THE ECOLOGICAL FOOTPRINT OF POVERTY ALLEVIATION: EVIDENCE FROM MEXICO'S OPORTUNIDADES PROGRAM

Jennifer Alix-Garcia
University of Wisconsin, Madison

Craig McIntosh
University of California, San Diego

Katharine R. E. Sims
Amherst College

Jarrod R. Welch
University of California, San Diego

February 27, 2010

Abstract

We study the consequences of poverty alleviation programs for environmental degradation in Mexico. We exploit the community-level eligibility discontinuity for a conditional cash transfer program to identify the impacts of income increases on deforestation, and use the program's initial randomized rollout to explore household responses. We find that additional income increases demand for resource-intensive goods. The corresponding production response and deforestation increase more detectable in communities with poor road infrastructure. These results are consistent with the idea that better access to markets disperses environmental harm and the full effects of treatment can only be observed where poor infrastructure localizes them.

Thanks to CONAFOR, Tania Barham, Tina Green, and Agustin Latapi for access to and advice on the use of data, and to Richard Carson, Paul Ferraro, Josh Graff-Zivin, Gordon Hanson, Jeff Vincent, and seminar participants at NBER, PACDEV, UCSD, Amherst College, and University of Wisconsin AAE for helpful comments.

1 Introduction

Is poverty alleviation likely to exacerbate or mitigate natural resource degradation? For policymakers pursuing the twin goals of sustainable development - raising human living standards and improving environmental quality – this is a crucial policy question. Deforestation represents an important type of environmental degradation. Forest resources are both a local resource, whose uses include fuel wood, fodder, timber, watershed protection and habitat, and a global public good. Global net forest cover is estimated to have fallen by 9.4 million hectares (just under one percent) per year during the 1990s (FAO 2000) and carbon emissions from deforestation are estimated at approximately 20% of the global total (IPCC 2007).

Whether higher household incomes increase or decrease pressure on forest resources depends on multiple factors (Barbier & Burgess 1996, Wunder 2001, Pfaff, Kerr, Cavatassi, Davis, Lipper, Sanchez & Timmins 2008) including prices of agricultural and pastoral goods (Pfaff 1999), demand for forest products (Baland, Bardhan, Das, Mookherjee & Sarkar 2007, Fisher, Shively & Buccola 2005, Foster & Rosenzweig 2003), credit constraints (Zwane 2007), returns to alternative household activities (Deininger & Minten 1999, 2002), agricultural intensification and extensification (Shortle & Abler 1999, World Bank 1992), and demand for environmental amenities (Cropper & Griffiths 1994). The complexity of the relationship between household incomes and deforestation means that research has generated few unambiguous theoretical predictions, and the search for sufficiently large, plausibly exogenous sources of income variation for empirical analysis has been a challenging one.

In this paper, we exploit the discontinuity in the community-level eligibility rule for the nation-wide implementation of Mexico's Oportunidades program, as well as random variation in the pilot phase of the program, to study the consequences of poverty alleviation programs for environmental degradation. Oportunidades represents an ambitious attempt to increase consumption among the poor in Mexico by building human capital. The program funnels large cash support payments to households conditional upon their children's school attendance and receipt of regular health checkups. The program has an annual budget of \$2.6 billion, or half a percent of GDP, and treats 40% of rural households, increasing per-capita income among recipients by an average of one-third. The program's rollout featured rigid, centralized eligibility thresholds at both the locality and the household level, with eligibility defined according to a marginality index. It therefore introduced a very large income shock in 1998-2000 that is discontinuous at the point in the income distribution where localities are defined as just "poor enough" to participate in the program.

We link spatial data on deforestation in Mexico from the period 2000-2003 to the location and eligibility of every locality in Mexico, and exploit this data structure to examine whether deforestation rates are affected by the program. While a relatively large literature exists using the household-level discontinuity in Oportunidades (Bobonis & Finan 2008, Angelucci & de Giorgi 2009), a paucity of locality-level data has hampered research using the community-level discontinuity. Exceptions to this trend are Barham (2009)'s paper on the impact of Oportunidades on child health and Green (2005)'s study of political impact. This structure provides us with an unusual ability to study economy-wide effects from the nation-wide introduction of a conditional cash transfer program in a large and diverse country.

We find that exposure to Oportunidades increases deforestation. Changes around the discontinuity imply a six-percent increase in the rate of deforestation among localities already deforesting, and an increase of nearly 40% in the probability that any deforestation occurs in a locality. To understand the micro-behavior that underlies this result we turn to household data from the randomized pilot phase of the program: the Progresa evaluation sample. The experimental data show that the additional household income significantly increases consumption, and recipient households shift strongly into resource-intensive items such as beef and milk. This suggests that the deforestation impacts might be caused indirectly as households shift demand from less land-intensive goods to more land-intensive goods, increasing their "ecological footprint" (Wackernagel & Rees 1996).

A critical feature of impact estimation where localities form the treatment unit is the issue of spillover effects. Given a set of localized demand shocks, better-integrated local markets allow this demand to be sourced from a broader set of producers. To the extent that new demand is satisfied by national or global markets we lose the corresponding link between local consumption increases and local environmental degradation.

This problem is analogous to previous literature on the effects of local rainfall shocks (Keller & Shiue 2008, Donaldson 2009) which suggests that as infrastructure improves, prices become less correlated with local shocks. This is a fundamental causal inference issue in the analysis of market-mediated impacts, but is impossible to disentangle in a sample where the size of markets is homogenous. The Oportunidades discontinuity not only provides a clean source of identification, but does so in a sample with tremendous variation in the access to transportation infrastructure (as measured by road networks), allowing us to investigate the spatial spillover issue using several distinct empirical techniques. We show theoretically that even when the true impact of treatment on production is constant, we are able to detect it only where infrastructure is poor, and thus the source of resources is geographically constrained.

Empirically, we find that consumption increases appear quite constant and are not driven by road infrastructure. The corresponding production increases on the other hand, display a more complex spatial pattern. Consistent with the idea that transportation infrastructure is a significant determinant of the spatial profile of market-mediated production impacts, we find larger deforestation effects in treated localities that have poor road infrastructure and thus are more isolated from outside markets. We also find corroborating evidence at the household-level of a production response (in this case, by richer, non-recipient households) only in treated localities which are more isolated. Finally, we investigate spatial spillovers of treatment using a new method for calculating spatial auto-correlation functions in a regression discontinuity context. This analysis shows the spatial contour of impacts to be flat where roads are good, and to be concentrated around the location of treatment where roads are bad. Overall, our results are consistent with the idea that broader markets may simply disperse environmental harm, and that only by examining places where these harms are localized by poor infrastructure may we capture the full effect of consumption on environmental degradation.

The paper is organized as follows: we begin in the next section by discussing our contribution to the literature on links between poverty and deforestation and the empirical problem introduced by the study of micro-interventions when agents may participate in market transactions on a broader spatial scale. Section 3 describes the Oportunidades program in more detail, and presents the estimation strategy and results of the discontinuity analysis. Section 4 seeks to disentangle the mechanisms through which this impact occurs by using household data from the randomized evaluation phase of the program. Section 5 presents the spatial analysis, and the final section concludes with a discussion of the policy implications of our findings.

2 Poverty, Deforestation, and Spatial Impact Analysis

Disentangling the relationship between poverty and deforestation involves careful examination of three distinct yet interrelated issues: existence of correlation or a causal link; understanding the relevant household decisions; and the role of local markets in mediating the relationship. Much theoretical work has examined the first two issues, but remains largely inconclusive. Reliable predictions regarding the sign of any causal link rely on a thorough understanding of the relevant household decision process. Unfortunately, many of the channels through which household decisions might affect deforestation lead to ambiguous predictions, resulting in the question becoming almost exclusively an empirical one. Empirical analysis of the relationship, however, introduces the third issue. When the household behavior change that drives any potential impact on deforestation passes first through local markets, detection of any such impact relies heavily on the extent to which local markets are connected to outside national and global markets. For the remainder of this section, we discuss in more detail the theoretical and empirical work that has examined the relationship between poverty and deforestation, the relevant household decisions, and our contribution to this literature, concluding with a discussion of the detection of the effects of income changes on local environmental outcomes when impacts can be dispersed.

2.1 Does alleviating poverty increase or decrease forest cover?

Conditional cash transfer programs that seek to alleviate household poverty and improve access to education or health are increasingly popular in developing countries, but may have unintended secondary effects. One possibility that has not received adequate previous attention is the potential for environmental consequences. It is not clear, *ex ante*, whether we should expect income increases to exacerbate or reduce environmental degradation: a large previous literature on the Environmental Kuznets Curve suggests the relationship is complex and non-linear (Stern 2004, Dasgupta, Laplante, Wang & Wheeler 2002, Panayotou 1997).

We focus here on forests as an environmental outcome of interest. Forests are a key local resource and global public good. Understanding how to prevent further deforestation would significantly contribute to efforts to limit greenhouse-gas emissions (Kaimowitz 2008, Stern 2008). However, even if we limit the scope to the relationship between income and deforestation, previous empirical results and theory are ambiguous (Pfaff, Kerr, Cavatassi, Davis, Lipper, Sanchez & Timmins 2008).

Initial work on the development-deforestation link focused primarily on the presence and shape of an Environmental Kuznets Curve (Cropper & Griffiths 1994, Pfaff 2000), positing that forest cover initially decreases as income rises but then recovers as income increases beyond some turning point. Subsequent work has shown both increases and decreases in forest cover as income increases. Foster & Rosenzweig (2003) use a general equilibrium framework to show that devotion of land to the production of forest products should rise as demand rises. They confirm this relationship using long-term changes in income and forest cover across Indian states. Deininger & Minten (1999, 2002) suggest that as countries grow richer, relative returns to off-farm labor would increase and reduce pressure on forests. They illustrate such a relationship in data from Mexico. Zwane (2007) finds that the relationship between income and deforestation in Peru is positive at low levels of income but may be negative at higher levels. Baland, Bardhan, Das, Mookherjee & Sarkar (2007) assesses the impacts of income growth on firewood collection in Nepal and find a net negative but very small effect.

The empirical literature on the relationship between income and deforestation has been hampered by concerns about the endogeneity of income growth. Rates of deforestation are clearly influenced by multiple factors which could be correlated with income shocks. These include population growth, agricultural returns, forest product prices, capital availability, technology, accessibility and institutional variables; see reviews by Angelsen & Kaimowitz (1999), Barbier & Burgess (2001). The endogeneity problem may be particularly severe for studies using cross-sectional variation to identify impacts. Conversely, in studies using panel variation in income (Zwane 2007, Baland, Bardhan, Das, Mookherjee & Sarkar 2007), the relatively small income changes observed in a short-term panel may not reflect true economic development - the magnitudes may not be large enough to correspond to realistic poverty-reduction goals. Also, these short-term fluctuations are different in nature than permanent income changes. Households are likely to respond differently to income changes that are perceived to be substantial and permanent versus small and temporary.

Exploiting Mexico's rollout of Oportunidades allows us to make two contributions to the existing empirical literature. First, the implementation of the Oportunidades program creates an exogenous source of variation in income, allowing for clean identification of causal effects. Second, the magnitude and duration of the program represents a substantial and durable increase in income for a large share of the households in poor communities. We are thus able to estimate impacts using a positive shock to income that is as large as is likely to be achievable by any actual poverty alleviation program.

