals of Payments for Ecosystem Services*

Seema Jayachandran
Northwestern University

April 4, 2022

Abstract

Conservation programs in low-income countries often have dual goals of protecting the environment and reducing poverty. This article discusses the tension between these two goals in payments for ecosystem services (PES) programs. Participants who undertake a pro-environment behavior receive a payment, which can be decomposed into two parts: the amount that compensates them for the cost of changing their behavior and the extra amount that is a “pure transfer” to them. To maximize the program’s environmental benefits, a policy maker would like to set the pure transfer component to zero, but the pure transfer is the only part of the payment that increases participants’ economic well-being. In practice, PES programs pay out some pure transfers, and the extent of the anti-poverty benefits depends on whether the pure transfers are de facto targeted to the poor. I lay out these points and then illustrate them with data from a randomized trial of payments for forest protection in Uganda. I provide evidence that the economic gains from participation in PES are indeed larger for those with low costs to fulfill the program’s conservation requirements. I also show that, in this context, poorer eligible households enjoyed more improvement in their economic well-being than richer ones did.

*Contact information: seema@alum.mit.edu. I thank Jake Gosselin for excellent research assistance and Devis Decet, Rebecca Dizon-Ross, Siddhartha Haria, and Kelsey Jack for helpful comments.

1 Introduction

As much as we wish it were not the case, protecting the environment and economic well-being are often in tension. An owner of forested land could conserve the forest or earn income by cutting down and selling the trees, for example. In low-income countries, the trade-off can be stark: further impoverish one's family or harm the environment.

Payments for ecosystem services (PES) are attractive as a way to ease this tension between economic security and environmental conservation. PES programs offer a payment to participants who undertake a specified pro-environment behavior, such as keeping their forest intact. The programs are voluntary, so if engaging in the pro-environment behavior would still make someone poorer or otherwise worse off, even with the compensation, she can decline to participate. Those who opt in and comply with the requirements do so because they believe it makes them better off.

PES lessens the stark trade-off in the sense that it enables additional people to protect the environment without exacerbating their family's poverty. But PES is often framed more strongly, as a win-win: it both protects the environment and *reduces* poverty. While PES indeed can achieve both goals, the dual goals are still, inherently, in tension: the more successful a PES program is in encouraging a participant to undertake pro-environment behaviors, the less effective it is in improving her economic well-being, all else equal.

To see this, it is useful to decompose a PES program's payment into two parts: the amount that compensates the participant for the cost of changing her behavior and the remainder, which is a "pure transfer" to her. A pure transfer is like an unconditional cash grant; it increases someone's financial resources with no strings attached. To maximize the program's environmental benefits, one would set the pure transfer component to zero. The compensation would be just high enough to induce the participant to undertake the pro-environment behavior. Setting the compensation higher would not induce more conservation from her, so any additional program funds would be better used to reach more participants. However, the pure transfer (also known as the inframarginal payment) is the only component of the payment that increases a participant's income. If the payment only compensates her for her costs of complying, the environment is better off and she is neither better nor worse off than she was without the program.

In practice, PES programs pay out some pure transfers, which is why there can be both environmental and poverty-reduction gains; policy makers cannot observe each person's cost

of complying, so cannot set a payment that matches each person’s costs. However, this win-win does not change the fact that, as pure transfers get closer to 0, the environmental benefits per dollar spent increase, but poverty alleviation becomes more modest.

The fact that a PES program’s economic benefits to participants depend on the difference between the payment level and their costs to fulfill the conservation requirements follows from standard economic reasoning, and it is not an insight that is novel to this article. The point is mentioned by Pagiola et al. (2005) and Jack et al. (2008) and discussed in more detail by Alix-Garcia et al. (2015). Rather, the contribution of this article is, first, to explain this insight in a way that can convey the intuition to a broad audience and, second, to present empirical analyses to illustrate it. To my knowledge, Alix-Garcia et al. (2015) is the only prior study that tests the prediction that the economic benefits of PES are larger for those with lower compliance costs. One of this study’s advances is to use machine learning techniques to construct a rich measure of each participant’s compliance costs. For the empirical analysis, I use data from a randomized trial of payments for forest protection conducted in Uganda (Jayachandran et al., 2017).

