Bigger is Better: Avoided Deforestation Offsets in the Face of Adverse Selection

Arthur van Benthem & Suzi Kerr

Paper presented at the 55th Annual National Conference of the Australia Agricultural & Resources Economics Society, Melbourne, Victoria, February 8-11, 2011

Bigger is Better: Avoided Deforestation Offsets in the Face of Adverse Selection*

Arthur van Benthem[†]
Stanford University

Suzi Kerr[‡]
Motu Economic and Public Policy Research

Running title: Bigger is Better: Offsets and Adverse Selection

*The authors would like to thank Adam Millard-Ball, Robert Heilmayr, Ruben Lubowski, Larry Goulder, and participants in seminars at Stanford University, NZARE, and the 2010 AIMES open science conference in Edinburgh for comments and suggestions. Suzi Kerr would like to thank the Program on Energy and Sustainable Development for support while she was at Stanford. Van Benthem was supported by a Stanford Graduate Fellowship.

[†]Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305, USA. E-mail: arthurvb@stanford.edu

[‡]Corresponding author. Motu Economic and Public Policy Research and Stanford University, Level 1, 97 Cuba Street, Wellington, New Zealand. Phone: (+64) 4 939 4250. Fax: (+64) 4 939 4251. E-mail: suzi.kerr@motu.org.nz

Abstract

Voluntary opt-in programs to reduce emissions in unregulated sectors or countries have spurred considerable discussion. Since any regulator will make errors in predicting baselines and participants will self-select into the program, adverse selection will reduce efficiency and possibly environmental integrity. In contrast, pure subsidies lead to full participation but require large financial transfers. We present a simple model to analyze this trade-off between adverse selection and infra-marginal transfers. We find that increasing the scale of voluntary programs both improves efficiency and reduces transfers. We show that discounting (paying less than full value for offsets) is inefficient and cannot be used to reduce the fraction of offsets that are spurious while setting stringent baselines generally can. Both approaches reduce the cost to the offsets buyer. The effects of two popular policy options are less favorable than many believe: Limiting the number of offsets that can be one-for-one exchanged with permits in a cap-and-trade system will lower the offset price but also quality. Trading ratios between offsets and allowances have ambiguous environmental effects if the cap is not properly adjusted. This paper frames the issues in terms of avoiding deforestation but the results are applicable to any voluntary offset program.

Keywords: deforestation, offsets, adverse selection, REDD, climate change policy, opt-in.

1 Introduction

Many reports (e.g. Stern, 2006) and key policy makers assert that avoiding deforestation is a key short-run climate mitigation option because of the apparently low abatement costs (Kindermann et al., 2008). Melillo et al. (2009) and Wise et al. (2009) both show that it is critically important to price carbon in forests, especially if there are positive incentives for biofuels. Current estimates of the forest carbon supply curve are based on either land use responses to commodity prices[1], or on estimates of the opportunity cost of land (e.g. Kindermann et al., 2006). These approaches do not take into account the difficulty of designing effective policies to address deforestation in developing countries, where most deforestation occurs (e.g. Andam et al., 2008; Pfaff et al., 2007). They assume the application of efficient price-based policies, yet these are hard to achieve. Offset programs have been shown to suffer from serious problems of spurious offsets and low effectiveness as a result of adverse selection.[2]

This paper formally models voluntary price-based policy to avoid deforestation, examining the implications of three key policy levers, project scale, price ("discounting" or trading ratios[3]), and baseline stringency on the policy's economic, environmental and distributional performance. We consider three inter-related specific criteria: *efficiency* is determined by whether land goes to its optimal use - land that yields high agricultural or timber returns should be cleared; land with returns lower than the positive environmental externalities from the forest should not;[4] *value* is concerned with the average cost to industrialized countries of climate mitigation through avoided deforestation; *quality* of offsets is measured as the percentage of offsets that are not spurious. If quality is not taken into account through a more ambitious cap or fund it will reduce 'environmental integrity' (defined as the degree to which real global environmental gains are achieved as a result of the policy).[5]

We use a microeconomic model of land use with a combination of analytical results and numerical simulations to show that (1) baseline uncertainty in a voluntary program leads to reduced

[1] Examples include Kerr et al. (2002) for econometric estimates of transition to agricultural land use in response to prices, the MIT EPPA global general equilibrium model as used in Melillo et al. (2009), or Sathaye et al. (2006) for a partial equilibrium approach.

[2] Adverse selection is caused by a combination of two factors: a *voluntary* element (i.e., agents can choose whether or not to opt in to the program) and *asymmetric information* about the baseline (i.e., the agents know more about their true baseline than the regulator). Montero (2000), Fischer (2005) and Arguedas and van Soest (2009) establish theoretical results for the effects of adverse selection on offsets programs. Montero (1999) gives the first empirical evidence in the case of the US acid rain program. He and Morse (2010) and Millard-Ball (2010) explore similar issues in the energy sector for the Clean Development Mechanism and for sectoral transportation caps respectively.

[3] In our model lowering the price is equivalent to requiring a trading ratio when offsets are used in a cap-and-trade system in which the cap can be adjusted to achieve the same global abatement. A 5:1 trading ratio is an 80% price discount.

[4] Higher efficiency will be associated with more avoided deforestation up to a limit. Avoided deforestation may also be of independent interest because of associated benefits such as flood protection, water quality and biodiversity.

[5] A lower quality of offset also has equity implications because a higher share of the gains from mitigation will go to actors who have not really mitigated.

efficiency and spurious offsets, but increasing the required scale of projects that can participate mitigates these problems; (2) 'discounting' offsets reduces the quality of offsets and reduces efficiency but generally lowers payments per hectare; therefore, the only rationale for offset discounting is to increase the value to industrialized countries per dollar of transfers to offset sellers; (3) more stringent baselines also reduce efficiency but generally improve the quality of offsets while also improving value for money.

Actual price-based policies for climate mitigation in developing countries are still mostly limited to offset programs. Examples include the payments for ecosystem services program in Costa Rica (Sanchez et al., 2007) and the Clean Development Mechanism, where actors are given credit for forest remaining above an estimated and assigned baseline, or for emission reductions below a baseline. Several designs have been proposed for an international program to reduce deforestation.[6] Some are beginning to be implemented on a wider scale - notably Norway's innovative contracts with Guyana, Brazil and Indonesia.[7] All proposed policies have elements of offsets in their design and face a tradeoff between efficiency and the desire of the funders of such programs to get the best value for the money they spend.[8] It has been suggested that programs should discount the price per hectare paid to landowners, or increase their assigned baseline, to "correct" for spurious offsets and the resulting loss of environmental integrity.

Our paper can be interpreted as an analysis of either adding avoided deforestation to a broader cap-and-trade market, or as an international fund used to pay for avoided deforestation to supplement a separate cap on other emissions. Both programs involve a baseline level of forest and provide rewards relative to that. In a cap-and-trade market these rewards would be offsets valued at the market price, whereas in the fund these rewards would be dollars. In both programs industrialized countries pay for the reductions that are achieved in developing countries. These two approaches are equivalent under the following assumptions. First, the rewards must be the same per unit of avoided deforestation. To set fund payouts that meet this assumption requires that the aggregate marginal cost functions of the forest landowners are known so that the market price in the cap-and-trade system can be predicted accurately. Second, the cap-and-trade market emissions cap and the level of the fund can be adjusted so that regardless of which approach is used, both the global environmental outcome and the permit price are identical. That is, the fund level would need to be set such that the environmental gains it achieved were equal to the difference in environmental

[6] These efforts are most recently referred to as REDD - Reducing Emissions from Deforestation and Degradation. A plethora of reports and edited volumes explore the issues associated with the design of REDD and its successor REDD+. See Angelsen (2008), Chomitz (2007), Plantinga and Richards (2008) for recent discussions of the challenges. Strand (2010) points out that offset programs can lead to increased emissions in the short run from countries that have not yet opted in, but use lax environmental standards to *increase* their baseline emissions.

[7] For example see Government of Kingdom of Norway and Government of the Republic of Indonesia (2010).

[8] Several studies provide evidence on the efficiency effects of adverse selection in the context of Costa Rican deforestation (Kerr et al., 2004; Robalino et al., 2008; Sanchez et al., 2007). Busch et al. (2009) focus on the global efficiency effects of different baselines (reference levels) in a deforestation program. Wara and Victor (2008) explore the extent of spurious credits in the Clean Development Mechanism.

gains between the environmental cap of the larger broad cap and trade system (including avoided deforestation) and the original cap and trade market (excluding avoided deforestation).[9]

Our presentation focuses on deforestation but the results are equally applicable to many other internationally funded mitigation options in developing countries, as well as wider applications of voluntary offset programs.

The remainder of this paper is organized as follows. Section 2 presents a simple model of voluntary deforestation policy that operates first at the level of individual plots, and then for larger scales. This demonstrates the trade-off between efficiency loss from adverse selection and the level of transfers, and analyzes how the three policy criteria are affected by the shapes of the distributions of land returns and observation errors. Section 3 discusses how the potential objectives are affected by three different *policy choices*: increasing the scale required for participation, changing the carbon payment (equivalent to "discounting offsets") and changing the generosity of the assigned baseline. Section 4 concludes and summarizes the main policy implications.

2 A Simple Model of Voluntary Opt-In

2.1 Efficient subsidies versus baselines with adverse selection

Consider a continuum of plots of forested land, indexed by i. Decisions on each plot are independent. Landowners decide to either clear fully or keep the forest. Landowners will clear their forest if the net return from deforesting r_i (e.g. agricultural plus timber revenues minus clearing costs) exceeds any payment p_c to maintain forest. Landowner i knows r_i with certainty. The marginal environmental externality from deforestation is defined as δ.[10] Returns r_i are distributed across i with density f_r.

The simplest policy would be to offer a subsidy equal to p_c per plot that remains forested, where $p_c = \delta$. All landowners with $r_i \leq p_c$ will accept the subsidy and not deforest but only landowners with $0 \leq r_i \leq p_c$ will actually change their behavior; landowners with $r_i > p_c$ will (efficiently)

[9] Suppose industrialized countries (ICs) have a joint emissions cap that requires them to undertake abatement of A. Total abatement cost (TAC) is the integral under the IC marginal abatement cost curve up to A. The market price of pollution equals p^*. ICs could use the fund to achieve n further units of abatement (and pay for m infra-marginal, or "spurious", units), at price per unit p_c which may be lower than p^*. Total global abatement would be $A + n$, where n is a function of p_c.