2.2 The ecological footprint of poverty alleviation

In the set of empirical studies discussed above, several potential mechanisms are proposed to explain how changes in household income affect deforestation. Many of these could apply in the case of programs designed to alleviate poverty by improving incomes. Foster & Rosenzweig (2003) propose that higher incomes will increase demand for forest products which will induce a supply response by households or communities where there is clear ownership of forest resources. In this case we would expect a conditional cash transfer program to result in less deforestation. Deininger & Minten (1999, 2002) suggest that income increases which occur through increased returns to off-farm labor would reduce agricultural land use and ease pressure on land, also reducing deforestation. Although a conditional cash transfer program might not directly raise off-farm wages, it could raise the opportunity cost of leisure, and therefore discourage on-farm production through a similar mechanism.

Other researchers have suggested that income increases could spur capital improvements or technological adoption, which would facilitate agricultural intensification and reduce pressure on forests (Shortle & Abler 1999, World Bank 1992). If poverty alleviation programs also reduce credit constraints, this mechanism would be relevant. Zwane (2007)'s model proposes different deforestation effects of income increases at high and low initial levels in part because of borrowing constraints. Across two periods of decision-making, an exogenous increase in income decreases borrowing in the first period, thereby increasing the money available to purchase agricultural inputs in the second period, and hence the value of cleared land. At low incomes, relaxing the credit constraint increases deforestation while at higher incomes there is an offsetting increase in the marginal utility of leisure which may result in less deforestation.

An advantage of using Oportunidades as a case study is that it was preceded by a randomized pilot program, Progresa. This experimental design, along with the rich household surveys that were conducted as a part of the program, allow us to unpack the household decisions corresponding to observed aggregate deforestation impacts. The evidence in the Mexican case leads us to propose a new indirect mechanism by which higher incomes increase deforestation. It is similar to Foster & Rosenzweig (2003) in that higher incomes increase demand for consumption goods. However, here increased demand increases deforestation because the production of the relevant goods requires more cleared land rather than more forested land. We assume that food goods produced by low-income households are inferior relative to other goods. As household income increases, households substitute away from these inferior goods (e.g. beans) to normal goods (e.g. beef). If the normal good (beef) is more resource intensive than the inferior good (beans) then households will increase their "ecological footprint" as they become richer, resulting in additional deforestation.

2.3 Estimating responses when shocks can be dispersed

In order to test our hypothesis that income changes lead to consumption driven impacts on deforestation, we must address an issue that is fundamental to the estimation of all market-mediated impacts: there is by no means a one-to-one mapping between the location of the consumption change and the location of the corresponding adjustment in production. Particularly, when the treatment unit (and therefore the source of variation in demand) is small relative to the geographic coverage of the program, the extent to which production impacts spill over will determine what is measured by comparing treated and untreated localities. In trying to understand how these locality-level shocks to income alter market demand and supply of forest-intensive resources, we can draw an analogy with the literature estimating the effect of localized rainfall shocks on prices. A well-established result from this literature is that as infrastructure improves, prices become less correlated with localized rainfall shocks and more correlated with the rainfall shocks of adjacent areas (Donaldson 2009, Keller & Shiue 2008). This effect occurs because demand within a given area is sourced from more distant producers when infrastructure is improved, and hence shocks are spread over a greater area.

When we measure market-level treatment effects from localized experiments (even randomized ones), this same phenomenon will generate observed heterogeneity in the measured treatment effect across infrastructure quality. This heterogeneity will be present even if the true, total treatment effect is constant. To see this, we can think of a market as a grouping of a set of units into a single price-setting mechanism, so that shocks to one unit within a market are transmitted to the other units. Let the number of units per market be given by η, which proxies for infrastructure quality. A treatment induces a constant increase in demand equal to τ per unit, and this increase in demand is sourced on average from itself and the $\eta - 1$ other members of the market.

The increase in outcomes within a unit as a function of its own treatment is the part of the effect that does not spill over, namely $\frac{\tau}{\eta}$. In addition to the direct effect of treatment, each unit will receive an expected spillover effect equal to the indirect treatment effect from the number of individuals within the market who were treated. Writing the share treated as σ, then $\sigma\eta$ units per market will be treated and the expected spillover effect will be $\sigma\eta\frac{\tau}{\eta} = \sigma\tau$. The average treatment effect is given by the difference between treated and untreated units, or

$$E(Y \mid T) - E(Y \mid C) = (\frac{\tau}{\eta} + \sigma\tau) - \sigma\tau = \frac{\tau}{\eta}.$$

This says that the experiment measures not the total effect of treatment but only the component of it that does not spill over to other members of the same market. Now if we think of infrastructure (in our case roads) as being an intermediating variable that determines the size of the market, it can be thought of as determining the number of units on to which the treatment effect τ spills. In environments where the road network is excellent, η moves towards infinity and we have a single national market where the measured difference between treatment and control units is zero. With poor road infrastructure, consumption is localized to the spatial unit of treatment, η goes to one and the estimated difference between treatment and control converges on the true total treatment effect, τ. If what we set out to do with our experiment was to measure the total environmental impact of the treatment, then the error, meaning the difference between the true total treatment effect and the result of the micro-experiment is given by $\tau(\frac{\eta-1}{\eta})$, which vanishes as markets become completely autarkic.

In a sample with variability over the quality of local infrastructure, we will observe heterogeneity in impacts even when the actual treatment effect is constant. The reason for this differential is that spatial arbitrage removes the difference between treated and control units when the pixel size of treatment is small and transport costs are low. Under the assumption of homogenous treatment effects, such an argument implies that we only get the correct estimated treatment effect when spatial arbitrage is shut off. This argument is consistent with the results of Foster & Rosenzweig (2003), who observe a positive feedback effect of higher income on forest reserves only in closed economies, but not in open ones. Presumably the reason for this heterogeneity is that closed economies do not arbitrage their increased demand for forest products across global markets, and hence they manifest the full treatment effect on internal markets. In what follows we investigate the heterogeneity in impacts across infrastructural quality and confirm that our largest observed treatment effects occur precisely where they are the most localized.

3 Oportunidades and Deforestation: Overall Impact

3.1 Program description

The intention of Oportunidades is to increase school attendance and health care among poor families in Mexico. The financial scope of Oportunidades is large. The annual budget is approximately $2.6 billion a year, about half of Mexico's anti-poverty budget. It treats some four million households providing cash transfers conditional on health care provision and school attendance. On average the transfers are about one-third of total income in these poor households, clearly meaningful income changes. The program has been widely studied and lauded for its success in achieving these objectives (Schultz 2004, Fernald, Gertler & Neufeld 2008, Skoufias & McClafferty 2001). The transparent nature of its enrollment criteria and benefits has contributed to the admiration of the program, and it is currently being replicated in various other countries.

The program was implemented in stages. The initial implementation of the program (beginning in 1997) was randomized, and combined with detailed household-level data collection. The full rural roll-out of the program occurred mainly in 1998-2000. This phase was not randomized, but was targeted to localities based on a marginality index; this created the discontinuity in treatment which we use. Eligible villages were first selected according to their level of marginality, and then surveys were conducted within villages to determine who would receive payments. Villages without nearby primary schools or health clinics were not eligible to receive the program. Our analysis focuses on the implementation from 1998-2000, as this is the period with the most useful variation and clearly defined eligibility rules, and because it precedes available deforestation outcome measures. We also leave out villages with more than 2,500 inhabitants – the threshold for "urban" communities in Mexico.

3.2 Data description

To conduct the analysis we merge information on localized deforestation with the program evaluation sample of Progresa, the full national eligibility and rollout data for Oportunidades, and a variety of other sources. Our unit of analysis is the locality. Locality-level eligibility for the program is based upon marginality indices calculated by CONAPO, which were created for 105,749 of the approximately 200,000 localities[1].

The spatial coordinates of each village in Mexico, along with the population and marginality index numbers for 1995, are from the National Institute of Geography and Statistics in Mexico (INEGI), and the data describing the roll-out of Oportunidades comes from the Oportunidades office. Although we have information on enrollment by village through 2003, we exclude villages enrolled after 2000, since after this point the rollout of the program in rural areas was largely finished, and the eligibility rules changed.

To measure deforestation at the locality level we rely on data from the Mexican National Forestry Commission (CONAFOR). The data is based on mosaics of Landsat satellite images from 2000 and 2003 (30 m resolution) and was created by CONAFOR under a mandate to accurately measure and monitor deforestation across the whole country (Monitoreo Nacional Forestal). Due to the large areas that must be covered, the classification of changes is based on changes in the Normalized Difference Vegetation Index (NDVI) values across time using comparisons during the dry season. NDVI is an indicator of vegetation cover and is used worldwide to measure changes in forest cover. Although NDVI change is the best available indicator of changes in forest cover, we note that the measure can have some errors due to weather shocks such as unusually high rainfall or drought conditions. These errors are in the dependent variable but are unlikely to be correlated with variation in treatment, conditional on regional fixed effects. In addition, because CONAFOR was primarily concerned with identifying areas of new deforestation, the 2000-2003 analysis does not include information on which areas might have afforested, so our deforestation variable is censored at zero.

To measure baseline forest, we use the National Forest Inventory (NFI) from 2000. These data are based on a combination of remote sensing using Landsat images and field sampling to verify the

[1] Ninety-three percent of the villages for which there is no marginality index had fewer than 25 inhabitants in 2000. The index was created using a principal components analysis based on seven variables from the 1995 Conteo (short census) and 1990 census, including illiteracy rates, dwelling characteristics, and proportion of the population working in the primary sector (Skoufias et al. 1999).

classification system. Because it is not possible to have deforestation without first having forest, and because there tends to be larger measurement error of deforestation when the areas are smaller, we exclude localities which start with less than 10 hectares of baseline forest in the year 2000. Figure 1 shows the distribution of forest across Mexico in 2000.

Program eligibility was defined at the locality level. We have point data on their locations, however data on the boundaries of the more than 150,000 localities does not exist. In order to assign each part of the landscape to a unique locality, we use the method of Thiessen polygons. These assign land to localities based on the closest locality point. This method relies on the assumption that localities are responsible for the land that is closest to them and has the advantage of avoiding the problem of double counting caused by other shapes such as circles around each locality. Figure 2 shows a zoomed in picture of land use in 2000 along with the locality boundaries assigned by the Thiessen polygons method.

3.3 Empirical strategy

We observe a cross-sectional relationship between enrollment in Oportunidades by the year 2000, and suspected deforestation between 2000 and 2003. One way to estimate the effect would be to apply OLS to the equation:

$$\Delta f_i = \alpha + \delta t_i + \beta' X_i + \varepsilon_i \tag{1}$$

where Δf_i represents the change in forest cover in polygon i over the period 2000-2003, t_i is equal to one if the locality associated with the polygon was enrolled in the program by 2000, X_i represents a vector of locality-level characteristics which might also affect deforestation, including poverty, and ε_i are unobserved factors affecting deforestation. If the program had been randomly distributed, then this would be an appropriate way to measure its effect on environmental outcomes. However, it is not randomly distributed, it is distributed to those who are poor, and who may be likely to have higher rates of deforestation even in the absence of the program. In addition, since enrollment in the program is voluntary, it is possible that those communities where enrollment is very high are systematically different than those where enrollment is very low – i.e., that selection problems could bias the estimates of the parameters in equation 1.