2 Conceptual framework

Policy makers implement PES to encourage people to pursue a particular pro-environment behavior that entails private costs but generates positive environmental externalities. If the environmental benefits of the behavior are larger than the private costs to undertake it, societal welfare would be higher if people undertook it. PES aims to align the individual’s incentives with society’s by rewarding her for undertaking the behavior.

Consider a PES program that pays a participant M if she undertakes the required behavior, for example, keeping her primary forest intact. The participant incurs a cost, C , to undertake the behavior. The cost could be a monetary outlay (e.g., for kerosene to replace firewood) or a time cost (e.g., patrolling her land so that interlopers do not clear it). A large part of the cost is often opportunity costs. For example, someone who keeps her forest intact forgoes income from selling timber.

In some cases, $C < 0$. Here, the person enjoys enough private benefits from the activity that she would undertake it even without the PES program. The premise of PES, however, is that $C > 0$ for at least some participants. PES would not be needed if everyone were undertaking the pro-environment behavior under the status quo.

An eligible person falls into one of three categories:

(1) Does not undertake the behavior. This case applies if $C > M$.¹ The payment is not enough to offset the costs and make the behavior in the person's private interest. The program makes no payment to her and generates no environmental benefits from her.

(2) Undertakes the behavior and would have done so even without the PES payment. This applies if $C \leq 0$. The program pays out M and generates no additional environmental benefits from the participant because she would have undertaken the pro-environment behavior even without the program. Her economic status improves by M . The entire payment is a pure transfer.

(3) Undertakes the behavior only because of the PES payment. This applies if $0 < C \leq M$. The program pays out M and generates additional environmental benefits from the participant. Her economic status improves by $M - C$. The amount C compensates her for her costs of compliance, and the remainder $M - C$ is a pure transfer to her.

The trade-off between the environmental and anti-poverty goals of PES can be seen by comparing cases (2) and (3). In case (2), there is no environmental benefit but the largest increase in the participant's economic well being. In contrast, in case (3) PES generates environmental benefits, but the economic gains to the participant are smaller than in case (2).

Moreover, when the program has environmental benefits, i.e., in case (3), the environmental benefits per dollar spent are maximized if the pure transfer is set to 0. Suppose everyone eligible has a compliance cost of C . When $M = C$, everyone complies, and the cost per complier is C . Consider an alternative payment $M' > C$. Compliance remains 100% but the cost per complier is higher. For the same total budget, a program could enroll more people at $M = C$ and thereby achieve more conservation. But when $M = C$, there is no pure transfer to the participant. Thus, the PES design that achieves the most environmental benefits for a given budget is the one that results in participants' economic status being identical to what it would have been without the program.

In reality, participants will differ in their compliance costs. A program could attempt to elicit the payment level that would just induce each person to comply and then price discriminate, setting person-specific payment levels. However, most programs use a uniform payment for reasons of simplicity and fairness. They might pay a fixed amount per year

¹I assume that, when indifferent, a person chooses to protect the environment. I also assume no extra cost to sign up for and meet the program requirements, beyond the cost of the pro-environment behavior. I discuss these potential extra costs in the next section.

per hectare of forest that is kept intact, for example. This means that PES programs, in practice, do pay out some pure transfers.

The efficacy of the pure transfers in reducing poverty depends on how poor the recipients of them are. This is determined by how poor those eligible are and whether, among them, costs of compliance are higher or lower for poorer people. If compliance costs are lower for the poor, then the pure transfers have a pro-poor tilt. If the costs are higher for the poor, then the targeting is relatively regressive and not what one would choose in a dedicated anti-poverty program.

3 Evidence from PES program in Uganda

This section uses data from the PES program evaluated in Jayachandran et al. (2017) to provide evidence on the ideas discussed above. The PES program ran for two years in western Uganda. Participating forest owners received USD 28 (in 2012 USD) per year per hectare of primary forest they owned if they kept it intact. The program was implemented in 60 villages randomly selected from the 121 villages in the study sample. The study enrolled 1,099 forest owners. The main outcome was deforestation measured using satellite imagery. In addition, household characteristics and other outcomes were collected through a baseline and endline household survey.

Jayachandran et al. (2017) reported that, on average, 28% of eligible households enrolled in and complied with the program, leading to an additional 5.5 hectares of intact forest per treatment village, on average. There was no strong evidence that the program improved economic well-being for the sample overall. This is consistent with the PES payments offsetting forest owners' compliance costs. Under the status quo, many forest owners cut trees and sold them to charcoal or timber dealers or cleared forest to use the land for cultivation. They needed to forgo this income to comply with the PES contract.