Analogous to the fund, ICs could purchase $n+m$ offsets from developing countries (DCs) at price p_c. This, however, would not be a fair comparison. Under the fund, the global abatement equals $A + n$. Using offsets, and with no trading ratio, global abatement will be $A - m$. The environmental outcome is worse than without offsets (and p_c would be lower). To correct this, ICs must increase their joint abatement target to $A + n + m$. This ensures that, after $n+m$ offsets are purchased from DCs, the IC mitigation effort is back at A and the pollution price at p^*. Global abatement is now also $A + n$.

If a trading ratio $t:1$ is applied, under offsets global abatement will be $A + (t-1)n - m$, where n and m are now functions of t. This could be higher or lower than $A + n$. Again an adjustment to the joint abatement target would be needed to make them equivalent.

[10] We implicitly assume that the amount of carbon per hectare of forest is constant. This could be relaxed with no loss of generality.

deforest.[11] The change in economic surplus ΔS_{eff} from this efficient policy relative to no policy equals

$$Efficiency\ gain = \Delta S_{eff} = \int_0^{p_c} (p_c - r)\, f_r(r)\, dr \tag{1}$$

This achieves efficient deforestation but requires a large transfer of resources

$$Total\ transfer = TT = p_c \int_{-\infty}^{p_c} f_r(r)\, dr \tag{2}$$

The total amount of avoided deforestation is

$$Avoided\ deforestation = AD = \int_0^{p_c} f_r(r)\, dr \tag{3}$$

The average cost ($AC = TT/AD$) to industrialized countries of climate mitigation through avoided deforestation summarizes the *value* of the program to ICs. Under the subsidy, the value is low (average cost is high) if many plots of land have negative returns and so would not have been cleared even without the subsidy.

To avoid large transfers, a second policy option is a voluntary deforestation program that will pay participants an amount p_c for each hectare of forest exceeding an assigned baseline.[12] Landowners know their true forest baselines BL_i:

$$BL_i = \left\{ \begin{array}{ll} 1 & if\ r_i \leq 0 \\ 0 & if\ r_i > 0 \end{array} \right\} \tag{4}$$

If the regulator observes r_i, the efficient solution is achieved by assigning each landowner i the true baseline $BL_i(r_i)$. If $BL_i = 1$ (no deforestation), no payment will be made and the forest will remain intact. If $BL_i = 0$ (full deforestation) and $0 \leq r_i \leq p_c$, the landowner will opt in and choose not to deforest. If $BL_i = 0$ and $r_i > p_c$, the landowner will deforest and forego the payment p_c. If

[11] The r_i may be interdependent. General equilibrium effects mean one landowner's decision whether to deforest will alter returns for others. This could operate through leakage where a landowner who does not deforest reduces supply of timber and/or food which affect prices for those. It could also occur if clearing involves investment in local infrastructure, or induces local service provision or labor supply that make clearing more attractive for neighboring parcels. These effects could also occur if a local government is the entity avoiding deforestation. For example, a farmer education program to raise yields on existing crop land with a new technology or practice could spill over to more intensive production in neighboring communities if the information spreads. This could either increase or decrease r_i. f_r could be thought of as an ex-post distribution of returns when a new set of equilibrium land uses is reached.

[12] If it were practically feasible, a policy that sets $p_c = r_i$ would reduce transfers even further. In a recent paper, Mason and Plantinga (2010) describe a model in which the regulator has the option to provide landowners with a menu of two-part contracts, which consist of a lump-sum payment from the landowner to the regulator and a "per unit of forest" back to the landowner. Under certain conditions, these are type-revealing, where an ex-ante unobserved "type" corresponds to a marginal opportunity cost curve of keeping a fraction of the land forested. A similar approach to maximize the benefits to the developed country funders in an environmental transfer program was developed in Kerr (1995). Our model does not consider this option.

$p_c = \delta$, the remaining deforestation is *efficient*. Efficiency and avoided deforestation are the same as in (1) and (3) but the total transfer is lower by the amount in (5) and hence the average cost (value) is lower (higher). This policy dominates the subsidy if transfers are costly.

$$Decrease\ in\ TT\ relative\ to\ subsidy = p_c \int_{-\infty}^{0} f_r(r)\, dr \tag{5}$$

In practice, however, the regulator cannot observe r_i, but instead observes $\widehat{r}_i = r_i + \varepsilon_i$. The observation error ε_i has density $f_\varepsilon \sim (0, \sigma_\varepsilon)$, is assumed to be symmetric around 0 and independent of f_r. The *predicted* baselines are

$$\widehat{BL}_i = \left\{ \begin{array}{ll} 1 & if\ \widehat{r}_i \leq 0 \\ 0 & if\ \widehat{r}_i > 0 \end{array} \right\} \tag{6}$$

What happens if the government assigns baseline $\widehat{BL}_i$? When $(r_i > 0, \widehat{r}_i > 0)$ or $(r_i \leq 0, \widehat{r}_i \leq 0)$, the assigned baseline coincides with the true baseline. The landowner will make the socially efficient decision. However, if $(r_i > 0, \widehat{r}_i \leq 0)$, the assigned baseline is 1 but the true baseline is 0. The landowner would have deforested the plot in the true baseline, but gets assigned an unfavorable "no deforestation" baseline. Hence, the landowner will not participate in the scheme. This leads to an efficiency loss if $0 \leq r_i \leq p_c = \delta$, since the landowner will now deforest while he would not have done so had his baseline been correctly assigned and he had participated in the scheme. Relative to the efficient outcome in (1) the efficiency loss caused by adverse selection equals

$$\int_{0}^{p_c} (p_c - r) \left(\int_{-\infty}^{-r} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr \tag{7}$$

The amount of avoided deforestation will fall by

$$\int_{0}^{p_c} \left(\int_{-\infty}^{-r} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr \tag{8}$$

Finally, consider the case where $(r_i \leq 0, \widehat{r}_i > 0)$. These landowners would have kept their forest, but now get assigned a full deforestation baseline. This will not affect their behavior, but it implies an additional infra-marginal transfer p_c. The total transfer (TT) is now lower than the subsidy amount (2). TT is given by the sum of *marginal transfers* (MT) and *infra-marginal transfers* (IT). The former are the payments made to landowners that change their decision as a result of the policy and do not deforest. The latter are payments to landowners that would not have deforested without the policy, but get assigned a favorable full deforestation baseline and will therefore opt in.[13]

$$
\begin{aligned}
TT &= MT + IT \\
&= p_c \int_0^{p_c} \left(\int_{-r}^{\infty} f_\varepsilon (\varepsilon)\, d\varepsilon \right) f_r (r)\, dr + p_c \int_{-\infty}^{0} \left(\int_{-r}^{\infty} f_\varepsilon (\varepsilon)\, d\varepsilon \right) f_r (r)\, dr
\end{aligned} \tag{9}
$$

The amount of avoided deforestation is reduced relative to both the subsidy and the full information voluntary program (both given by (3)). Total transfers are lower than under the subsidy (2), but can be either higher or lower than under the full information program (5).[14] The effect of adverse selection on average cost is theoretically ambiguous relative to the subsidy but clearly higher relative to the full information voluntary program.

To obtain intuition for this ambiguity, we use the decomposition in (9) to write AC as

$$
p_c \left(1 + \frac{\int_{-\infty}^{0} \left(\int_{-r}^{\infty} f_\varepsilon (\varepsilon)\, d\varepsilon \right) f_r (r)\, dr}{\int_{0}^{p_c} \left(\int_{-r}^{\infty} f_\varepsilon (\varepsilon)\, d\varepsilon \right) f_r (r)\, dr} \right) = p_c \left(1 + \frac{OS}{AD} \right) \tag{10}
$$

where OS denotes the amount of infra-marginal forest credited, or offsets that are “spurious”. Moving from a subsidy to a voluntary program reduces OS but also lowers AD. For most realistic distributions (described in Section 2.2) the reduction in OS is larger than the reduction in AD, so AC would fall. We use the fraction of offsets that are spurious ($FOS = OS/AD$) as a measure of the offset quality. The cases described above are summarized in Figure 1.

2.2 The impacts of observation error distributions on policy objectives

The tradeoff between efficiency and value depends on the distributions of observation errors. We now analyze the impact of observation error variance on our three policy objectives: economic efficiency, value (AC) and offset quality (FOS; a measure of the environmental integrity of the program).

Equation (7) shows that any change in $f_\varepsilon(\varepsilon)$ that increases the probability mass in the range $[-\infty, -r_i]$, where $0 \leq r_i \leq p_c$, will increase the efficiency loss from adverse selection (assuming $p_c = \delta$) and decrease avoided deforestation. A mean preserving spread such that $F'_\varepsilon(x) \geq F_\varepsilon(x)$ $\forall x < 0$ is sufficient. If the distribution of errors is normal, an increase in variance will generate such a mean preserving spread.

[13] In a cap-and-trade program, infra-marginal transfers would be spurious or non-additional credits.

[14] From (2) and (9), it follows trivially that $TT(baseline) < TT(subsidy)$. However, $TT(baseline)$ is unsigned relative to $TT(full\ information)$, because $IT(baseline) > IT(full\ information) = 0$, but $MT(baseline) < MT(full\ information)$. Generally, $TT(baseline) > TT(full\ information)$. However, if, for example, $f_r(r)$ has no density below 0, $TT(baseline) < TT(full\ information)$.

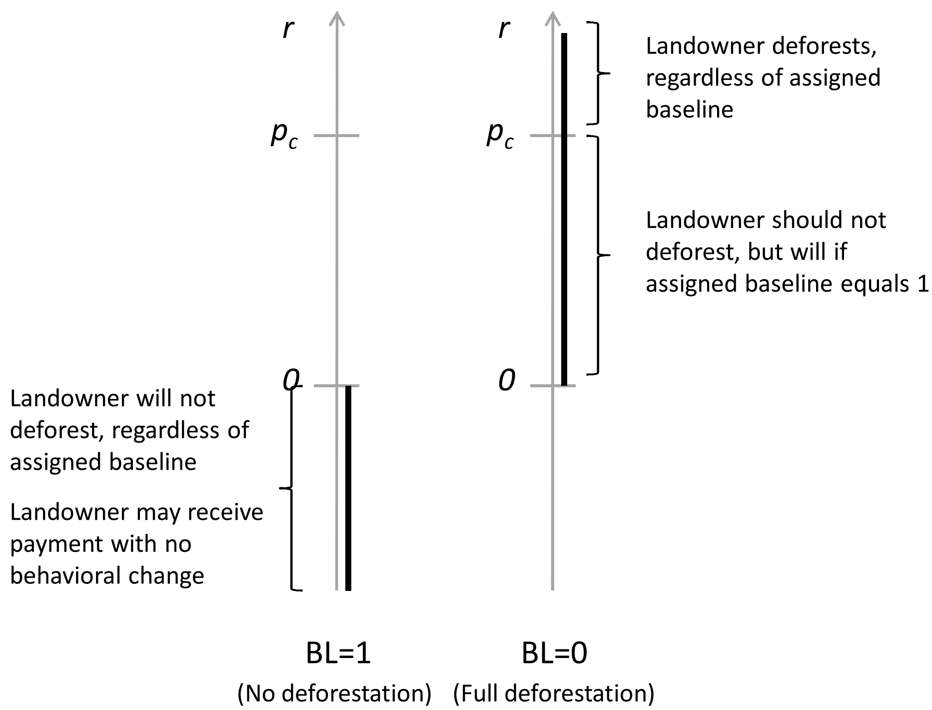

Figure 1: Adverse selection causes efficiency loss in the range $0 \leq r_i \leq p_c = \delta$. It increases average cost in the range $r_i \leq 0$. Both are caused by assigning landowners in these ranges an incorrect baseline.