If the discontinuity is sharp, meaning that the rule for eligibility perfectly predicts treatment, then one can simply include the eligibility rule as a proxy for the treatment itself. In our case, this would be a dummy variable (E_i) equal to one if the locality's marginality index exceeds -1.2 – the point at which enrollment begins to rise very sharply. This corresponds to the boundary between "medium" and "low" levels of poverty, as classified by the index. We use this approach in several specifications, understanding that it is only an imperfect proxy for treatment. It measures intention to treat, rather than the actual effect of being treated.

Figure 3 examines the discontinuity visually, retaining the marginality index as the X-variable, but then plots the two critical dimensions of the discontinuity structure on the vertical axes. On the right axis we see the scatterplot of the proportion of localities enrolled in Oportunidades by 2000 according to bins across the distribution of the marginality index. It is important to note that the number of observations in each bin varies considerably across bins, with few observations in the extreme bins and many more per bin towards the middle. We do not use information on the presence of schools or clinics in these villages, but we observe that even without using this data to exclude villages, we see a sharp increase in enrollment above values of -1.2 on the marginality index. The proportion enrolled remains high for intermediate values of the marginality index and then is lower at high levels of marginality; we suspect that the decreases in enrollment at very high levels of marginality may be related to the fact that the very poorest villages may not have been eligible as a result of their lack of infrastructure.

The solid line in Figure 3 shows the smoothed deforestation rate by marginality bin, estimated with a break at the lower end of the discontinuity window, along with a 95% confidence interval[2]. Deforestation rates, while sloping upward across this part of distribution, appear to jump by around 50% precisely at the discontinuity. Note that because income is decreasing as we move to the right, a treatment that increases income is effectively pushing households to the left on this figure. The implication is that while the cross-sectional data are supportive of a Kuznets-style relationship (deforestation highest in the middle part of the distribution) the eligibility discontinuity lies above this value, and so if we took the Kuznets relationship to be causal, we would have expected an income

[2]We replicate this picture using the 1994 *levels* of forest cover within the Thiessen polygons as a falsification test. Unfortunately, the data on 1994 forest areas is missing large tracts of data in northwest Mexico and in parts of the state of Guerrero. However, given the available data, nearly 30,000 observations, there appears to be no difference in 1994 forest levels (measured in percent of polygon in forest) at the point of the discontinuity either visually or statistically.

increase in this part of the poverty distribution to decrease deforestation. This would appear to provide another piece in the already substantial body of evidence suggesting that cross-sectional Kuznets relationships do not depict a causal link between income and environmental changes.

Figure 4 zooms in on the same picture, showing the sample we will use in our discontinuity analysis. Deforestation rates in the 2000-03 interval average just under half a hectare on the richer end of the discontinuity, but once a locality becomes just poor enough to qualify for Oportunidades average deforestation jumps to nearly one and a half hectares. The data range includes marginality levels from -2 to 1, which constitutes 85% of the total sample with baseline forest and populations less than 2,500. The assumption behind the identification strategy is that households which are very close to each other in poverty measures will be so similar that the only difference in deforestation over this range will come from the receipt of Oportunidades payments. As a robustness check, we also include results from a sample restricted to the range -1.6 to -.4 on the poverty index, which constitutes only 30% of the total sample. This sample is shown in Figure 5. In both cases, there are significant differences between deforestation rates before and after localities become eligible for the program.

Our situation differs from a sharp discontinuity in two ways. First, enrollment is not one hundred percent beyond any threshold. Second, there is a range over which the probability of enrollment increases. Presented with the first problem, some authors have used the eligibility criteria to predict the probability of enrollment[3]. In our case, the discontinuity also shows a differential probability of enrollment over the range between -1.2 and -.9. Because of this, we use a fuzzy discontinuity strategy, following closely the methodology of Green (2005) and Jacob & Lefgren (2004). Nonlinear combinations of the eligibility rule and the marginality index are used as instruments in a system of equations given by:

$$\Delta f_i = \alpha + \delta T_i + \gamma M_i + \beta' X_i + \varepsilon_i \tag{2}$$

$$T_i = \omega + \tau_1 E_i + \tau_2 E_i I_i + \tau_3 M_i + \tau_4 M_i I_i + \mu I_i + \Gamma' X_i + \nu_i \tag{3}$$

where I_i represents the value of the marginality index in locality i, M_i is equal to one over the zone where enrollment increases rapidly (from -1.2 to -.9) and zero otherwise, and the other variables

[3] For a review of regression discontinuity approaches, see Imbens & Lemieux (2008).

are as defined above. The vector X_i may also include, depending upon our specification, the size of the polygon in kilometers squared, the population in 1995, the percentage of the polygon that was forested in 2000, and regional dummy variables.

This specification assumes that the underlying relationship between poverty as measured by I_i and deforestation is linear over the range that we consider. While this may be a reasonable approximation over the narrow range around the discontinuity, we also experiment by including quadratic and higher-order terms to control for potential non-linear effects.

Table 1 presents some simple statistics from the two subsamples comparing average deforestation levels in the eligible and marginal zones for the program. In the "full" sample, there is a positive and significant difference in deforestation between marginal and definite groups, while in the restricted sample this difference is positive but not significant. These simple comparisons of means across the running variable seem to indicate the presence of a jump in deforestation around the discontinuity. They do not, however, control for the underlying relationship between poverty and deforestation, nor do they control for any other covariates which might be correlated with both of these.

3.4 Results

3.4.1 Simple approach

We first present results from the simplest reasonable approach – using the eligible localities as a proxy for the treated localities. Table 2 shows the results of a Tobit estimation where the dependent variable is equal to the natural log of one plus the area (in kilometers squared) of land suspected of deforestation between 2000 and 2003 within Thiessen polygon associated with a locality. The first four columns show results from the full sample, and the last two from the restricted sample. The simplest specification includes just the eligibility criterion, the marginality index, polygon area, baseline percentage of that area in forest, and population. In this case the effect of eligibility is positive and significant. Adding more covariates reduces the size of the coefficient somewhat, but it is still positive and significant. A squared term of the index renders the eligibility variable insignificant, but a cubic term returns its significance and results in an increase in the point estimate. It is possible that the correlation of these terms with the slope of the "marginal" zone creates this effect. The last two columns show positive and significant effects of eligibility on deforestation in the restricted sample.

Among non-eligible localities, the probability of deforestation is 4%, and the value of the dependent variable among those observations is around .09 (.10 for the restricted sample). Considering the marginal effects of eligibility in this light reveals a large effect of the program – it increases the probability of deforestation by 1.5 percentage points – around 37 percent. The increase in deforestation among deforesters fluctuates between .003 to .006. This change constitutes an increase in deforestation among the deforesters of around 6 percent.

3.4.2 Fuzzy discontinuity

Results from the instrumental variables discontinuity approach are presented next. We begin by examining the predictive power of the instruments and then show the impact estimation results. Table 3 shows the results of the first stage OLS regressions of a dependent variable equal to one if the locality was treated by 2000. The first four columns test the power of the set of fuzzy discontinuity instruments on the full sample, and the last three columns for the restricted sample. The variables have the expected signs – being eligible for the program (in the zone above -1.2) increases the probability of enrollment, as does being in the marginal zone. The slope of the increase in probability of enrollment in the marginal zone is given by the interaction of the marginality index with the marginal zone, and is positive and significant as predicted. Estimations 3 – 4 and 6 – 7 include nonlinear terms of the marginality index. F-tests of the set of instruments show that these specifications are somewhat weaker than those including just a linear term of the index. This confirms the suspicion suggested above that the nonlinear terms are correlated with the instruments.

Table 4 shows estimations of the systems of equations given by equations 2 and 3, using as the dependent variable the natural log of 1 plus the area deforested of the polygon, as measured in kilometers squared. The estimates are similar in quality to those of the simplest approach – participation in the program increases both the probability and the amount of deforestation – although they are somewhat larger. This suggests that the results are robust to using the cleanest source of exogenous variation given the eligibility rules of the program.

3.5 Heterogeneity in treatment

As we have discussed above, land use change is a multi-faceted process, driven by changes in returns to land as well as household dynamics. We have no reason to believe that the impact of participation in Oportunidades will be uniform across localities. In fact, we expect that the impact will vary significantly depending upon the profitability of existing forest land in alternative uses and the quality of local transport infrastructure. The most basic dimension at which we might expect the effect to vary is the underlying suitability of land with respect to a production supply response. Table 5 examines this question by looking for heterogeneity across a locality's propensity to be deforested. Column 1 of this table shows the estimation of predicted deforestation, which is conducted using only the ineligible localities of the restricted sample. This estimation shows that deforestation increases with marginality (though not significantly), with the baseline amount of forest, population, and area of the polygon. It decreases with the slope of the land in the polygon. We then use this estimation to predict the probability of deforestation for the whole sample, and interact this predicted value with eligibility for the program. These results are shown in columns 2 – 5, and support our hypothesis: having a higher risk of deforestation (higher quality land) increases the likelihood of deforestation significantly more in localities which are eligible for the program. This result is robust to including non-linear terms of the marginality index. We do not show the parallel results from the restricted sample, although they are nearly equivalent.

Next, we consider the possibility that transport infrastructure might also affect the impact of the program on deforestation. The first three columns of table 6 show the differential impact of eligibility at different categories of road density, where road density is calculated as the kilometers of roads within a ten kilometer buffer around each locality. This number ranges from zero to 139, and the three columns are the sub-sample divided into equal groups according to road density. Here we observe that the program only has a significant positive impact on deforestation where road densities are very low. The second three columns of the table show the effect of land quality – the propensity to deforest – in the given road density categories. Again we observe that deforestation is much more likely to occur where predicted deforestation is high, and that being eligible for Oportunidades increases overall deforestation, but only where road networks are very limited.

In sum, the results show that Oportunidades is associated with an acceleration of deforestation.

Localities that received treatment show greater deforestation than localities with very similar poverty levels that did not receive treatment. However, treatment effects vary across land quality and road networks. The problem of estimating responses when shocks can be dispersed through spillovers as described in Section 2.3 suggests that our failure to detect deforestation impacts in places with good road networks comes not from a lack of environmental harm there, but from a lack of correspondence between the places where consumption increases and the places where production rises to meet that consumption. We cannot test for true heterogeneity in impacts across the quality of road networks jointly with the presence of this spillover heterogeneity, but market spillover would suggest that we move towards the correct deforestation impacts in places with more localized markets. To the extent that some spillovers exist across localities even with the worst road infrastructure in the sample, even these impacts are lower bounds on the true environmental harm. Similarly, the fact that our data structure necessitates the use of a 'baseline' forest cover from two years after the treatment began suggests that the estimates provided here may give a lower bound for the true link between income increases and deforestation.

The strength of the identification provided by the eligibility discontinuity lies not in its ability to disentangle micro channels but rather in the ability to estimate economy-wide impacts of a major national program being implemented at scale. In order to try to understand the household-level changes that underly these broader impacts, we turn to the evaluation data from the randomized pilot of the program.