Below, I investigate the compliance costs and economic benefits of the PES program in greater detail. Forgone income from selling trees serves as the measure of compliance costs, C , and a self-assessment of overall economic well-being serves as the measure of economic benefits, $M - C$.

First, I present evidence that the compliance costs and economic gains from participation in PES are negatively correlated, consistent with the inherent trade-off I laid out in the previous section. Then I investigate how compliance costs and economic benefits of participation vary with the (pre-program) socioeconomic status of the participant. This

relationship bears on whether the pure transfers made by program were de facto pro-poor.

Are the economic benefits of PES higher when compliance costs are lower?

I test the hypothesis that people who had to pay more costs to comply with the PES contract enjoy less of an economic benefit from participating. Compliance costs are mostly opportunity costs in this context, namely forgone income from selling trees and cultivating more land. The challenge is that one does not observe forgone income; it is a counterfactual of what income would have been for forest owners in treatment villages had they not been offered PES. Fortunately, the forest owners in the control group offer a way to estimate this counterfactual, as they are ex ante (nearly) identical to those in the treatment group. I use baseline characteristics to predict forest owners' endline income from clearing forest. I then use the model estimates to construct the predicted forest income, absent the program, for the treatment group. A household's predicted forest income serves as a measure of its PES compliance costs.

I focus on income from selling trees as the measure of opportunity costs. First, this source of income was more lucrative than the other opportunity cost, forgone agriculture income from cultivating newly cleared land. Second, while the endline survey collected data on income from cultivation, it is difficult to attribute the portion that was from newly cleared land.

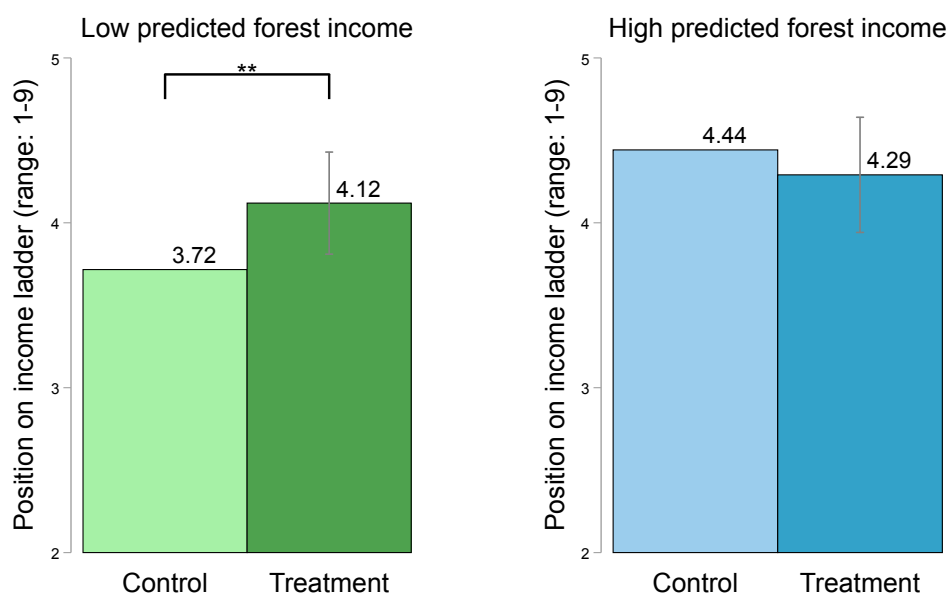
Using the control group, I use LASSO to predict income from selling trees, or more precisely, income per hectare of forest owned.² (The PES payments were per hectare of forest, so the relevant opportunity costs are per hectare.) The potential predictors include variables from the baseline balance table in Jayachandran et al. (2017), baseline forest cover (satellite-based measure), and additional variables that, intuitively, seemed like they might predict deforestation, such as owning a saw and the household's distance to a main road.³

Economic well-being was not a main focus of the original study, so the endline survey did not include a comprehensive income or expenditure module. One proxy for total income, however, is a survey question that asked, "Imagine a 9-step ladder where on the bottom, the first step, stand the people in your community who earn little money, and on the highest step, the 9th, stand the people who earn the most money in your community. On which step

²To increase predictive accuracy, I take the inverse hyperbolic sine transformation of the (skewed) forest income variable, which is similar to a logarithmic transformation but allows for 0's. Through the analysis, I top-code forest income at the 95th percentile of the distribution to reduce the influence of extreme outliers.