Under the same assumptions and for $f_\epsilon(\epsilon)$ symmetric around 0, AC will increase.[15] More landowners with $r_i < 0$ will now get assigned $\widehat{BL}_i = 0$ and receive the payment p_c, but they do not provide additional deforestation and funders pay more for less benefit.

Numerical illustration

To provide more intuition for the results, we now assume a parametric form for the distribution of net agricultural returns $f_r(r)$ on forested land and the baseline prediction error $f_\varepsilon(\varepsilon)$. In the remainder of this paper, we will focus mostly on return distributions $f_r(r)$ for which $F_r(0) > 0.5$ and that are downward sloping at 0. The first assumption reflects the reality in key countries that most forested land is not at risk of deforestation. Landowners have previously chosen not to clear the remaining forest so only land on which relative returns have recently risen will still be forested but be at risk of clearing. The second assumption implies that there is a higher probability mass

[15] OS, defined in (10), will increase because $OS(f'_\varepsilon) = \int\limits_{-\infty}^{0} \left(\int\limits_{-r}^{\infty} f'_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr = \int\limits_{-\infty}^{0} (1 - F'_\varepsilon(-r))\, f_r(r)\, dr =$ (by symmetry of f_ϵ) $\int\limits_{-\infty}^{0} F'_\varepsilon(r) f_r(r)\, dr \geq$ (by the mean preserving spread) $\int\limits_{-\infty}^{0} F_\varepsilon(r) f_r(r)\, dr = OS(f_\varepsilon)$. Since AD decreases and OS increases, AC increases.

for returns just below zero than for returns just above zero, which intensifies the tradeoff between efficiency and reducing transfers and, in particular, infra-marginal rewards.

With no shocks, all land with positive returns would already have been cleared without any policy while no land with negative returns would have been cleared. Hence, there will be positive probability mass below zero and no mass above zero and the assumption trivially holds. Deforestation occurs because the returns distribution shifts over time. If this shift, driven by, for example, technology and local infrastructure change, has both a common and an idiosyncratic (e.g. normal unbiased shock to each plot) element we would still expect the second assumption to hold.[16] The density above zero will tend to be lower than below zero, since the tail of the normal distribution implies a negative slope.

We consider $f_r(r) \sim N(-1, 1)$, $f_\varepsilon(\varepsilon) \sim N(0, \sigma_\varepsilon)$ and $p_c = \delta = 0.5$ as our central case. Figure 2 plots the various policy objectives as a function of the standard deviation of the observation error σ_ε: the efficiency loss from adverse selection (7) relative to potential efficiency (1), AC and FOS.

Naturally, the efficiency loss is 0 if the observation error standard deviation $\sigma_\varepsilon = 0$. The efficiency loss is increasing in σ_ε. As σ_ε grows large the assignment of baselines becomes random. Participation, efficiency and avoided deforestation all fall toward 50% of their maxima (at $\sigma_\varepsilon = 0$). Figure 1 shows that efficiency losses only result from landowners with $0 \leq r_i \leq \delta$. These will make the inefficient decision to deforest if and only if they get assigned $\widehat{BL}_i = 1$, which happens with probability approaching 0.5 as σ_ε increases. Figure 1 also shows that offsets that are spurious are given out only to those with $r_i < 0$. As σ_ε increases from zero, the fraction of offsets that are spurious rises rapidly. Combined, the fall in AD and rise in FOS have dramatic implications for AC: AC quickly rises from the efficient value of 0.5 (the environmental externality δ), as FOS becomes large. For σ_ε of 0.3, AC doubles and 50% of the offsets are spurious.

This section has shown that a mean preserving spread that increases the tails of the observation error distribution (which in a normal distribution would be implied by an increased variance) unambiguously has (weakly) negative effects on all three policy objectives. Any improvement in our ability to observe returns, or equivalently predict deforestation, would reduce the tradeoff between efficiency and transfers.

2.3 The impact of different marginal costs of avoiding deforestation on policy objectives

How do the policy objectives depend on marginal abatement cost? While this is largely a function of geography and economic factors and so not easily affected by international deforestation policy, it will influence where efforts to develop programs to avoid deforestation will be most effective. In

[16] The common shock will generate a probability mass of forested land above zero return up to the size of the shock; the idiosyncratic shock will also move some land to higher returns (and some to lower) leaving lower probability mass in the upper tail of the returns distribution.

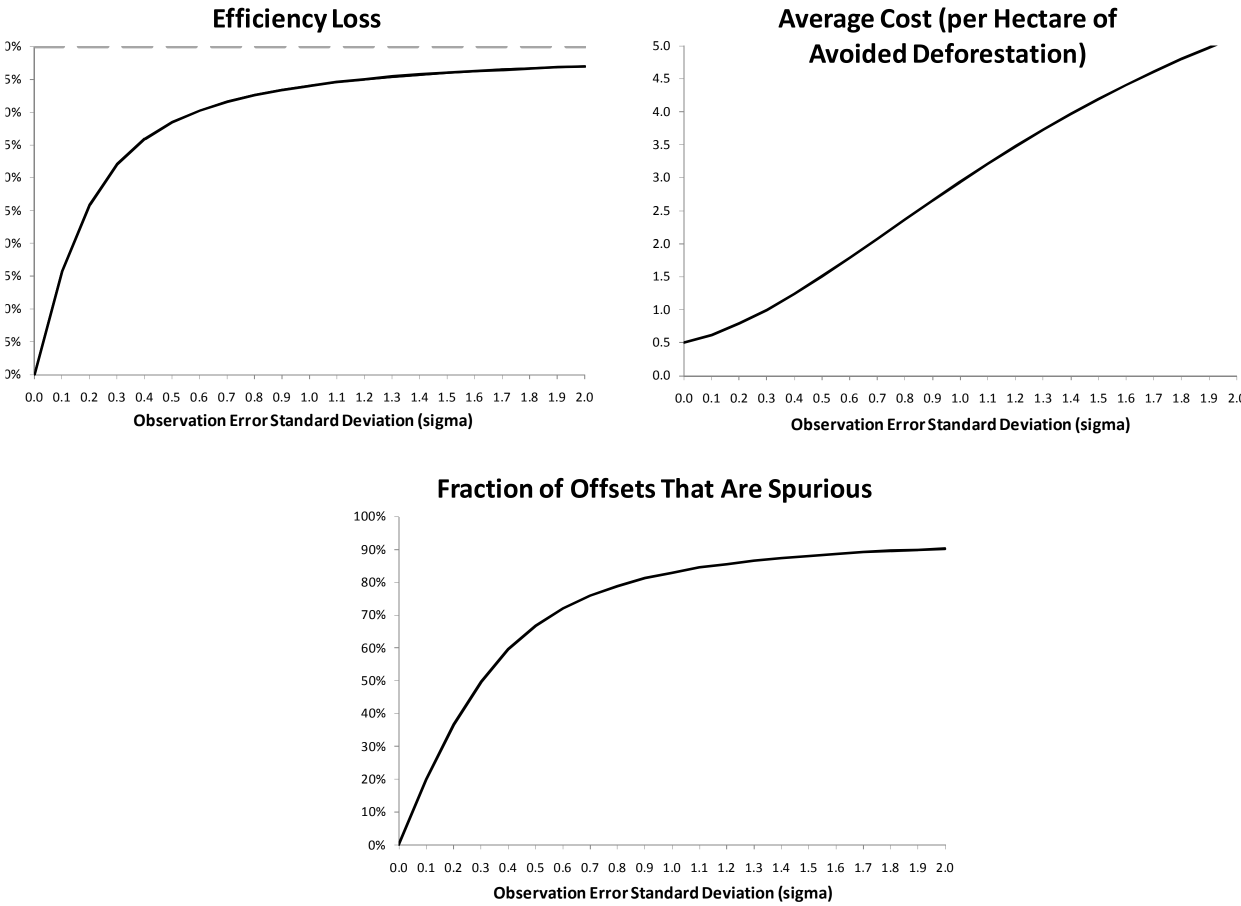

Figure 2: Efficiency loss, AC and FOS as a function of observation error standard deviation σ_ε ($p_c = \delta = 0.5$).

this model, abatement costs are represented by the foregone net return from deforestation r and the marginal abatement cost curve depends on the distribution $f_r(r)$.

We first consider which distributions $f_r(r)$ lead to the largest efficiency gain from voluntary avoided deforestation policy. We abstract from observation errors and adverse selection for now. The efficiency gain relative to no policy (1) depends on $f_r(r)$ through two effects. First, a higher probability mass of returns between $[0, \delta]$ increases the efficient level of AD. Second, a higher probability mass of very small positive returns between $[0, \epsilon << \delta]$ relative to returns between $[\delta - \epsilon, \delta]$ increases efficiency. Therefore, the first condition is not sufficient for an overall efficiency gain.

Figure 3 illustrates this by using three different returns distributions to generate returns distributions that imply different marginal abatement cost curves. As the distribution changes from case 1 ($N(-1, 0.5)$) to case 2 ($N(-1, 1)$), f_r increases for all r between $[0, \delta]$. This increases the deforestation response at every positive price, by unambiguously lowering marginal abatement cost, and hence increases the efficiency gain of the policy. The efficiency gain increases fourfold, while AD increases fivefold. However, moving from case 2 to case 3 ($Uniform(-3.6, 1.6)$), f_r increases for r close to δ, but decreases for small r. AD increases by seven percent, but the efficiency gain *decreases* by three percent. Hence, the relationship between avoided deforestation and the potential efficiency gain is ambiguous.