4 Progresa and Deforestation: Household Channels

4.1 The Progresa data

The initial, experimental phase of Oportunidades was known as Progresa. This pilot phase featured a three-year period during which the intervention was directly randomized at the locality level. Of the pool initially identified for participation in the program (poor and very poor according to the 1995 Conteo), 506 localities were randomized into 320 "treatment" (initial intervention) and 186 "control" (delayed intervention) groups. The experiment included several baseline and evaluation surveys that have been used in previous studies (see Skoufias (2005), Section 3 for a description of the evaluation design). For our analysis, we combine the 1997-98 baseline surveys with the 2000 followup which occurred at the end of the experimental phase. This evaluation design provides a unique opportunity to study the micro-foundations of the household production and consumption decisions that underly the observed deforestation impacts.

We wish to examine both demand and supply-side impacts of the program. A large and careful literature exists on the consumption impacts of Progresa, and the program is found to have increased the intake of meat and animal products (Hoddinott & Skoufias 2004). Given the well-documented significant increase in the resources required to supply an animal-intensive diet (White (2000), Gerbens-Leenes & Nonhebel (2002), Bouma et al. (1998)) and the intense competition between cattle-rearing and forest resources in Mexico (Barbier & Burgess (1996), Kaimowitz (1995)) this seems a natural place to look for a demand-driven increase in pressure on forest cover. To this end, we examine changes in consumption of beef and milk products. We might also suspect that there would be an increase in demand for forest products. Since the survey does not contain direct measures of timber demand, we use a proxy, namely home improvements.The survey does not contain measures of timber demand, with the exception of new housing construction in the form of the number of rooms in the home.

On the production side, we assess changes in the number of cattle owned, number of plots, and plot size. Evidence exists to show that Mexican households substitute into the consumption of home produced goods during periods of economic recession, and hence we may expect that the sourcing of production moves away from the households of beneficiaries (Hicks 2008). In terms of changes in production inputs, we also consider the impact of the program on child labor. Given that this type of labor is disproportionately used on the family farm, this provides an additional reason why households eligible for Progresa/Oportunidades may produce less on their own household farms and consume more goods produced elsewhere. Finally, we investigate whether the ineligible households in treatment localities display spillover effects in production that are not observable in the ineligible households in control localities. Of primary interest is not the previously-established intention to treat effect of the program, but the heterogeneity of this effect across the quality of local transportation infrastructure.

Since the data in this section are generated by a randomized experiment, we could use OLS on the simple difference in difference equation below, restricting the sample to eligible households:

$$y_{it} = \beta_0 + \beta_1 T_i + \beta_2 P_t + \beta_3 T_i P_t + \varepsilon_{it} \quad (4)$$

where y_{it} is the household-level outcome variable related to consumption, land use, or child labor, T_i equals 1 if the household is in a treated locality, P_t is equal to one in the post-treatment period, $T_i P_t$ is the interaction of T_i and P_t, and ε_{it} is the household specific error. Because randomization was at the locality level we use clustered standard errors that are robust to correlation between households within a given locality. β_3 is the treatment effect of interest. To test for heterogeneity of treatment effects across the quality of local transportation infrastructure, we include a second specification for each outcome variable which examines the interaction between treatment effect and the inverse road density in the locality.

Effects on all household level outcome variables with the exception of child labor and land plots used in the eligible sample will be estimated using negative binomial MLE. Child labor is a binary outcome, so for this we show the results of a linear probability model, and in the case of land plots used by household members in the eligible sample we use a Poisson MLE. Note that for all count variables, the choice of OLS, or Poisson or negative binomial MLE does not significantly change the estimated marginal effects; however, use of the count data models should provide us with more accurate standard errors[4].

4.2 Progresa results

The experimental household data present a straightforward narrative as to the micro-level impacts of treatment. Table 7 shows no increase in the direct demand for timber products in the context of the home improvements proxy, but does show confirmation of the strong increase in consumption of resource-intensive goods that results from receipt of Progresa transfers. Meat and milk consumption and the ownership of cows all surge under the treatment; the treatment effects represent increases relative to the baseline mean of of 32%, 24%, and 18%, respectively. These large average

[4]The application of OLS would provide us with consistent estimates of the parameters of interest, however, this would not be the ideal approach. A problem arises in that many of the household-level variables are count variables; in this case maximum-likelihood estimation using a count-data distribution such as Poisson or negative binomial would be more efficient. Poisson MLE is a useful starting point for the analysis of count data, but is often inadequate due to the well-known equality of mean and variance in this distribution. Often count data are overdispersed; in other words, the variance exceeds the mean. A test of overdispersion proposed by Cameron & Trivedi (2005) (p. 671) and applied to the household evaluation data indeed rejects the null hypothesis of equidispersion for all count variables except the number of plots of land used by household members in the eligible sample.

consumption-side impacts prove to be invariant to the quality of local road networks; everyone eats more protein regardless of local infrastructure.

Table 8 presents production-side results. Despite a significant 3 percentage point decrease in the share of children who work in the home, there is not a significant fall in the number of children on the labor force, and there is no increase in either the number of plots or the total area cultivated by recipient households. Therefore, the program does not appear to relieve credit constraints to production in any simple way (or at least, if it does so then this effect is counterbalanced by an increase in the marginal utility of leisure). In addition, the program has no detectably differential effects across the distribution of road density. Therefore Progresa does not appear to provoke a substantial increase in the extensive margin of agricultural production among beneficiary households.

If consumption is increasing in this large treated population and production is not increasing in these households, from where is this additional consumption being sourced? To answer this question Table 9 turns to the analysis of indirect effects; namely those experienced by households that reside in eligibly poor localities but who do not themselves qualify as poor (i.e. spillover effects). Within this group we observe that while the program does not have significant effects on production in this group overall, in road-poor areas there is a significantly stronger increase in the number of hectares under cultivation and in the number of cows being grazed on that land.

These results can easily be cast in the framework introduced in Section 2. Progresa induces greater consumption of resource-intensive goods everywhere, and hence increases pressure on resources regardless of network quality. Since treatment does not increase output among recipient households, this additional demand is put into the marketplace. With low road density, the demand must be met locally and so the spatial distribution of forest pressure maps closely onto the locality of treatment. Where infrastructure is better, the large increase in demand for animal protein (as well as the diversification of fruit and vegetable consumption found in Hoddinott & Skoufias (2004)) will be met through a marketplace that serves up more consumer variety sourced from a greater variety of locations. In such circumstances, a large component of the treatment effect spills over into other places and hence the observed difference between treatment and control localities is small. Importantly, in this very simple framework the *true* total treatment effect is the one observed in the most isolated locations. The environmental damage from consumption is not necessarily lower in places with better infrastructure, it is simply undetectable.

5 Spatial analysis

5.1 Spatial ACFs in a RD framework

Having motivated the idea of a marketplace as a spatial unit within which the impacts of treatment effects are dispersed, we move to examine the spatial contour of the response to treatment in a more direct way. The empirical approach tries to mirror the discontinuity analysis, which is built on the logic that while the distribution of outcomes may be endogenous across the broader distribution of the eligibility score, it is exogenous within a window around the discontinuity. In order to estimate the spatial distribution of impacts directly, we adapt techniques introduced by Conley & Topa (2002) to the regression discontinuity framework. Specifically, since treatment is linked directly with marginality through the locality-level eligibility criterion, spatial auto-correlation functions calculated in the standard way will be endogenous. We therefore construct areas (0-10 km, 10-20, and so on up to 40 km) around each locality, and within each area we count the number of localities within a the band around the discontinuity, and the number of these localities actually treated. This provides a conservative way of using "as if random" saturation in the intensity of treatment in the window around the discontinuity to measure spillover effects.

The underlying information used here is the same as that used in the discontinuity analysis, but the structure of the data is slightly different. Here we divide the country in a grid of equally-sized cells 10 km square. Each cell with forest in it constitutes a potential observation. If deforestation occurs in cell i between 2000 and 2003, then $d_i = 1$. For each cell we also calculate an "intensity of treatment," which is composed of a ratio where the denominator is the number of "study" localities in the cell, and the numerator is the number of villages out of the study villages that receive Oportunidades. We define a study village as one which is in the subsample that we used for the discontinuity analysis, i.e., one which is located between -1.6 and -0.4 on the poverty index. For a given cell, s_{i0} represents this ratio, which we refer to as a saturation. For each cell we also calculate this ratio for all of the neighboring cells, excluding the own cell. This is the saturation at 10 kilometers, s_{i10}, and proceed outwards, calculating saturation in successive rings around a given cell up to 40 kilometers. We also calculate the density of road networks in the 40 kilometers surrounding each cell. We call this variable c_i and interact it with each of the saturation variables to help us understand how road access might affect the probability of deforestation. For areas which have no "study" localities in them, we include a dummy variable equal to one when there are no localities, and for these observations include zeros in the saturation observations. We then drop all cells with no baseline forest cover and estimate:

$$d_i = \alpha + \sum_{k=0,10,20,30,40,50} [\beta_k s_{ik} + \theta_k s_{ik} c_i] + \Gamma X_i + \epsilon_i, \quad (5)$$

where X_i are control variables and ϵ_i is the error term. We calculate standard errors using bootstrapping in order to avoid the problem of spatial autocorrelation of error terms. In estimations of this type, it is quite reasonable to worry about spatial correlation of the error terms. In a standard regression model, this correlation does not bias the coefficient estimates but does create inefficiency. In order to avoid the problem of bias in the coefficients, we run a linear probability model using OLS. We bootstrap the estimation to obtain standard errors for the distribution of the coefficients (for a discussion of spatial autocorrelation in the probit, tests, and estimation strategies, see Pinkse & Slade (1998)). Our theory tells us that deforestation should be most strongly correlated with nearby treatment intensity where infrastructure is poorest.

5.2 Spatial analysis: Results

The results from the spatial regression are shown in Table 10. The table contains only partial results – in all cases, the mean poverty level in each buffer is included, along with the variables indicating zero observations in a buffer. The last column also includes fixed effects at latitude, which capture spatial variation in ecosystem, as well as cultural heterogeneity, to the extent that it varies geographically in Mexico. The variable capturing infrastucture quality is a dummy variable equal to one in the case where there are less than 150 kilometers of road within a 30 kilometer buffer around the locality. In the simplest specification, which does not include interactions of saturations with road density, having low road density significantly decreases the probability of deforestation. In the two versions where interactions are included, however, we observe that road density is very important in determining the effect of program concentration on deforestation. In particular, in very remote areas (those with low road density), the probability of deforestation as a result of Oportunidades recipients nearby increases.

Figure 6 graphs out the reported coefficients from column (3) by distance, for the subsample with high road density and that with density less than 150 kilometers. This provides a visual image of the effect of the program on deforestation according to distance, and shows that for cells which are very isolated, the deforestation effect is highly localized. For well-connected cells, on the other hand, the effect of having more Oportunidades recipients nearby is not significantly different from zero, which corresponds with our hypothesis that good infrastructure may help spread the impacts of the program to the point where they are non-detectable locally.