³The tuning parameter was chosen with 5-fold cross validation. The model was estimated using the *civlasso* command in the software package, Stata.

Figure 1: PES increases economic-well being only for those with low costs of compliance



Notes: The outcome is self-assessed position on an income ladder. Low predicted forest income refers to households with a below-median predicted value of income from forest. Statistical inference is based on a linear regression that allows for non-independence of errors within a village and, for precision, controls for baseline characteristics: household head has no more than 8 years of education, house has a grass or bamboo roof, and house is made of mud and poles. Asterisks denote a statistically significant difference between the groups: * $p < .10$, ** $p < .05$, *** $p < .01$.

are you today?” I use this ladder variable as a continuous measure of overall income. The sample mean is 4.15, and the standard deviation is 2.14.

When testing how economic benefits of PES vary with compliance costs, for ease of presentation, I split the sample into those with below- or above-median compliance costs (where predicted forest income per hectare is the measure of compliance costs). Figure 1 shows the results, comparing economic well-being for the treatment and control groups, by low versus high compliance costs. On the left are those with low costs of compliance. This is the group for whom more of the PES payments should be pure transfers. The mean self-assessed position on the economic ladder in the control group is 3.72. The treatment group reports a higher economic well-being of 4.12. This difference has a p-value of 0.012. The PES program was financially beneficial for this group.

The bars on the right analyze households with high compliance costs. The control group has an average ladder position of 4.44. The treatment group has a somewhat lower

but statistically indistinguishable value of 4.29. The PES program did not change economic well-being among those with high compliance costs. For people who had to sacrifice income to comply, the program made them no better or worse off economically, on average.

The figure also shows that, in the control group, economic well-being is significantly higher among those with high predicted forest income. The difference (4.44 versus 3.72) has a p-value < 0.01 . This link between predicted forest income and overall income could be causal: earning income from selling trees makes the high-predicted-forest-income households richer. Alternatively, certain correlates of being richer, such as being nearer to a major road, could facilitate earning income by selling trees. In either case, this pattern is a preview of the analysis below, which focuses on whether the poorer or richer PES-eligible households enjoy more pure transfers.

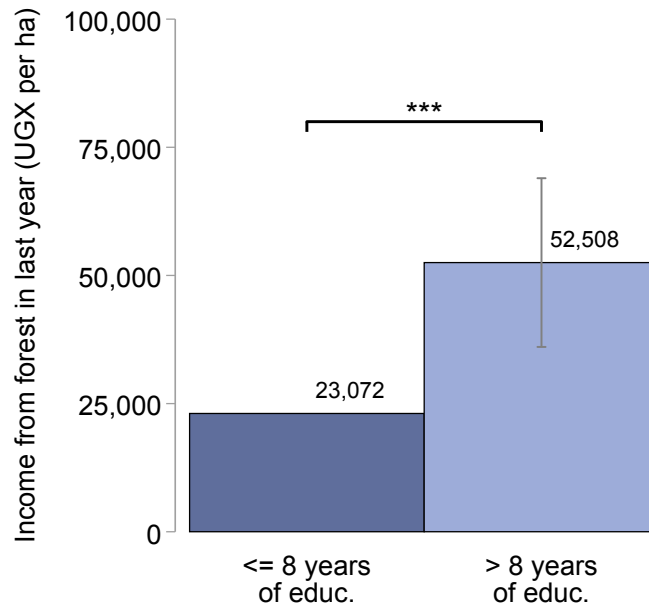
How do PES compliance costs vary with the household's economic status?

With heterogeneity in people's compliance costs and a uniform reward for compliance, a PES program inevitably makes some pure transfers. The welfare benefits of these pure transfers, which are not serving the purpose of protecting the environment, depend on how targeted to the neediest households they happen to be. The answer to that question depends on how compliance costs vary with poverty.

As a proxy for the poorer households in the sample, I use an indicator for the household head's education being below-median for the sample, which maps to having no more than 8 years of education. To test whether compliance costs are lower or higher for poorer households, one can focus on the control group, which represents the status quo without PES. Figure 2 shows that (actual) income per hectare from forest products is much lower among poorer households. This suggests that if poorer households participate and comply with PES at a similar or higher rate than richer households, the pure transfers are progressively targeted.