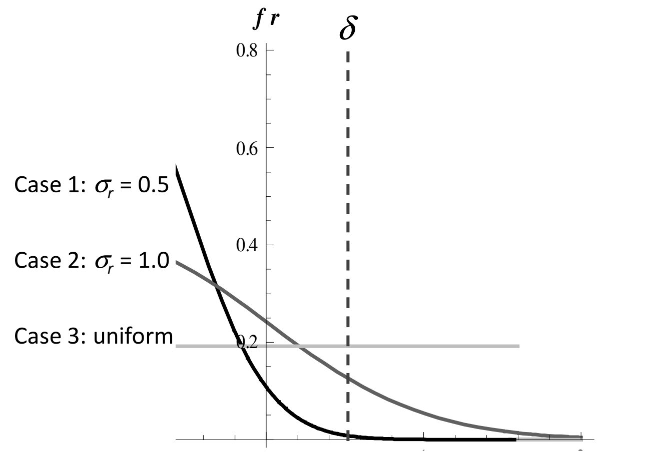

Notes: case 1: $f_r(r) \sim N(-1, 0.5)$; case 2: $f_r(r) \sim N(-1, 1)$; case 3: $f_r(r) \sim Uniform(-3.6, 1.6)$. $\delta = p_c = 0.5$. $f_\varepsilon(\varepsilon) \sim N(0, 0.5)$.

Figure 3: The ambiguous relationship between returns distributions, avoided deforestation and efficiency.

Proposition 1. A returns distribution f_r that generates more AD at $p_c = \delta$ than f_r' does not necessarily generate a higher efficiency gain.

Proof. By counterexample (Figure 3).[17] ■

[17] A more general counterexample can be constructed as follows. Consider a distribution f_r that is downward sloping in the interval $r \in [0, p_c = \delta]$, and a distribution f_r' such that $f_r' = f_r(p_c - r)$ for this interval and $f_r'(r) = f_r(r)$ elsewhere in the domain. First, note that $\int_0^{p_c} f_r'(r)\,dr = \int_0^{p_c} f_r(p_c - r)\,dr = \int_0^{p_c} f_r(r)\,dr$. Second, since $\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon$ is increasing in r, $AD(f_r') = \int_0^{p_c}\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right) f_r'(r)\,dr > \int_0^{p_c}\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right) f_r(r)\,dr = AD(f_r)$. For example, consider $f_\epsilon(\epsilon)) \sim Uniform(-k, k)$ with $k > p_c$. In that case, $\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon = \frac{1}{2}\left(1 + \frac{r}{k}\right)$ for $r \leq p_c$. Hence, $(p_c - r)\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right)$ is decreasing in r. Therefore, $\Delta S(f_r') = \int_0^{p_c}(p_c - r)\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right) f_r'(r)\,dr = \int_0^{p_c}(p_c - r)\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right) f_r(p_c - r)\,dr < \int_0^{p_c}(p_c - r)\left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\,d\varepsilon\right) f_r(r)\,dr = \Delta S(f_r)$, since $f_r'(r)$ is increasing in r while $f_r(r)$ is decreasing in r and $f_r'(0) = f_r(p_c)$. Hence, AD can increase while efficiency decreases.

Note that sufficient conditions for efficiency to increase are $f_r'(r) > f_r(r)\ \forall r \in [0, p_c = \delta]$, or - somewhat weaker - $\int\limits_0^{p_c} f_r'(r)\,dr > \int\limits_0^{p_c} f_r(r)\,dr\ \forall r \in [0, p_c = \delta]$.

Proposition 1 shows that stronger assumptions on f_r and f_r' are needed to ensure an increase in efficiency than an increase in AD: an increase in the (observation error-weighted) probability mass between $[0, \delta]$ is sufficient for AD to increase, but not to guarantee increased efficiency. In other words, a return distribution that leads to a higher amount of optimal avoided deforestation does not necessarily lead to a greater increase in efficiency.

With observation errors, a change in the returns distribution also affects the likelihood of spurious offsets: a less negatively (more positively) sloped distribution around zero yields fewer spurious offsets. The combined effects on AD and spurious offsets determine the effect on AC.

Numerical illustration

Figure 4 illustrates these effects using a numerical example similar to the previous one with $f_r(r) \sim N(-1, \sigma_r)$, $f_\varepsilon(\varepsilon) \sim N(0, \sigma_\varepsilon)$, $\sigma_\varepsilon = 0.5$, $p_c = \delta = 0.5$ and three different σ_r which alter the relevant part of $f_r(r)$. Case 3 now corresponds to a $N(-1, 2)$ returns distribution. Marginal abatement cost unambiguously falls between case 1 and 2 while in case 3 it is higher than 2 for some units and lower for others.

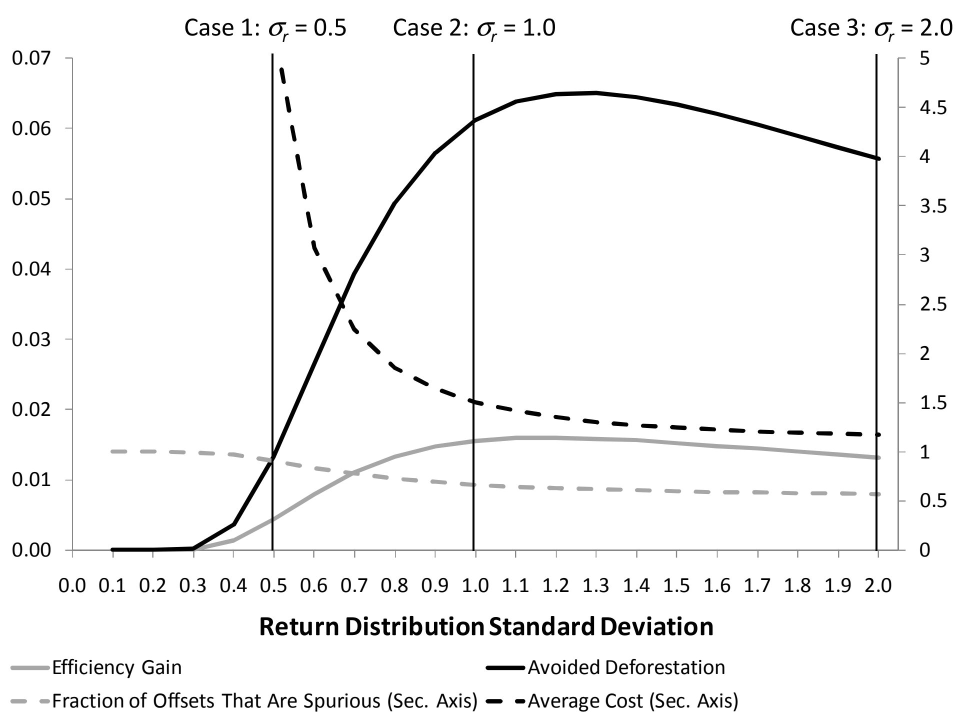

Figure 4: Impact of changing $f_r(r)$ on the policy objectives, with $p_c = \delta = 0.5$ and $\sigma_\varepsilon = 0.5$.

Moving from case 1 to case 2 unambiguously raises efficiency and lowers AC. This follows from the statement below Proposition 1, since $f'_r(r) > f_r(r)$ for all r between $[0, \delta]$. It corresponds to a downward movement in the marginal cost curve. The fraction of offsets that are spurious also falls from 91% to 67%. In contrast, moving from case 2 to case 3, efficiency falls slightly; AC does also. The probability mass of returns between $[0, \delta]$ decreases slightly, limiting the potential efficiency gains and AD. The distribution becomes almost flat in the region $[0, \delta]$. This means that the density close to zero (where abatement costs are low) falls relative to the density close to δ (where abatement costs are high). This flatness however also means that the ratio of land with returns at risk of infra-marginal payments (r just below 0) to returns with potential efficiency gains (r between $[0, \delta]$) is lower: the fraction of spurious offsets keeps decreasing (from 67% to 57%), as does AC.

A shift in the returns distribution that implies consistently lower marginal abatement costs in the relevant price range and reduces the density of returns just below zero relative to those above zero will improve efficiency, value (reduced AC) and quality (lower FOS). Governments may try to achieve such a shift in the return distribution (as perceived by landowners) by adopting policies complementary to the voluntary program that address information failures or non-carbon externalities and hence increase the attractiveness of keeping low productivity land forested (e.g. for tourism, or a sustainable form of selective logging) or that reduce the attractiveness of agriculture on marginal land (Angelsen, 2010).

3 The Impact of Policy Choices

Governments have several policy options at their disposal to design a voluntary avoided deforestation program. We analyze three policy options: increasing the project scale, offset price discounting and changing the assigned forest baseline.

3.1 Policy 1: increasing the project scale

A first policy to consider is to increase the scale of each project. So far, we have considered a small-scale policy in which landowners get assigned plot-specific baselines and can opt in separately with each individual plot. While some forest carbon programs in practice are indeed small-scale, other proposals feature baselines for larger areas (e.g., a region or a country).[18] Larger programs devolve responsibility for changing individual landowners' behavior from the industrialized country offset buyer to large local entities that may, in addition to the benefits of scale discussed in this section, have more authority and better information to enable efficient developing country policy. Section 2

[18]The Costa Rican Payments for Ecosystems services program is an example of a small scale system. Norway's recent performance based agreement with Brazil sets up a large scale system: http://www.norway.org/ARCHIVE/policy/environment/regnskogen_i_brasil_en/

showed that observation errors in voluntary programs reduce efficiency and avoided deforestation, value and quality. This section shows that increasing the required scale of each project in the program mitigates these adverse consequences.

3.1.1 A multiple-plot model

We now consider a single entity (a large landowner or alternatively, a region or country) which controls N 1-hectare plots. Each plot j has a return from deforestation r_j. We initially assume that these returns are distributed i.i.d. over plots with density f_r. Without the program, the entity will clear all plots for which the return r_j exceeds zero. Hence, the true baseline is

$$BL_N = \sum_{j=1}^{N} BL_j \ where \ BL_j = \left\{ \begin{array}{ll} 1 & if \ r_j \leq 0 \\ 0 & if \ r_j > 0 \end{array} \right\} \tag{11}$$

The government observes each r_j with error ε_j: $\widehat{r}_j = r_j + \varepsilon_j$. Assume that ε_j is i.i.d. across j. This means that $\widehat{r}_j$ has a distribution with mean μ_r and variance $\sigma_r^2 + \sigma_\varepsilon^2$. The distribution of $\widehat{r}_j$ is more dispersed than $f_r(r)$. The government could compute an unbiased prediction of the baseline $\widehat{BL}_N$ as the sum of the expectation of the random variables for the plot-specific baselines. From its point of view, the true baseline for a specific plot is a Bernoulli random variable with mean p_{1i} and variance $p_{1i}(1-p_{1i})$, where $p_{1i} = Pr(r_j < 0|\widehat{r}_j) = Pr(BL_j = 1|\widehat{r}_j)$.[19] Since these are non-identically but independently distributed across j, the central limit theorem yields that for $N \to \infty$

$$\widehat{BL}_N = \sum_{j=1}^{N} \widehat{BL}_j \ \xrightarrow{d} N\left(\sum_{j=1}^{N} p_{1j}, \sum_{j=1}^{N} p_{1j}\left(1-p_{1j}\right)\right) \tag{12}$$

where $\widehat{BL}_N$ is a cumulative baseline for all N plots.