6 Conclusions

This paper conducts an analysis of the impact of large income transfers on deforestation, taking advantage of the discontinuity created by the eligibility rule for Oportunidades. We find that the income transfer increases deforestation, at least in the population that is just below the marginality level required to be able to receive payments. We then use household data to disentangle the mechanism behind this increase in forest loss. Here we observe that households increase their consumption of two relatively land-intensive goods – beef and milk. We do not detect a corresponding increase in consumption of a good that might increase forest cover through increasing demand for forest products– housing construction. Nor do we detect changes on the production side triggered by exposure to Progresa, and hence we conclude that the observed deforestation effects of the program arise from consumption changes, in other words through an expansion of each household's "ecological footprint" of resource use.

Average household income increases by one-third as a result of the transfers, which leads the probability of deforestation to increase by 30 percent and the rate of deforestation among deforesters to increase by nearly 6 percent. These increases are significant in the entire sample, but are strongest in two subgroups: places that were already at high risk of deforestation, and places with poor infrastructure. These results underline the importance of considering spatial spillovers in the analysis of micro-experiments, and provide no support for the argument that increasing incomes will translate into improved environmental outcomes.

In recent years the use of local average treatment effects in the analysis of development program impacts has come under fire for answering small questions using a non-representative sample, and for obfuscating important sources of heterogeneity in outcomes (Deaton 2009). Although we estimate local average treatment effects in this paper, our use of the national rollout means that we have a very large and heterogeneous sample at the discontinuity. Therefore we are able to exploit the jump in program participation to cleanly identify impacts of poverty reduction but also to investigate a critical source of heterogeneity. Furthermore, the eligibility cutoff that we use for identification in this paper is precisely the extensive margin of the actual program, and hence measures the exact impact of expanding the current program, as in Karlan & Zinman (2009). Hence we submit that the treatment effect estimated in this paper is both policy relevant and has substantial richness in terms of the analysis of heterogeneity.

In terms of the generalizability of these results, it is important to recognize the dimensions in which impacts of a CCT program may not reproduce the dynamics of a more endogenous long-term increase in income. Most obvious is the conditionality; it explicitly seeks to alter the prices faced by households in the use of one input to production, child labor. The program also features conditionality on regular health checkups for beneficiary children, and this increase in focus on their health may lead to dietary changes that would not be replicated with a simple increase in income. Further, Oportunidades payments are made monthly and hence provide a cash flow that may be more suited to consumption than investment. It is quite possible, for example, that an alternative program delivering the same total amount of cash to beneficiary households in one lump sum would have seen more investment and less consumption, particularly if credit markets are imperfect. Finally, no particular household receives Oportunidades payments for longer than they have children of eligible age, and so the program features a rolling beneficiary pool and is not likely to generate the real wealth effects that would be seen if permanent income had increased. Despite these caveats, CCT programs have emerged as a major policy tool in the fight against global poverty, and so to the extent that they present one of the most obvious policy levers for decreasing poverty our results are relevant even if we interpret impacts as limited to these programs.

Our findings, particularly the spatial contours of treatment effects, motivate the idea that transportation infrastructure plays a critical role in determining the "footprint" of environmental impacts. This underlines the empirical issues generated by spatial spillover effects when we examine the production response to market-mediated increases in local demand. A well-established result in the literature on rainfall shocks and on famines is the idea that infrastructure decreases the correlation between localized shocks and local market prices. Extended to a program evaluation context, this logic suggests that when treatment is administered at small spatial units, market-driven spillovers cause an underestimation of the true harm from treatment. By this logic, the strong deforestation impacts seen in isolated parts of Mexico when treated with Oportunidades is deeply troubling, because it is precisely in these environments that we are closest to capturing the full impact of treatment. We see these results not as a criticism of poverty-alleviation programs but rather as a cautionary tale. Should we wish to achieve increases in wealth simultaneously with improvements in environmental quality, our study suggests that carefully designed environmental management schemes should accompany poverty alleviation programs.

References

Angelsen, A. & Kaimowitz, D. (1999), 'Rethinking the Causes of Deforestation: Lessons from Economic Models', *World Bank Res Obs* **14**(1), 73–98.

Angelucci, M. & de Giorgi, G. (2009), 'Indirect effects of an aid program: how do cash injections affect ineligibles' consumption', *American Economic Review* **99**, 486–508.

Baland, J.-M., Bardhan, P., Das, S., Mookherjee, D. & Sarkar, R. (2007), 'The Environmental Impact of Poverty: Evidence from Firewood Collection in Rural Nepal'.

Barbier, E. B. & Burgess, J. C. (1996), 'Economic analysis of deforestation in mexico', *Environment and Development Economics* **1**(02), 203–239.

Barbier, E. B. & Burgess, J. C. (2001), 'The Economics of Tropical Deforestation', *Journal of Economic Surveys* **15**(3), 413–433.

Barham, T. (2009), 'A Healthier Start: the Effect of Conditional Cash Transfers on Neonatal and Infant Mortality in Rural Mexico'. Working Paper, University of Colorado, Boulder.

Bobonis, G. & Finan, F. (2008), 'Neighborhood Peer Effects in Secondary School Enrollment Decisions', *Review of Economics and Statistics* (4).

Bouma, J., Batjes, N. H. & Groot, J. J. R. (1998), 'Exploring land quality effects on world food supply', *Geoderma* **86**(1-2), 43 – 59.

Cameron, A. C. & Trivedi, P. K. (2005), *Microecnometrics*, Cambridge University Press.

Conley, T. & Topa, G. (2002), 'Socio-Economic Distance and Spatial Patterns in Unemployment', *Journal of Applied Econometrics* **17**(4), 303–327.

Cropper, M. & Griffiths, C. (1994), 'The Interaction of Population Growth and Environmental Quality', *The American Economic Review* **84**(2, Papers and Proceedings of the Hundred and Sixth Annual Meeting of the American Economic Association), 250–254.

Dasgupta, S., Laplante, B., Wang, H. & Wheeler, D. (2002), 'Confronting the Environmental Kuznets Curve', *The Journal of Economic Perspectives* **16**(1), 147–168.

Deaton, A. (2009), 'Instruments of Development: Randomization in the Tropics, and the Search for the Elusive Keys to Economic Development'. NBER Working Paper No. 214690.

Deininger, K. & Minten, B. (2002), 'Determinants of Deforestation and the Economics of Protection: An Application to Mexico', *American Journal of Agricultural Economics* **84**(4), 943–60.

Deininger, K. W. & Minten, B. (1999), 'Poverty, Policies, and Deforestation: The Case of Mexico', *Economic Development and Cultural Change* **47**(2), 313–344.

Donaldson, D. (2009), 'Railroads of the Raj: Estimating the Impact of Transportation Infrastructure'. Working paper, MIT Department of Economics.

Fernald, L. C., Gertler, P. J. & Neufeld, L. M. (2008), 'Role of cash in conditional cash transfer programmes for child health, growth, and development: an analysis of Mexico's Oportunidades', *The Lancet* **371**(9615), 828–837.

Fisher, M., Shively, G. E. & Buccola, S. (2005), 'Activity Choice, Labor Allocation, and Forest Use in Malawi', *Land Economics* **81**(4), 503–517.

Foster, A. D. & Rosenzweig, M. R. (2003), 'Economic Growth And The Rise Of Forests', *The Quarterly Journal of Economics* **118**(2), 601–637.

Gerbens-Leenes, P. W. & Nonhebel, S. (2002), 'Consumption patterns and their effects on land required for food', *Ecological Economics* **42**(1-2), 185 – 199.

Green, T. (2005), 'Do Social Transfer Programs Affect Voter Behavior? Evidence from PROGRESA in Mexico, 1997 - 2000'.

Hicks, D. L. (2008), 'Relative Consumption and Income Volatility in Emerging Market Economies'. Unpublished Dissertation.

Hoddinott, J. & Skoufias, E. (2004), 'The Impact of PROGRESA on Food Consumption', *Economic Development and Cultural Change* **53**, 37–61.

Imbens, G. W. & Lemieux, T. (2008), 'Regression discontinuity designs: A guide to practice', *Journal of Econometrics* **142**(2), 615–635.

IPCC (2007), 'Fourth Assessment Report'. Intergovernmental Panel on Climate Change. http://www1.ipcc.ch/.

Jacob, B. A. & Lefgren, L. (2004), 'Remedial Education and Student Achievement: A Regression-Discontinuity Analysis', *Review of Economics and Statistics* **86**(1), 226–244.

Kaimowitz, D. (1995), Livestock and deforestation in central america in the 1980s and 1990s: a policy perspective, Technical Report 9.

Kaimowitz, D. (2008), 'The prospects for Reduced Emissions from Deforestation and Degradation (REDD) in Mesoamerica', *International Forestry Review* **10**(3), 485–495.

Karlan, D. & Zinman, J. (2009), 'Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts', *Review of Financial Studies* .

Keller, W. & Shiue, C. (2008), 'Tariffs, Trains, and Trade: The Role of Institutions versus Technology in the Expansion of Markets'.

Panayotou, T. (1997), 'Demystifying the environmental Kuznets curve: turning a black box into a policy tool', *Environment and Development Economics* **2**(04), 465.

Pfaff, A. (2000), *From Deforestation to Reforestation in New England, USA*, Global Prospects of Deforestation and Forest Transition, Kluwer, Helsinki.

Pfaff, A., Kerr, S., Cavatassi, R., Davis, B., Lipper, L., Sanchez, A. & Timmins, J. (2008), *Effects of Poverty on Deforestation*, Vol. 25 of *Economics of Poverty, Environment and Natural Resource Use*, Springer, Netherlands, pp. 101–115.

Pfaff, A. S. P. (1999), 'What Drives Deforestation in the Brazilian Amazon?: Evidence from Satellite and Socioeconomic Data', *Journal of Environmental Economics and Management* **37**(1), 26–43.

Pinkse, J. & Slade, M. (1998), 'Contracting in space: an application of spatial statistics to discrete choice models', *Journal of Econometrics* **85**, 125–154.

Schultz, T. P. (2004), 'School subsidies for the poor: evaluating the Mexican Progresa poverty program', *Journal of Development Economics* **74**(1), 199–250.

Shortle, J. & Abler, D. (1999), *Agriculture and the Environment*, Handbook of Environmental and Resource Economics, Edward Elgar, Cheltenham, UK, pp. 159–176.

Skoufias, E. (2005), Progresa and its impacts on the welfare of rural households in Mexico, Research Report 139, International Food Policy Research Institute (IFPRI).

Skoufias, E., Davis, B. & Behrman, J. R. (1999), An Evaluation of the Selection of Beneficiary Households in the Education, Health, and Nutrition Program (PROGRESA) of Mexico, Technical report, International Food Policy Research Institute.

Skoufias, E. & McClafferty, B. (2001), Is Progresa Working? A Summary of the Results of an Evaluation by IFPRI, Technical Report FCND Discussion Paper No. 118, IFPRI.

Stern, D. I. (2004), 'The Rise and Fall of the Environmental Kuznets Curve', *World Development* **32**(8), 1419–1439.

Stern, N. (2008), 'The Economics of Climate Change', *The American Economic Review* **98**, 1–37(37).

Wackernagel, M. & Rees, W. (1996), 'Our Ecological Footprint', *Green Teacher* **45**(1996), 5–14.

White, T. (2000), 'Diet and the distribution of environmental impact', *Ecological Economics* **34**(1), 145 – 153.