To simplify the analysis, I have ignored some subtleties. First, the forgone income from selling trees in a given year overstates the opportunity costs because the participant retains the trees. A more accurate measure would be the trees' risk-adjusted expected value in a year multiplied by the annual discount rate (Jayachandran, 2013). Second, if someone's foregone income is sufficiently high that she would not partake in PES, then it does not matter, for the purposes of analyzing potential pure transfers, if her forgone income is even higher. Thus, a more precise test would top-code the opportunity costs at the level of the

Figure 2: Poorer households earn less income from selling trees, absent PES



Notes: The figure uses data from the control group. The confidence interval is calculated using a regression that allows for heteroskedastic errors. Asterisks denote a statistically significant difference between the groups: * $p < .10$, ** $p < .05$, *** $p < .01$.

potential PES payments. While the exact value to use would vary across people and is hard to know, the pattern in Figure 2 is robust to top-coding forest income per hectare at various values: poorer households continue to have lower opportunity costs to protect their forest.

Are the economic benefits of PES higher or lower for poorer households?

The result above would suggest that poorer households should enjoy more economic benefits of PES. However, the opportunity costs of protecting the environment are not the only determinant of compliance. There might be other barriers to participating, for example in signing up or meeting program requirements. The wedge between the costs of protecting the environment and of meeting the PES requirements could advantage or disadvantage the poor. For example, as reported in Jack and Jayachandran (2019), lack of prior experience with written contracts is associated with 10 percentage points lower sign-up, which might be disproportionately excluding poorer households from participation. Other barriers to sign-up could go in the opposite direction, however.

In practice, poorer households signed up for and complied with the PES contract at the same rate as richer households, as shown in Table 1, column (1). Their lower opportunity

Table 1: Poorer households forgo less income when they comply with PES, so enjoy more economic benefits from PES

	Enrolled in and complied with PES program (1)	Forest income in last year (in 100,000's UGX per ha) (2)	Self-assessed position on income ladder (3)
≤ 8 years of education	-0.031 [0.035]	-0.287*** [0.071]	-1.037*** [0.172]
Treated		-0.214*** [0.080]	-0.001 [0.187]
Treated \times ≤ 8 years of education		0.213** [0.093]	0.191 [0.251]
Number of observations	592	967	1,099
Observations included	Treated only	All	All

Notes: Asterisks denote significance: * $p < .10$, ** $p < .05$, *** $p < .01$. Each observation is a household. Standard errors allow for non-independent errors within a village.

costs suggest they should sign up and comply at a *higher* rate, so this results hints of participation barriers disadvantaging the poor, on average. Nonetheless, on net, poorer and richer household receive PES payments at an equal rate.

Next, I analyze (actual) endline income from selling forest products (column 2). The interpretation of the negative and significant coefficient for ≤ 8 years of education is that, in the control group, poorer households earn less income from the forest. This is the result seen above in Figure 2. The negative coefficient on *Treated* means that the program reduced richer households' (the omitted group's) forest income. This is consistent with them refraining from cutting and selling trees and, instead, earning PES payments. The net effect of the program for poorer households is the sum of *Treated* and its interaction with the low-education indicator. The net effect is close to 0: Poorer households did not lose out on forest income from participating, because they would have earned negligible forest income absent the program.

Finally, column 3 analyzes economic well-being, using the income ladder survey question. As expected, in the control group, the low-education subsample has lower self-assessed total income. The zero coefficient on *Treated* means that the program made high-education households neither better nor worse off economically. The PES payments appear to have

just offset their forgone forest income. (This is like the $M = C$ scenario discussed in section 2.) The net effect of the program for poorer households is the sum of the last two coefficients. It is not statistically different from 0 but is positive, consistent with the program lifting up the economic position of poorer participants a bit.

To summarize the results in this section, the economic benefits of PES for eligible households are higher for those with low compliance costs, as one would expect from simple economic reasoning. Economic theory does not have a general prediction about whether poorer or richer eligible households should enjoy more economic benefits, but the empirical answer is important for assessing how valuable the pure transfer part of PES payments is: Even if the pure transfers are not advancing an environmental goal, are they going toward the poorest of the poor, whom an anti-poverty program would target? In this context in western Uganda, the compliance costs were lower for poorer households. Because they participated in the program at the same rate as richer households, we therefore expect more of the pure transfers to be flowing to them. Indeed, there is suggestive evidence that the program led to more economic gains for poorer households.