3.1.2 Increasing scale and efficiency

With the N-plot baseline, the entity that controls the area (which could be a local or national government) must decide whether or not to opt in with his entire forest area, or not participate. The difference with the single-plot model is illustrated by Figure 5.

Figure 5 contrasts the single plot with the multiple plot case. In the single plot case, an inefficiency occurs when the true baseline is 0, but the government assigns a baseline of 1. In the

[19] Note that $p_{1i} \neq Pr(\widehat{r}_j < 0)$, except if $f_r(r)$ is symmetric around zero. If the government naively assumed that r and $\widehat{r}$ have the same distribution, it would calculate $p_{1i} = F_\varepsilon(-\widehat{r}_j) =$ (if f_ε is symmetric) $1 - F_\varepsilon(\widehat{r}_j)$. This would lead to a biased estimate of the baseline. Consider $f_r(r) \sim N(-1,1)$ and $f_\varepsilon(\varepsilon) \sim N(0,1)$. In that case, $f_{\widehat{r}}(\widehat{r}) \sim N(-1,2)$. The probability that $\widehat{r} > 0$ exceeds the probability that the true return $r > 0$. Therefore, if the government used a bottom-up plot-level to estimate $\widehat{r}$ and assigned a zero baseline for all plots with positive $\widehat{r}$, the baseline would be biased downwards.

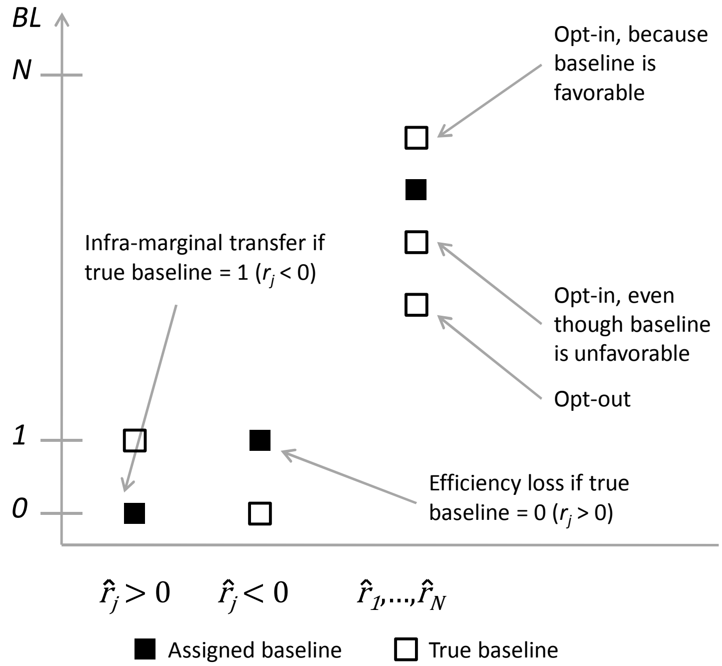

Figure 5: Single versus multiple plot policy.

multi-plot case, assigning a more favorable baseline ($\widehat{BL}_N < BL_N$) will lead to guaranteed opt-in and infra-marginal payments. However, if (($\widehat{BL}_N > BL_N$), the entity has two options. If it opts in, it will clear all plots with returns exceeding $p_c = \delta$, but forego clearing plots with returns between 0 and p_c. Let N_{p_c} be the number of plots with $r < p_c$. Hence, opting in is favorable if and only if

$$p_c \left(N_{p_c} - \widehat{BL}_N \right) > \sum_{j | r_j \in [0, p_c]} r_j \tag{13}$$

Hence, for some assigned baselines $\widehat{BL}_N$, the entity will still opt in, but for assigned baselines exceeding a threshold value, the entity will opt out with all of its N plots. There are cases in which scale increases efficiency: even if the baseline is too stringent, the entity will still opt in *with all N plots*. Hence, all plots with returns between 0 and p_c will remain forested. This leads to a higher efficiency gain than plot-specific baselines, in which some plots with returns between 0 and p_c will get assigned a baseline equal to 1, and opt out. However, in some cases the efficiency of the new system is lower than with plot-specific baselines. This happens when the baseline is so unfavorable that the entity opts out *with all N plots*. In the single plot program, some critical plots with returns between 0 and p_c will receive correct baseline and some deforestation will be efficiently avoided.

By the law of large numbers, as $N \to \infty, = \frac{\widehat{BL}_N}{N} \to \frac{BL_N}{N}$: the standard error of the average baseline *per plot* goes to zero. However, the standard error of $\widehat{BL}_N$ does not converge to zero.

Therefore, it is possible that the entity gets assigned a baseline that it so unfavorable that it decides to opt out with all N plots. Since this standard error only grows at rate $\sqrt{n}$ while the expected benefit from program participation grows at rate n, the probability of opt-in approaches 1 as $N \to \infty$ and the efficient solution will be obtained.

In the limit, larger scale will lead to the same efficient outcome as under the full information voluntary program. However, real-world programs can only be scaled up to a finite number of plots.[20] We therefore explore the effects of moderate increases in scale numerically in the next section.

3.1.3 Numerical simulations of increased scale

This section presents numerical simulations to illustrate the differences between a single-plot versus a multiple plot program. Throughout this section, we assume $f_\varepsilon(\varepsilon) \sim N(0, 0.5)$, unbiased assigned baselines ($E[(\widehat{BL}_N] = BL_N)$, and $p_c = \delta = 0.5$. The central case returns distribution is $f_r(r) \sim N(-1, 1)$, but we also consider alternative distributions. Table 1 demonstrates what happens to the policy objectives as N increases.

Table 1: The impact of increasing the required project scale: $f_r(r) \sim N(-1, 1)$.

	1-plot	2-plot	10-plot	100-plot	Maximum efficiency: $\sigma_\varepsilon = 0$
	(1)	(2)	(3)	(4)	(5)
Efficiency gain	1.33	1.42	1.82	2.47	2.53
AC	0.91	0.87	0.68	0.51	0.50
FOS	45.29%	42.53%	26.83%	1.86%	0%
AD	5.34	5.54	6.60	8.95	9.18
TT	4.88	4.82	4.51	4.56	4.59
Opt-in	4.88%	18.30%	58.42%	97.47%	100%

Notes: Each row is based on 10,000 random draws from the probability distributions. $\delta = p_c = 0.5$. $f_\varepsilon(\varepsilon) \sim N(0, 0.5)$. Efficiency gains, AD and TT are all normalized per 100 plots. Assigned baselines are unbiased.

Table 1 shows that increasing the project scale has dramatic consequences for its performance. For the central case, N-plot baselines increase efficiency and AD and reduce AC as N increases.

[20] The next section illustrates that, for *finite* N and certain unusual land return distributions, it is possible that efficiency initially *decreases* with N.

100 plots are enough to approach the efficient solution. The reason is that the observation error *normalized per plot* decreases as N grows, and the probability of opt-in becomes very high (97.47%).[21] This high opt-in rate signals efficiency though Table 1 shows that most gains are achieved through the first 5% of plots that participate. Hence, scale mitigates adverse selection for the central case returns distribution.

The central case returns distribution reflects that in most developing countries, the majority of the forested land is not at threat of deforestation, at least in the short to medium run. Still, we test the robustness of the result by analyzing the effects of project scale for two other return distributions: a $N(0,1)$ distribution (which implies that 50% of the forest will be cleared absent any policy) and a symmetric bimodal normal $BMN(0.5, 0.1)$ distribution with modes at -0.5 and 0.5 and standard deviation $\sigma_r = 0.1$. The latter distribution is unlikely to represent reality, but illustrates that - for finite N - efficiency does not monotonically increase in N.

Figure 6 summarizes the effects on efficiency and FOS of increasing project scale and compares it to the efficient ($\sigma_\varepsilon = 0$) solution. For the $N(0,1)$ distribution, the effects are similar to the central case distribution. Both efficiency and FOS improve with scale. The $BMN(0.5, 0.1)$ distribution demonstrates that increasing scale does not monotonically increase efficiency for all distributions. The intuition is that there are many plots with returns around $p_c = \delta$ (as well as returns close to $-\delta$). Therefore, the $BMN(0,5,0.1)$ distribution has many realizations for which (13) holds only if the baseline is correct or more favorable. A slight baseline error will cause the entity to opt out with all N plots. This effect dominates for small N: AD and efficiency *decrease* with N between N = 1 and 10. However, such distributions are highly stylized and unlikely to represent true returns distributions. Average cost and the fraction of offsets that are spurious decreases with scale in each case although they are also theoretically ambiguous for finite N.[22]

We now test the robustness of these conclusions in another way. An important assumption has been that $f_r(r)$ and $f_\varepsilon(\varepsilon)$ are i.i.d: both returns and observation errors are independent across plots. In reality, there may be a high degree of spatial correlation in both returns and errors. We introduce spatial correlation across plots in the following stylized way:

$$\begin{aligned} r_j &= \rho_r r_{j-1} + u_r \\ \varepsilon_j &= \rho_\varepsilon \varepsilon_{j-1} + u_\varepsilon \end{aligned} \tag{14}$$

where u_r and u_ε are i.i.d. with variances σ^2_{ur} and $\sigma^2_{u\varepsilon}$, such that $\sigma^2_r = \sigma^2_{ur}/(1-\rho^2_r)$ and

[21] In some realizations, landowners efficiently opt out: their return exceeds δ.

[22] Consider the following example with three plots. For plot 1: $BL = 1$, $\widehat{BL} = 0$ (spurious offset). For plot 2: $BL = 0$, $\widehat{BL} = 1$ and $r < p_c = \delta$. For plot 3: $BL = 0$ and $\widehat{BL} = 0$ and $r < p_c = \delta$. Under a single-plot policy, 1 and 3 opt in, leading to 1 spurious credit and 1 real offset. Now consider a policy in which the participation decision needs to be made for plots 2 and 3 together; plot 1 remains standalone. Under this larger-scale policy, plots 2 and 3 get an assigned baseline $\widehat{BL}_2 = 1$. If $r_2 + r_3 > p_c$, the entity will opt out with both plots. Plot 1 still opts in. This means there are only spurious offsets now: scale adversely impacts both FOS and AC.