World Bank (1992), 1992 World Development Report, Technical report, World Bank.

Wunder, S. (2001), 'Poverty Alleviation and Tropical ForestsWhat Scope for Synergies?', *World Development* **29**(11), 1817–1833.

Zwane, A. P. (2007), 'Does poverty constrain deforestation? Econometric evidence from Peru', *Journal of Development Economics* **84**(1), 330–349.

7 Figures

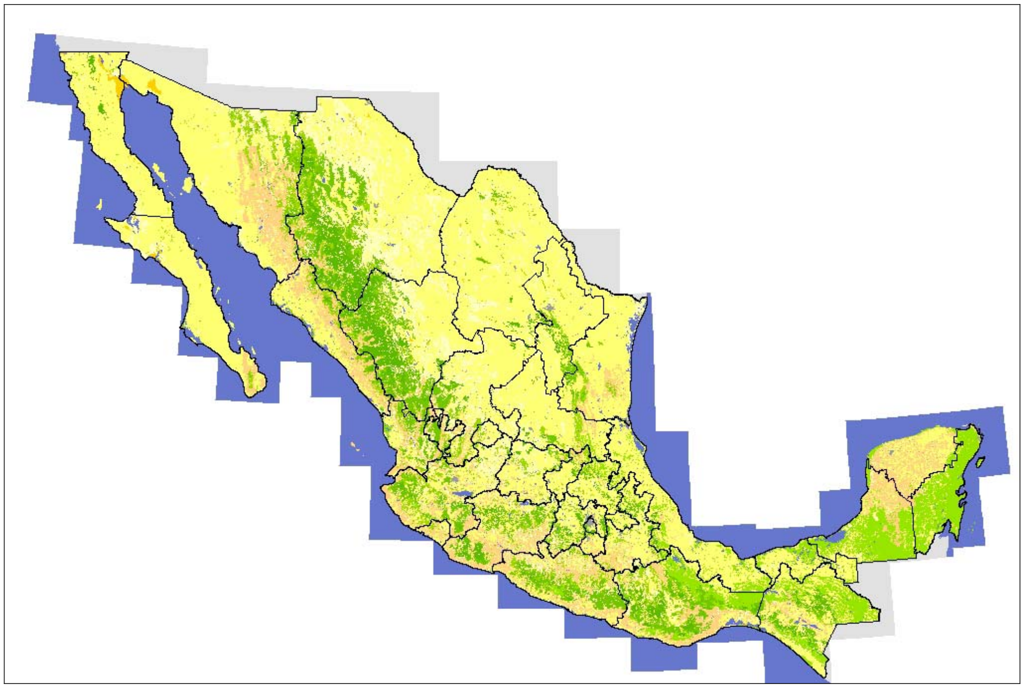

Figure 1: Forest Cover in Mexico, 2000

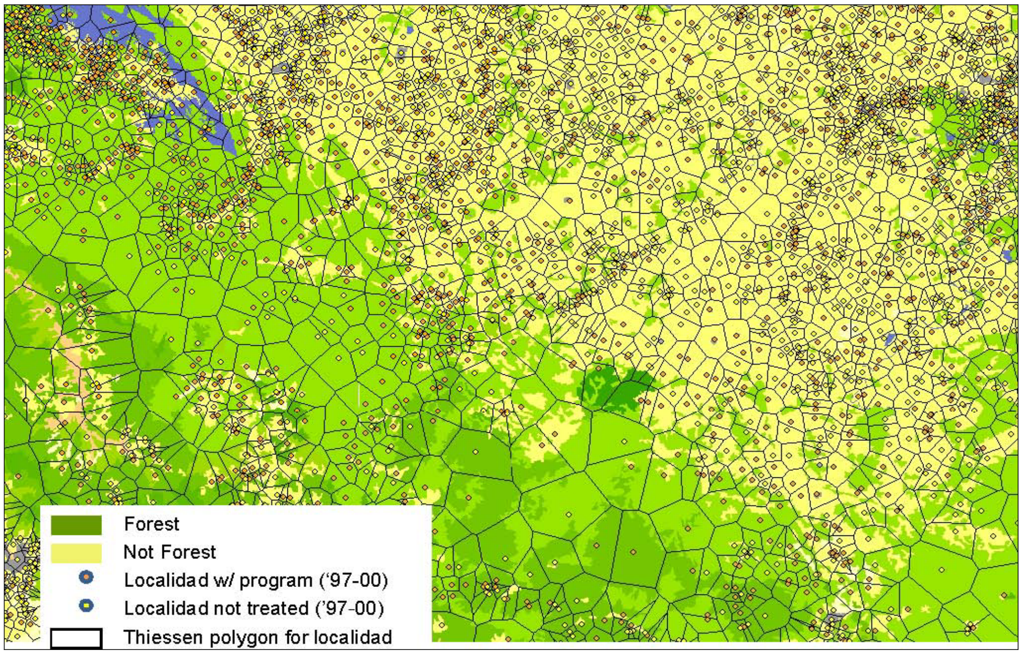

Figure 2: Thiessen Polygons

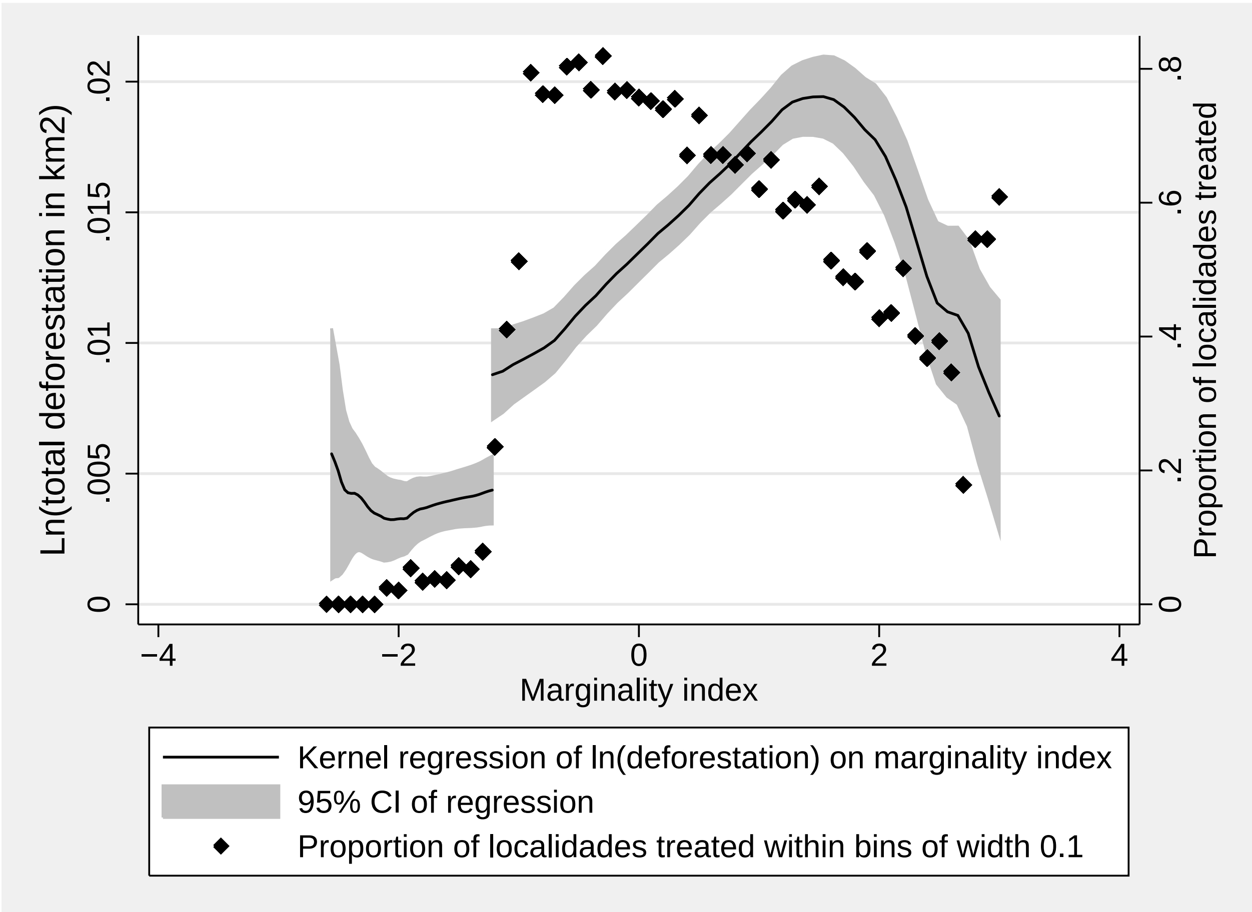

Figure 3: Entire sample minus observations with index > 3 (42 observations missing)