For data reasons, the analysis of the economic gains from PES focused on the short-run effects among eligible households. Another important determinant of how progressive the pure transfers are is whether the eligibility criteria exclude poorer households from even participating (Pagiola et al., 2005). Anecdotally, poorer households are less likely to own forest, but systematic data are not available. In addition, the long-run effects of PES on poverty might differ from the immediate effects. For example, poorer households might be more credit-constrained, so their pure transfers might enable them to invest in profitable opportunities that increase their long-run income. In addition, if a long-term PES program causes some forest owners to move out of agriculture, their economic well-being will depend on how successful they are in their new occupation.

4 Policy implications

The points discussed in this article have several implications for the design of PES programs. First, it is valuable at the design stage to assess how compliance costs vary with participants' economic status, in order to understand how pro-poor the pure transfers will be. These data do not need to be collected for every household; self-reported costs for a small, representative sample would suffice to assess the de facto targeting of the economic benefits of PES participation. A more ambitious approach would be to use an incentive-compatible

elicitation of compliance costs, as done in Jack (2013).

A second related implication is that funding models for PES could be more creative in light of the dual environmental and poverty-reduction benefits. The part of PES payments that are needed to compel people to protect the forest or plant trees should be eligible for carbon financing, for example. However, conceptually, the pure transfers should not be charged to the buyer of a carbon credit. Some PES programs might only be able to offer carbon credits at the prevailing price in the market if the pure transfers are excluded from the cost base. These pure transfers, though, are often achieving a different set of actors' equally important goal of reducing poverty. In such a case, the PES program could be funded through a combination of revenue from carbon credits and aid money from a development agency, for example. The development agency (or philanthropists focused on poverty reduction) could perhaps provide just enough funding to make the carbon credits competitive in the market, assuming that doing so is justified by the PES program's poverty reduction benefits.

A third implication is that to fundamentally change the environment-economic tradeoff discussed in this article, one needs to permanently lower households' cost of undertaking pro-environment behaviors. With this goal in mind, PES programs sometimes make their payments in kind, for example offering inputs for an alternative livelihood. However, such efforts can undermine PES effectiveness if participants value the in-kind payment less than they would its cash equivalent. A more promising option would be a complementary policy to encourage innovation in the market for alternatives to charcoal or alternatives to wood as a building material. Viable substitutes for tree products would lower the price of charcoal or lumber and, thereby, reduce people's forgone income from protecting the forest. If PES programs were paired with such efforts, then in the short run, PES incentives could offset households' compliance costs, and in the long run, the compliance costs would drop. This would allow the PES payments to be lowered over time or might even make them unnecessary.

References

- ALIX-GARCIA, J. M., K. R. SIMS, AND P. YAÑEZ-PAGANS (2015): “Only One Tree From Each Seed? Environmental Effectiveness and Poverty Alleviation in Mexico’s Payments for Ecosystem Services Program,” *American Economic Journal: Economic Policy*, 7, 1–40.
- JACK, B. K. (2013): “Private Information and the Allocation of Land Use Subsidies in Malawi,” *American Economic Journal: Applied Economics*, 5, 113–35.
- JACK, B. K. AND S. JAYACHANDRAN (2019): “Self-Selection into Payments for Ecosystem Services Programs,” *Proceedings of the National Academy of Sciences*, 116, 5326–5333.
- JACK, B. K., C. KOUSKY, AND K. R. SIMS (2008): “Designing Payments for Ecosystem Services: Lessons From Previous Experience With Incentive-Based Mechanisms,” *Proceedings of the National Academy of Sciences*, 105, 9465–9470.
- JAYACHANDRAN, S. (2013): “Liquidity Constraints and Deforestation: The Limitations of Payments for Ecosystem Services,” *American Economic Review*, 103, 309–313.
- JAYACHANDRAN, S., J. DE LAAT, E. F. LAMBIN, C. Y. STANTON, R. AUDY, AND N. E. THOMAS (2017): “Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation,” *Science*, 357, 267–273.
- PAGIOLA, S., A. ARCENAS, AND G. PLATAIS (2005): “Can Payments for Environmental Services Help Reduce Poverty? An Exploration of the Issues and the Evidence to Date From Latin America,” *World Development*, 33, 237–253.