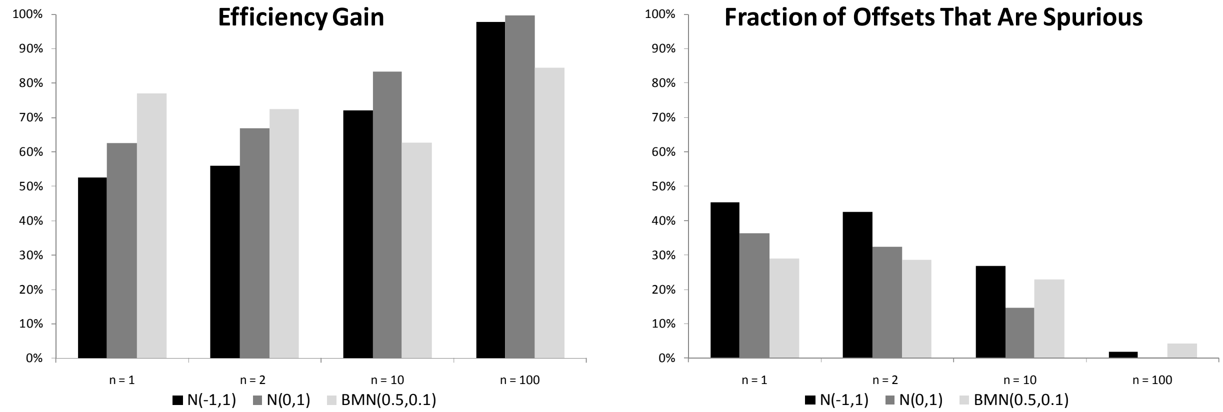

Figure 6: The impact on efficiency (left panel) and FOS (right panel) of increasing the project scale for alternative return distributions, for $N = 1$, 2, 10 and 100.

$\sigma_\varepsilon^2 = \sigma_{u\varepsilon}^2/(1-\rho_\varepsilon^2)$. Table 2 summarizes the main findings for the central case.[23]

Table 2 shows that as the correlation across plots and errors increases, efficiency and AD decrease, and AC and FOS increase. The intuition is that observation errors do not cancel out across plots, but are persistent. High spatial correlation reduces the probability of participation for a given N and therefore adversely impacts the policy objectives. A larger required project scale would mitigate the effects of spatial correlation but will not eliminate it if the correlation is driven partly by common unobservable factors at the national scale.

3.2 Policy 2: discounting the payment per hectare

In the analysis above, the payment per hectare p_c was assumed to be equal to the marginal externality from deforestation δ. This section analyzes what happens to the three project objectives if we vary p_c. Reducing p_c is equivalent to the practice of "offset discounting" sometimes observed in practice in offset systems (Kollmuss et al., 2010). Under a system of discounting, fewer offsets are awarded than the environmental gains represented by the difference between the baseline and the actual forest level. Many people promote this as a way to correct for spurious offsets.[24] First, we discuss the impact of changing p_c in a single-plot model. Then, we analyze how these results change in a multiple-plot model using numerical simulations.

It is straightforward that, independent of scale, any $p_c \neq \delta$ is less efficient if $\sigma_\varepsilon = 0$. All entities get assigned the true baseline, and paying less than δ reduces efficiency, because entities with average returns $((\underline{r}|0 \leq r \leq \delta) > p_c)$ would opt out. Paying more than δ (a "premium"

[23] Results for the $N(0,1)$ distribution are similar. The efficiency loss from high spatial correlation is relatively lower compared to the central case.

[24] Schneider (2009a and 2009b), Chung (2007), Environmental Defense Fund (2007), and Greenpeace and Papua New Guinea (2008) all propose discounting as a way to reduce non-additionality or spurious offsets.

Table 2: The impact of spatially correlated returns and observation errors: $N = 1$ and $N = 100$. $f_r(r) \sim N(-1, 1)$.

	1-plot	100-plot	100-plot $\rho_r = \rho_\varepsilon = 0$	100-plot $\rho_r = \rho_\varepsilon = 0.5$	100-plot $\rho_r = \rho_\varepsilon = 0.9$	$\sigma_\varepsilon = 0$ $\rho_r = \rho_\varepsilon = 0.99$
	(1)	(2)	(3)	(4)	(5)	(6)
Efficiency gain	1.33	2.47	2.41	2.02	1.59	2.53
AC	0.91	0.51	0.52	0.61	0.80	0.50
FOS	45.29%	1.86%	3.54%	17.76%	37.69%	0%
AD	5.34	8.95	8.72	7.27	5.72	9.18
TT	4.88	4.56	4.52	4.42	4.59	4.59
Opt-in	4.88%	97.47%	94.96%	78.40%	37.28%	100%

Notes: Each row is based on 10,000 random draws from the probability distributions $\delta = p_c = 0.5$. $f_\varepsilon(\varepsilon) \sim N(0, 0.5)$. Efficiency gains, AD and TT are all normalized per 100 plots. Assigned baselines are unbiased.

rather than a "discount") reduces efficiency because some entities will opt in even though their private gains from deforestation exceed the full environmental cost. This is inefficient from an economic perspective. In the single-plot model with full and symmetric information, the change in efficiency relative to a no policy case was given in (1). A simple application of Leibniz' Rule yields that efficiency is maximized when $p_c = \delta$. We will now investigate if this result changes with asymmetric information - i.e. when $\sigma_\varepsilon > 0$.

3.2.1 Discounting in the single-plot model

In the single-plot model, the introduction of observation error does not change the conclusion that the most efficient payment is $p_c = \delta$. The efficiency change relative to no policy equals

$$\Delta S(p_c) = \int_0^{p_c} (\delta - r) \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr \tag{15}$$

Proposition 2. In the single-plot model, efficiency is maximized for $p_c = \delta$, regardless of $f_\varepsilon(\varepsilon)$.

Proof. The first order condition is given by $\frac{d(\Delta S(p_c))}{dp_c} = (\delta - p_c) \left(\int_{-p_c}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(p_c)$, using Leibniz' Rule. Since $\int_{-p_c}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon > 0$ for any $f_\varepsilon(\varepsilon)$ and $f_r(p_c) \geq 0$, efficiency is maximized when $p_c = \delta$. ■

We now investigate what happens to the other policy objectives as the payment p_c varies.

Proposition 3. AD, MT, IT, and TT are globally (weakly) increasing in p_c; FOS is globally (weakly) decreasing in p_c.

Proof. $AD = \int_0^{p_c} \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr$. The derivative of AD w.r.t. p_c is $\left(\int_{-p_c}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(p_c) \geq 0\ \forall p_c$, proving the first statement. The derivative of MT (first term in (9)) w.r.t. p_c is $\int_0^{p_c} \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr + p_c \left(\int_{-p_c}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(p_c) \geq 0\ \forall p_c$. The derivative of IT (second term in (9)) w.r.t. p_c is $\int_{-\infty}^{0} \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr \geq 0\ \forall p_c$. Hence, the derivative of TT w.r.t. p_c is weakly greater than zero $\forall p_c$. Finally, using (9), $FOS = IT/TT = \int_{-\infty}^{0} \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr / \int_{-\infty}^{p_c} \left(\int_{-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr$. Since the denominator is monotonically (weakly) increasing in p_c, FOS is monotonically (weakly) decreasing in p_c. ■

Proposition 3 demonstrates that, contrary to the intended effect, the fraction of offsets that is spurious, FOS, increases when the offset price is discounted (i.e., reduced): as p_c falls the share of offsets that is spurious rises toward 1.

The effect of changing p_c on AC is ambiguous. Since AD is bounded, very high values of p_c will lead to increasing AC. For intermediate values of p_c, a small increase in p_c can either lead to almost no additional deforestation, or a large increase in avoided deforestation, depending on the specification of the return distribution $f_r(r)$. For instance, if $f_r(r) = 0$ for $r \in [0, \underline{p}]$, then AC will be infinite for $p_c \leq \underline{p}$ and achieve a global minimum for some $p_c > \underline{p}$. Hence, AC can either be increasing or decreasing in p_c.

We conclude that, in the single-plot model, efficiency is maximized by paying $p_c = \delta$. Paying more reduces FOS, leads to more AD, but requires higher transfers. The effect on AC is ambiguous for low values of p_c, but eventually AC must increase.

3.2.2 Discounting in the multi-plot model

In the multi-plot model, $p_c = \delta$ no longer unambiguously maximizes efficiency. The intuition is as follows. Raising p_c above δ has two countervailing effects on efficiency. First, it will increase the opt-in probability. This increases efficiency because it helps prevent deforestation of plots with returns below δ. Second, it causes certain forest to be *inefficiently* prevented from deforestation. The relative strength of these channels determines whether a higher p_c can be more efficient than $p_c = \delta$. A lower p_c will never increase efficiency, since it will both reduce opt-in and cause inefficient deforestation. The effects go in the same direction. Hence, $p_c \geq \delta$ maximizes efficiency in the multiple-plot model. Figure 7 illustrates this when $\delta = 0.5$ and $N = 10$ or 100, and also shows the impact of discounting on other criteria.

Figure 7 shows that raising p_c above δ can increase efficiency. For $N = 10$, raising p_c above δ (to $p_c = 0.6$) slightly increases efficiency. Hence, efficiency is no longer maximized at $p_c = \delta$. However, when $N = 100$, the opt-in probability at $p_c = \delta$ is already almost efficient at 97.47%. Raising p_c to 0.6 increases opt-in only slightly to 98.93%. Hence, we find that the most efficient solution is sometimes achieved for $p_c > \delta$. This increased efficiency coincides with higher AC, however. The figure also shows the beneficial effects of increased scale.

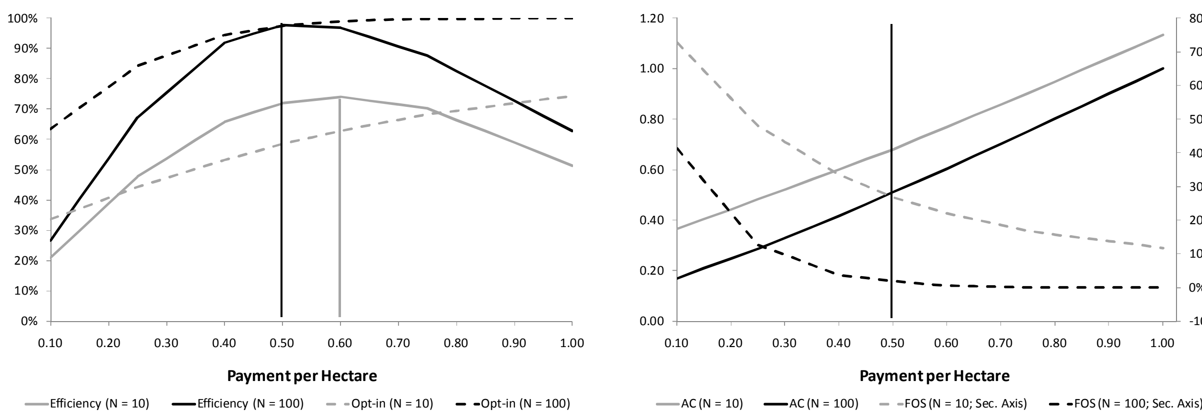

Figure 7: The impact of changing p_c and project scale on efficiency and opt-in (left panel) and on AC and FOS (right panel), for $f_r(r) \sim N(-1, 1)$, $\delta = 0.5$, $N = 10$ and 100.