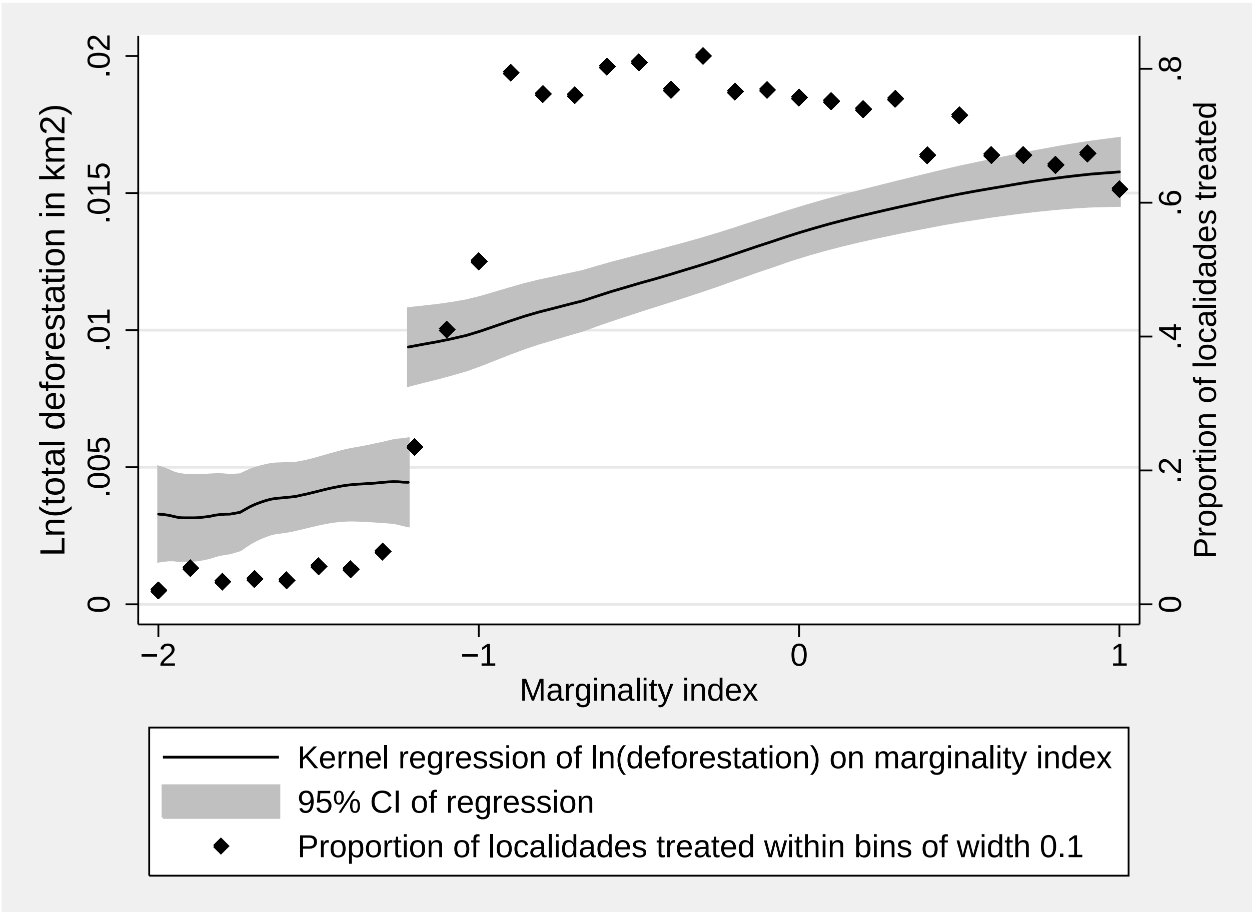

Figure 4: Estimation sample

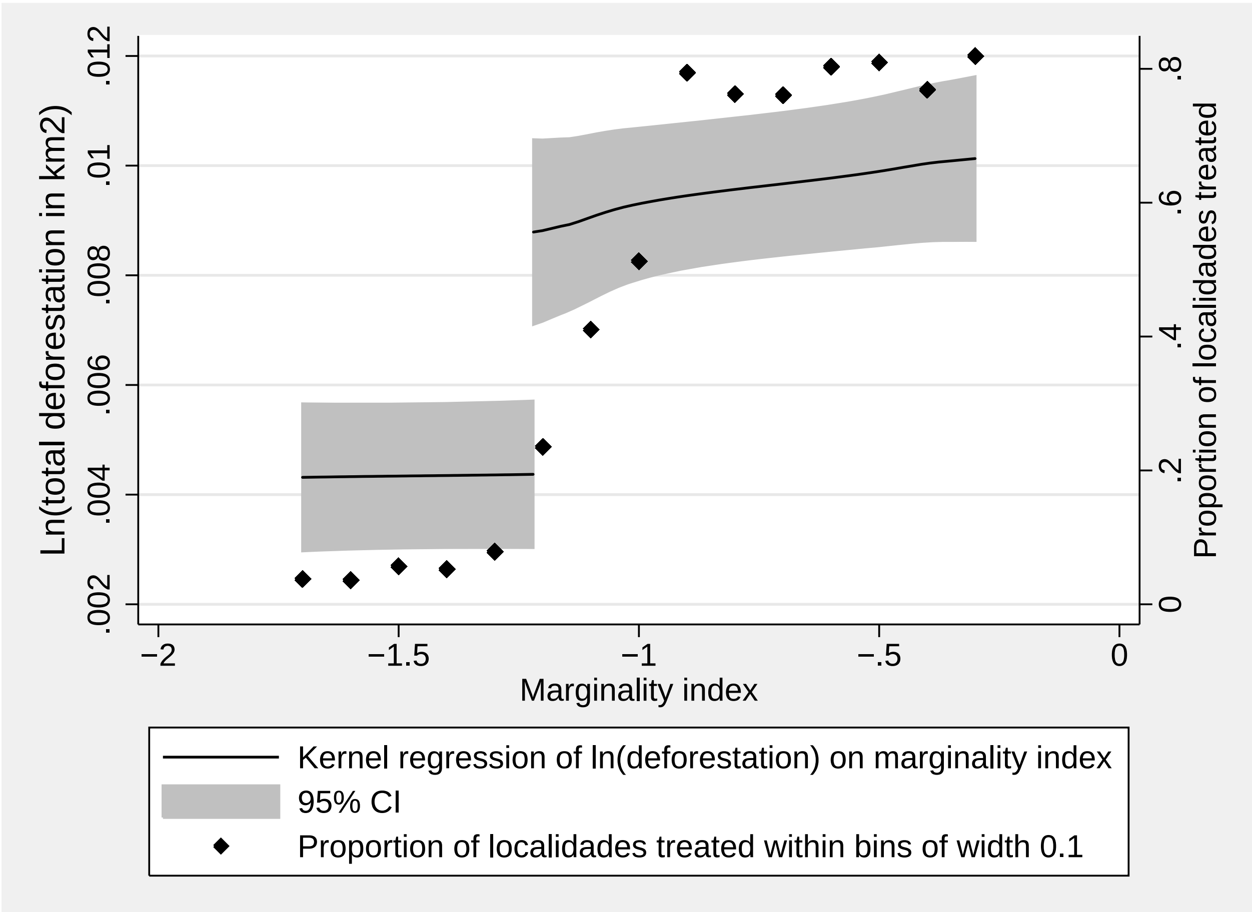

Figure 5: Restricted estimation sample

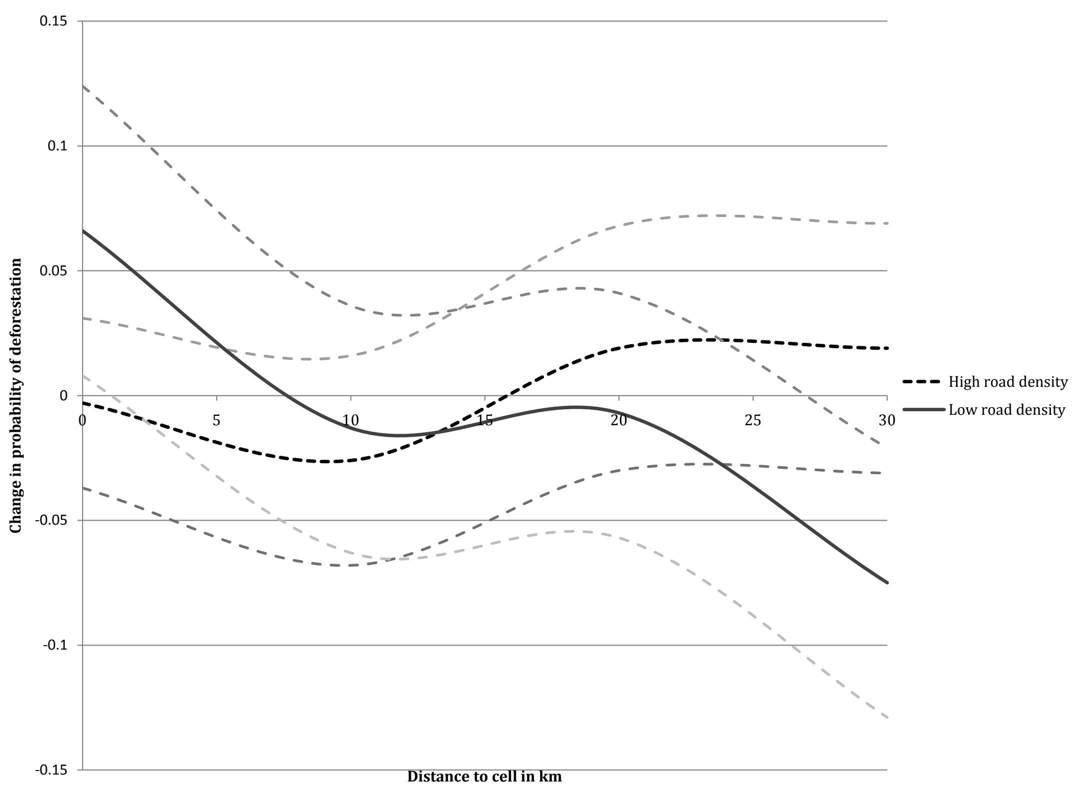

Figure 6: Own deforestation probability as a function of treatment within distance bands

8 Tables

Table 1: Tests of difference in deforestation across marginality zones

	Non-eligible <-1.2 (1)	Eligible $>= -1.2$ (2)	Test of difference (1) vs (2)	Marginal (3)	Test of difference (1) vs (3)	Definite > -.9 (4)	Test of difference (3) vs (4)
Full sample							
Ln(km2 deforested)	.0039	.0146	7.38	.0079	3.47	.015	4.08
Observations	2827	37392		2014		35378	
Restricted sample							
Ln(km2 deforested)	.0043	.0096	3.88	.0079	2.81	.010	1.41
Observations	2166	8626		2014		6612	

Table 2: Tobit Estimation
Dependent variable: log-deforestation

	Full sample				Restricted sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Eligible	.040	.027	.020	.053	.046	.039
	(.017)**	(.016)*	(.023)	(.029)*	(.025)*	(.024)
Marginality index	.075	.045	.079	.090	.044	.014
	(.005)***	(.004)***	(.006)***	(.008)***	(.026)*	(.025)
Index squared			-.008	-.003		
			(.006)	(.006)		
Index cubed				-.011		
				(.006)*		
% polygon in forest, 2000	.114	.194	.113	.113	.194	.230
	(.010)***	(.011)***	(.010)***	(.010)***	(.021)***	(.023)***
Ln(polygon area, km2)	.094	.094	.095	.095	.075	.086
	(.003)***	(.003)***	(.003)***	(.003)***	(.006)***	(.006)***
Ln(Population in 1995)	.034	.020	.034	.034	.024	.013
	(.002)***	(.002)***	(.002)***	(.002)***	(.004)***	(.004)***
Ln(1+slope)		-.055				-.055
		(.003)***				(.008)***
Obs.	40219	40219	40219	40219	10792	10792
Uncensored observations	3844	3844	3844	3844	711	711
Log-likelihood	-8579	-7392	-8578	-8577	-1771	-1522
Ecoregion controls	no	yes	no	no	no	yes
Marginal effects						
on pr($y > 0$)	.017	.011	.008	.021	.015	.011
	(.006)***	(.006)*	(.010)	(.010)*	(.008)**	(.006)*
on $y > 0$	.006	.004	.003	.008	.006	.005
	(.002)***	(.002)**	(.003)	(.004)*	(.003)*	(.003)*

Tobit estimation. Standard errors in parentheses. * significant at 10%; ** significant at 5%;