In summary, we find that efficiency considerations never justify discounting the payment p_c below the environmental damage δ. Setting $p_c > \delta$ can be justified if opt-in at $p_c = \delta$ is below 100%, but this efficiency increase comes at the expense of higher average cost. Since increasing scale leads to full opt-in in the limit, $p_c = \delta$ always becomes the most efficient payment as N approaches infinity. There always exists a p' ($p' > \delta$) such that efficiency falls and transfers rise for any $p > p'$. In that price region, efficiency decreases while transfers increase. For choices $p_c > \delta$, the tradeoff between efficiency and average cost remains.

Analogous to Proposition 3 for the single-plot model, FOS (and opt-in) unambiguously become more favorable as p_c increases. If an entity had a favorable baseline at p_c, it would have chosen to opt in. At a higher price $p_c + \epsilon$, the entity will still opt in but no additional spurious offsets will be generated. Hence, raising the price does not increase the number of spurious offsets while it does increase AD, leading to a reduced FOS.

Hence, we conclude that restoring environmental integrity of offsets is *not* a valid reason to advocate offset discounting. Discounting does however, almost always, reduce the average cost of offsets. If the discounted price is achieved by limiting demand (e.g., limiting the number of offsets that can enter the market), so that buyers pay less than the market price for a unit that is then fully fungible with other units (i.e., a 1:1 trading ratio without downward adjustment of the cap, as is the case in the Clean Development Mechanism), buyers will reap gains and the environmental outcome will be negative. If however the gains to industrialized countries are spent on additional mitigation (as it is with $(t:1, t>1)$ trading ratios) this could have a positive environmental effect even if the cap is not adjusted.[25] This might however be more efficiently achieved through changes in baselines.

3.3 Policy 3: changing the generosity of the assigned baseline

Another policy choice for the regulator is to set a baseline that is, in expectation, too high or too low. In other words, the government assigns the following baselines for plot i

$$\widehat{BL}_i = \left\{ \begin{array}{ll} 1 & if\ \widehat{r}_i \leq r^* \\ 0 & if\ \widehat{r}_i > r^* \end{array} \right\} \tag{16}$$

where r^* is a specified return set by the government. The government, aware of adverse selection, may try to pay only landowners who are most likely to deforest in the baseline, for instance by choosing $p_c > r^* > 0$. Assuming $p_c = \delta$, we analyze the impact of this policy change on the various criteria: efficiency, AD, AC and FOS. To provide intuition, we first discuss the impact in the context of the single-plot model. Then, we present numerical simulations of the multiple-plot model.

3.3.1 Changing baselines in the single-plot model

Proposition 4. More generous baselines (weakly) increase efficiency and AD, but require a (weakly) higher TT.

Proof. The efficiency change relative to no policy equals $\Delta S(r^*) = \int\limits_0^{p_c} (p_c - r) \left(\int\limits_{r^*-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr$. By Leibniz' Rule, this expression is globally weakly decreasing in r^*. Hence, efficiency is maximized if $r^* \to -\infty$. AD is given by $AD(r^*) = Pr(0 \leq r \leq p_c, \widehat{r} > r^*) = \int\limits_0^{p_c} \left(\int\limits_{r^*-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr$, which is globally weakly decreasing in r^*. TT is given by $TT(r^*) = p_c \int\limits_{-\infty}^{p_c} \left(\int\limits_{r^*-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr$, which is also globally weakly decreasing in r^*. ■

[25] Note that, if trading ratios are used without adjusting the cap, the environmental effect could be negative even for large t. A straightforward example is a returns distribution with positive probability mass below zero, but no probability mass between 0 and p_c.

The effect on FOS is theoretically ambiguous and depends on the baseline error distribution. Using (9) and (10), $FOS = OS/AD = IT/TT$ is decreasing in r^* if and only if MT/IT is increasing in r^*. We can write $MT/IT = \int_0^{p_c} \left(\int_{r^*-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr / \int_{-\infty}^{0} \left(\int_{r^*-r}^{\infty} f_\varepsilon(\varepsilon)\, d\varepsilon \right) f_r(r)\, dr = \int_0^{p_c} (1 - F_\varepsilon(r^* - r))\, f_r(r)\, dr / \int_{-\infty}^{0} (1 - F_\varepsilon(r^* - r))\, f_r(r)\, dr$. If this expression is increasing in r^*, FOS is decreasing in baseline stringency. This condition will certainly hold if the baseline error is bounded from above.

The fact that efficiency increases as the baseline becomes more generous is not surprising, since in the limit this is equivalent to assigning a no-forest baseline or a subsidy of p_c per hectare of forest standing. As discussed in Section 2, such a subsidy is indeed efficient but requires a large infra-marginal transfer.

Using (10) and making OS and AD functions of r^* we can see that the effect of r^* on AC is also ambiguous. OS, the amount of spurious offsets, is decreasing in r^*, but so is AD. The shape of AC is dependent on the return distribution $f_r(r)$.

Numerical illustration

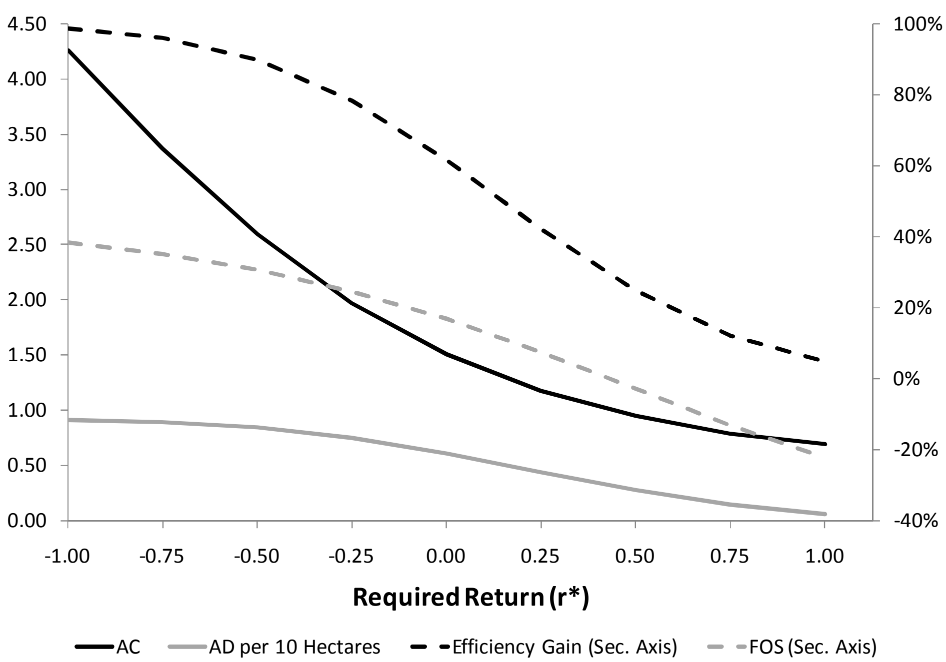

Figure 8: The impact of baseline generosity on the project objectives for the central case.

Figure 8 shows that for reasonable returns distributions like our central case, FOS and AC are both decreasing in r^* (as baselines become more stringent). For very stringent baselines FOS

becomes negative: the environmental gains are greater than the number of traded offsets. Efficiency, lower average cost and offset quality are conflicting policy aims for this policy option also: efficiency requires setting r^* low (generous baseline), while minimizing average cost and maximizing offset quality requires setting r^* high (stringent baseline).

3.3.2 Changing baselines in the multiple-plot model

The conclusions from the single-plot model also hold in the multiple-plot model. Figure 9 illustrates the effect of assigning baselines that are too (un)favorable in expectation for the central case returns distribution. The true baseline equals 84 (84 out of 100 plots will remain forested in absence of a policy). The figure shows that increasing the baseline (i.e., making it less favorable) unambiguously reduces efficiency and AD, but also reduces AC and FOS.

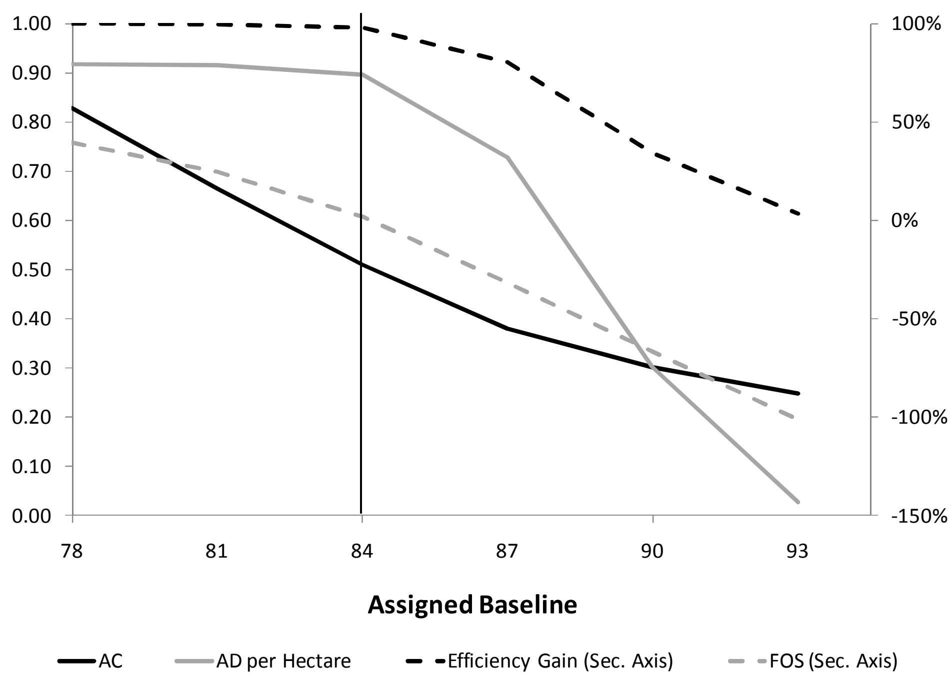

Figure 9: The impact of changing baseline generosity on the project objectives, for $f_r(r) \sim N(-1,1)$ and $N = 100$. 84 is an unbiased baseline.