Table 3: First stage regressions

Dependent variable = 1 if locality received Oportunidades in 1998 or 1999

	Full sample				Restricted sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Eligible	.586	.796	.922	1.013	.624	.797	.434
	(.032)***	(.050)***	(.051)***	(.090)***	(.071)***	(.247)***	(.400)
Marginal	1.453	1.368	1.304	1.272	1.338	1.261	1.315
	(.112)***	(.095)***	(.095)***	(.098)***	(.094)***	(.139)***	(.146)***
Marginal x index	1.718	1.595	1.470	1.427	1.490	1.399	1.454
	(.106)***	(.092)***	(.093)***	(.098)***	(.093)***	(.153)***	(.160)***
Eligible x index	-.163	.043	.217	.306	.011	.180	-.121
	(.020)***	(.032)	(.037)***	(.080)***	(.053)	(.236)	(.351)
Marginality index	.072	-.046	-.189	-.283	.120	-.165	-.123
	(.020)***	(.032)	(.036)***	(.083)***	(.049)**	(.391)	(.393)
Index squared			-.045	-.054		-.098	-.491
			(.005)***	(.008)***		(.133)	(.355)
Index cubed				.009			-.186
				(.007)			(.157)
Percent polygon	-.108	-.027	-.027	-.027	.016	.016	.016
in forest, 2000	(.007)***	(.007)***	(.007)***	(.007)***	(.013)	(.013)	(.013)
Ln(polygon area)	-.027	-.033	-.033	-.033	-.012	-.012	-.012
	(.002)***	(.002)***	(.002)***	(.002)***	(.003)***	(.003)***	(.003)***
Ln(total population		.203	.203	.203	.153	.153	.153
in 1995)		(.001)***	(.001)***	(.001)***	(.002)***	(.002)***	(.002)***
Ln(1+average slope)		.019	.018	.018	.015	.015	.015
		(.002)***	(.002)***	(.002)***	(.004)***	(.004)***	(.004)***
Obs.	40219	40219	40219	40219	10792	10792	10792
Log-likelihood	-23850	-15030	-14983	-14983	-3276	-3275	-3275
Adjusted R-squared	.156	.456	.457	.457	.557	.557	.557
Ecosystem controls	no	yes	yes	yes	yes	yes	yes
F-test of instruments	2560.83	1720.61	508.55	403.31	298.95	189.02	81.01

Linear probability model estimation. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 4: Fuzzy discontinuity estimates
Dependent variable: log-deforestation

	Full sample		Restricted sample	
	(1)	(2)	(3)	(4)
Treated	.068	.042	.109	.092
	(.020)***	(.022)*	(.053)**	(.055)*
Marginality index	.064	.043	.004	-.020
	(.004)***	(.005)***	(.040)	(.041)
Percent polygon in forest, 2000	.106	.197	.186	.229
	(.010)***	(.011)***	(.021)***	(.023)***
Ln(polygon area)	.098	.094	.077	.088
	(.003)***	(.003)***	(.006)***	(.006)***
Ln(total population in 1995)		.013		-.0006
		(.005)**		(.009)
Ln(1+average slope)		-.057		-.057
		(.003)***		(.008)***
Obs.	40219	40219	10792	10792
Log-likelihood	-32544	-22388	-7077	-4798
Ecosystem controls	no	yes	no	yes
Marginal effects				
on pr($y > 0$)	.029	.017	.037	.027
	(.008)***	(.008)**	(.017)**	(.015)*
on $y > 0$	.011	.006	.015	.012
	(.003)***	(.003)**	(.007)**	(.006)*

IV Tobit estimation. Standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 5: Predicted Deforestation risk and impact of Oportunidades
Dependent variable: log-deforestation

	(1)	(2)	(3)	(4)	(5)
Eligible		-.092	-.028	-.056	-.017
		(.028)***	(.029)	(.033)*	(.039)
Predicted Risk		3.116	3.138	3.109	3.183
		(.343)***	(.342)***	(.342)***	(.345)***
Eligible x Predicted Risk		.612	.808	.837	.763
		(.348)*	(.346)**	(.347)**	(.349)**
Marginality index	.022		-.046	-.042	-.030
	(.033)		(.004)***	(.005)***	(.008)***
Index squared				-.009	-.005
				(.006)*	(.006)
Index cubed					-.011
					(.006)**
Percent polygon in forest, 2000	.060				
	(.018)***				
Ln(total population in 1995)	.006				
	(.003)*				
Ln(polygon area)	.025				
	(.004)***				
Ln(1+average slope)	-.008				
	(.006)				
Obs.	2166	40219	40219	40219	40219
Log-likelihood	363	-7543	-7488	-7487	-7485
Adjusted R-squared	.036				
Ecosystem controls	yes	no	no	no	no

Tobit estimation. Standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 6: Road density and impact of Oportunidades

Dependent variable: log-deforestation

	Road density			Road density		
	Low	Medium	High	Low	Medium	High
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	.136	.051	.032			
	(.053)**	(.036)	(.033)			
Eligible				-.092	.065	-.039
				(.062)	(.050)	(.037)
Predicted risk				2.234	4.192	3.347
				(.698)***	(.584)***	(.520)***
Eligible x predicted risk				1.850	-.350	.790
				(.697)***	(.575)	(.507)
Marginality index	.046	.048	.025	-.046	-.039	-.057
	(.008)***	(.008)***	(.008)***	(.008)***	(.007)***	(.008)***
Percent polygon in forest, 2000	.208	.196	.167	-.072	-.073	-.094
	(.021)***	(.017)***	(.017)***	(.022)***	(.017)***	(.018)***
Ln(Population, 1995)	-.026	-.015	-.018			
	(.020)	(.013)	(.011)			
Ln(polygon area in km2)	.090	.079	.101			
	(.006)***	(.005)***	(.007)***			
Ln(1+slope)	-.048	-.057	-.054			
	(.006)***	(.005)***	(.006)***			
Obs.	13406	13406	13407	13406	13406	13407
Log-likelihood	-7831	-6522	-5139	-2890	-2592	-1849

Tobit estimation. The first three columns use instrumental variables, while the second three are standard tobits. Standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table 7: Household-level Consumption Impacts, Progresa

	Negative Binomial							
	Rooms in home		Days Ate Beef		Days Drank Milk		No. of Cows	
Treatment effect (d)	0.014	0.017	0.125***	0.124***	0.351***	0.347***	0.109**	0.087*
	(0.033)	(0.035)	(0.031)	(0.033)	(0.085)	(0.090)	(0.049)	(0.051)
Treatment x inverse road density		-0.032		0.016		0.115		0.134
		(0.120)		(0.140)		(0.548)		(0.178)
Village chosen to receive Progresa (d)	0.000	0.002	-0.022	-0.028	-0.104	-0.112	-0.003	0.035
	(0.038)	(0.039)	(0.025)	(0.026)	(0.087)	(0.093)	(0.069)	(0.068)
Post treatment year (d)	0.053*	0.050*	-0.152***	-0.148***	-0.700***	-0.715***	-0.259***	-0.234***
	(0.028)	(0.029)	(0.027)	(0.028)	(0.072)	(0.077)	(0.045)	(0.047)
Inverse of road density		0.252*		-0.162*		0.038		0.870***
		(0.149)		(0.087)		(0.363)		(0.184)
Village x inverse road density		0.038		0.111		0.190		-0.098
		(0.207)		(0.150)		(0.499)		(0.271)
Post treatment x inverse road density		0.049		-0.082		0.202		-0.102
		(0.113)		(0.107)		(0.311)		(0.089)
Observations	23360	23360	33128	33128	33128	33128	34248	34248
Mean dependent	1.557		0.388		1.440		0.604	
variable in baseline	(0.930)		(0.661)		(2.367)		(2.304)	
Observations	23360	23360	33128	33128	33128	33128	34248	34248
Mean dependent	1.557		0.388		1.440		0.604	
variable in baseline	(0.930)		(0.661)		(2.367)		(2.304)	

** significant at 5%; *** significant at 1%

Table 8: Household-level Production Impacts, Progresa

	Poisson		Neg. Binomial		OLS			
	No. of Plots		Total Hectares		Child works in Home		Child has Job	
Treatment effect (d)	0.032	0.036	0.054	0.063	-0.030***	-0.028***	-0.009	-0.016
	(0.036)	(0.038)	(0.156)	(0.160)	(0.009)	(0.009)	(0.009)	(0.010)
Treatment x inverse road density		-0.136		0.322		-0.062		0.124
		(0.105)		(0.550)		(0.071)		(0.096)
Village chosen to receive Progresa (d)	0.013	0.032	-0.107	0.019	0.022***	0.021**	-0.006	-0.000
	(0.052)	(0.054)	(0.176)	(0.169)	(0.009)	(0.008)	(0.009)	(0.010)
Post treatment year (d)	-0.096***	-0.087**	0.438***	0.425***	-0.004	-0.003	-0.030***	-0.024***
	(0.032)	(0.034)	(0.125)	(0.123)	(0.005)	(0.006)	(0.008)	(0.008)
Inverse of road density		0.559***		2.488***		0.031		0.075
		(0.084)		(0.609)		(0.027)		(0.082)
Village x inverse road density		-0.144		-0.906		0.042		-0.104
		(0.182)		(0.661)		(0.066)		(0.105)
Post treatment x inverse road density		-0.107		-0.401		-0.020		-0.101
		(0.081)		(0.303)		(0.037)		(0.084)
Constant					0.030***	0.028***	0.149***	0.144***
					(0.005)	(0.005)	(0.007)	(0.007)
Observations	45087	45087	32631	32631	50071	50071	50071	50071
Mean dependent	0.824		1.724		0.044		0.145	
variable in baseline	(0.955)		(3.535)		(0.204)		(0.352)	

** significant at 5%; *** significant at 1%

Table 9: Local Spillover Impacts of Progresa
Impacts on Ineligible Households in Treatment Villages

	Negative Binomial Regressions					
	No. of Plots		Total Hectares		No. of Cows	
Spillover effect (d)	0.008	0.002	-0.247	-0.259	0.128	0.034
	(0.033)	(0.035)	(0.204)	(0.196)	(0.097)	(0.102)
Spillover x inverse road density		0.186		1.387**		1.308**
		(0.176)		(0.626)		(0.591)
Village chosen to receive Progresa (d)	0.036	0.073	0.027	0.236	-0.089	-0.024
	(0.047)	(0.048)	(0.259)	(0.243)	(0.161)	(0.160)
Post treatment year (d)	-0.214***	-0.217***	0.345**	0.301**	-0.699***	-0.641***
	(0.027)	(0.029)	(0.167)	(0.152)	(0.095)	(0.096)
Inverse of road density		0.625***		4.334***		2.185***
		(0.091)		(0.898)		(0.440)
Village x inverse road density		-0.660***		-1.916		0.134
		(0.162)		(1.277)		(1.127)
Post treatment x inverse road density		0.051		-0.436		-0.317
		(0.082)		(0.487)		(0.212)
Observations	40569	40569	30068	30068	31184	31184
Mean dependent	1.031		2.844		1.577	
variable in baseline	(1.667)		(5.322)		(4.675)	

** significant at 5%; *** significant at 1%

Table 10: Spatial regressions – dummy variable for low density
Dependent variable = 1 if deforestation

	(1)	(2)	(3)
Own saturation	.012	.004	-.003
	(.018)	(.018)	(.017)
Within 10-20 km	.012	.0009	-.026
	(.022)	(.022)	(.021)
Within 20-30 km	.090	.105	.019
	(.023)***	(.025)***	(.025)
Within 30-40 km	.104	.137	.019
	(.021)***	(.027)***	(.026)
Density $<$ 150 km	-.128	.011	-.011
	(.011)***	(.029)	(.029)
Baseline forest	.0005	.0005	.0008
	(.0001)***	(.0001)***	(.0001)***
Density x own saturation		.071	.078
		(.032)**	(.034)**
Density x 10-20 km		.008	-.003
		(.026)	(.024)
Density x 20-30 km		-.094	-.024
		(.031)***	(.029)
Density x 30-40 km		-.127	-.076
		(.041)***	(.036)**
Obs.	10977	10977	10977
R^2	.059	.062	.188
Lat-long fixed effects	no	no	yes

OLS with bootstrapped standard errors.
** significant at 5%; *** significant at 1%