This section has shown that only increasing project scale improves all objectives simultaneously (for "typical" returns distributions). Discounting offsets and changing baseline generosity affect the objectives in opposing directions. Discounting offsets reduces efficiency, AD and offset quality, but improves the value for money for funders. Making assigned baselines more stringent reduces efficiency and AD, but improves quality and value. This illustrates the conflicting nature of these policy objectives. It also illustrates that tightening the baseline should be favored over offset discounting if environmental integrity of offsets is a key policy concern and not enough of the gains to industrialized countries from discounting are spent on additional mitigation (e.g. through trading ratios and/or reducing the cap) to counteract the fall in offset quality.

Table 3 provides a summary of the impact of the various policy options on the policy objectives discussed in this paper.

Table 3: The effects of the various policy options on the policy objectives.

	Policy option		**Policy criteria**		
	Required scale of project (1)	Maximize efficiency (2)	Maximize offset quality (minimize FOS) (3)	Maximize value for money (minimize AC) (4)	Maximize avoided deforestation (5)
Increase scale	Finite N	(+)	(+)	(+)	(+)
	$N \to \infty$	+	+	+	+
Raise price above δ	$N = 1$	−	+	(−)	+
	$N > 1$	First +, then −	+	(−)	+
Lower price below δ ("discount")	Any	−	−	(+)	−
Generous baseline	Any	+	(−)	(−)	+
Stringent baseline	Any	−	(+)	(+)	−

Notes: + indicates a favorable effect; - indicates an unfavorable effect; round brackets indicate that the effect holds for "reasonable" distributions but is theoretically ambiguous.

4 Conclusion and Policy Implications

This paper built a model of landowner decisions and a voluntary avoided deforestation program. It demonstrated the tradeoff between three key policy criteria: efficiency, value (minimizing average cost to industrialized countries) and quality of offsets, when there is asymmetric information. It analyzed the effects of increasing the scale of required projects including when returns are spatially correlated. It then explored the effects on the three policy criteria of two other policy levers: offset discounting (reducing the payment per hectare to below the value of the environmental externality) and changing the baseline.

We have four main findings. First, under almost all circumstances, voluntary deforestation programs (or, in fact, general offset programs) will perform better with increased required scale of project. Second, offset discounting and setting more stringent baselines highlight the tradeoffs involved in policy design: efficient policy may involve high transfers that make the policy unattractive to the industrialized countries which will fund them. Both policies reduce efficiency but generally raise value. Moreover, there is still an efficiency gain relative to no policy. Third, discounting lowers the quality of offsets and does not improve the environmental outcome, at least if not accompanied by a sufficiently high trading ratio. In some cases, even a high trading ratio leads to a worse environmental outcome. Therefore, the main rationale for offset discounting is to raise the value of the policy to industrialized countries. Fourth, making baselines more stringent does increase the quality of offsets.

Our key messages for policy makers are three. First, make 'projects' as large as possible. Regional or national scale programs where funds or offsets are transferred to the government on the basis of aggregated regional or national monitoring data will be much more efficient and offer better value for money. Although baseline deforestation rates are still difficult to predict at the national or regional scale, the errors fall dramatically relative to small scale prediction.

Second, recognize that the primary purpose of offset discounting or below market prices is to reduce the cost to industrialized countries so that paying for avoided deforestation becomes an attractive mitigation option for them. It actually increases the share of funds that go to spurious offsets and reduces efficiency. Discounting can only be justified on environmental integrity grounds if the trading ratio with "regular" cap-and-trade credits compensates for the loss of offset quality. In that case, discounting can extract rents from sellers to pay for additional environmental protection. The use of baselines more stringent than business as usual typically does reduce the number and fraction of spurious offsets while simultaneously reducing the cost to industrialized countries, but reduces efficiency.

Third, invest in research to improve understanding of local and global deforestation drivers. This will allow more accurate assessment of returns distributions and their evolution and hence more accurate prediction of baseline deforestation. Moreover, this will help identify domestic policies to effectively control deforestation.

This paper has highlighted the tradeoffs involved in various policy design options for avoiding deforestation. In future work, we will present a framework in which industrialized and forest-covered developing countries can explore these trade-offs, subject to the restriction that the policy has to be individually rational for both parties. By defining the Pareto efficient bargaining set we can explore its determinants and give guidance to policy makers who may seek to expand the set and to negotiators who want to find agreements within the set (and in their countries' favor).

If countries can be encouraged to be more generous, by pushing less to lower the average cost of (real) offsets or by accepting more stringent baselines, it will be easier to create an efficient international framework to avoid deforestation. Combined with effective domestic policies that respond to the international incentives, this could meet the expectations of those who promote avoided deforestation as a key climate mitigation option in the short term.

References

- Andam, K., P. Ferraro, A. Pfaff, J. Robalino, and A. Sanchez. 2008. "Measuring the Effectiveness of Protected-Area Networks in Reducing Deforestation." *Proceedings of the National Academy of Sciences* 105(42): 16089-16094.

- Angelsen, A. 2010. "Policies for reduced deforestation and their impact on agricultural production", in Climate Mitigation and Agricultural Productivity in Tropical Landscapes, Special Feature, *Proceedings of the National Academy of Sciences*. Accessed online on 19 July 2010.

- Angelsen, A. (ed.). 2008. "Moving ahead with REDD: Issues, options and implications." CIFOR, Bogor, Indonesia.

- Arguedas, C. and D.P. van Soest. 2009. "On reducing the windfall profits in environmental subsidy programs." *Journal of Environmental Economics and Management* 58(2): 192-205.

- Busch J., B. Strassburg, A. Cattaneo, R. Lubowski, A. Bruner, R. Rice, A. Creed, R. Ashton, and F. Boltz. 2009. "Comparing climate and cost impacts of reference levels for reducing emissions from deforestation." *Environmental Research Letters* 4(4).

- Chomitz, K. 2007. "At Loggerheads? Agricultural Expansion, Poverty Reduction, and Environment in the Tropics." World Bank: Washington, DC.

- Environmental Defense Fund. 2007. "CDM and the Post-2012 Framework." Discussion Paper, Vienna, August 27-31, AWG/Dialogue.

- Fischer, C. 2005. "Project-based mechanisms for emissions reductions: balancing trade-offs with baselines." *Energy Policy* 33(14): 1807-1823.

- Government of the Kingdom of Norway and Government of the Republic of Indonesia. 2010. Letter of Intent on "Cooperation on reducing greenhouse gas emissions from deforestation and forest degradation."

- Greenpeace and Papua New Guinea. 2008. REDD proposals submitted to the United Nations. Available at http://unfccc.int.

- He, G. and R. Morse. 2010. "Making Carbon Offsets Work in the Developing World: Lessons from the Chinese Wind Controversy." Stanford University, Program on Energy and Sustainable Development Working Paper 90.

- Kerr, S. 1995. "Adverse Selection and Participation in International Environmental Agreements", in *Contracts and Tradeable Permit Markets in International and Domestic Environmental Protection*, Ph.D. Thesis, Harvard University.

- Kerr, S., A.S.P. Pfaff, and A. Sanchez. 2002. "The dynamics of deforestation: evidence from Costa Rica." Motu Manuscript. Available at http://www.motu.org.nz/files/docs/MEL0270.pdf.
- Kerr, S., J. Hendy, S. Liu and A.S.P. Pfaff. 2004. "Tropical forest protection, uncertainty and the environmental integrity of carbon mitigation policies." Motu Working Paper 04-03.
- Kindermann, G.E., M. Obersteiner, E. Rametsteiner, and I. McCallum. 2006. "Predicting the deforestation trend under different carbon prices." *Carbon Balance Management* 1(15).
- Kindermann, G.E., M. Obersteiner, B. Sohngen, J. Sathaye, K. Andrasko, E. Rametsteiner, B. Schlamadinger, S. Wunder, and R. Beach. 2008. "Global Cost Estimates of Reducing Carbon Emissions Through Avoided Deforestation." *Proceedings of the National Academy of Sciences* 105(30): 10302-10307.
- Kollmuss, A., M. Lazarus, and G. Smith. 2010. "Discounting Offsets: Issues and Options." Stockholm Environment Institute Working Paper WP-US-1005.
- Mason, C.F. and A.J. Plantinga. 2010. "The Additionality Problem with Offsets: Optimal Contracts for Carbon Sequestration in Forests." Working paper. Available at http://idei.fr/doc/conf /ere/papers_2010/paper_Mason.pdf.
- Melillo, J.M., J.M. Reilly, D.W. Kicklighter, A.C. Gurgel, T.W. Cronin, S. Paltsev, B.S. Felzer, X. Wang, A.P. Sokolov, and C.A. Schlosser. 2009. "Indirect Emissions from Biofuels: How Important?" *Science* 326: 1397-1399.
- Millard-Ball, A. 2010. "Adverse Selection in an Opt-In Emissions Trading Program: The case of sectoral no-lose targets for transportation." Stanford University, Program on Energy and Sustainable Development Working Paper 97.
- Montero, J.P. 1999. "Voluntary Compliance with Market-Based Environmental Policy: Evidence from the U. S. Acid Rain Program." *Journal of Political Economy* 107(5): 998-1033.
- Montero, J.P. 2000. "Optimal design of a phase-in emissions trading program." *Journal of Public Economics* 75(2): 273-291.
- Pfaff, A., S. Kerr, L. Lipper, R. Cavatassi, B. Davis, J. Hendy, A. Sanchez-Azofeifa. 2007. "Will buying tropical forest carbon benefit the poor? Evidence from Costa Rica." *Land Use Policy* 24(3): 600-610.
- Plantinga, A.J. and K.R. Richards. 2008. "International Forest Carbon Sequestration in a Post-Kyoto Agreement." Discussion Paper 08-11, Harvard Project on International Climate Agreements, Belfer Center for Science and International Affairs, Harvard Kennedy School.
- Robalino, J., A. Pfaff, G. A. Sanchez-Azofeifa, F. Alpizar, C. Leon, and C. M. Rodriguez. 2008. "Deforestation Impacts of Environmental Services Payments: Costa Rica's PSA Program 2000-2005." Environment for Development Discussion Paper 08-24. Washington, DC: Resources for the Future.
- Sanchez-Azofeifa, A., A. Pfaff, J.A. Robalino, and J.P. Boomhower. 2007. "Costa Rica's Payment for Environmental Services Program: Intention, Implementation, and Impact." *Conservation Biology* 21(5): 1165-1173.

- Sathaye J., W. Makundi, L. Dale, P. Chan, and K. Andrasko. 2006. "GHG mitigation potential, costs, and benefits in global forests." *Energy Journal* 27(3): 127-162.

- Schneider, L. 2009a. "Assessing the additionality of CDM projects: practical experiences and lessons learned." *Climate Policy* 9(3): 242-254

- Schneider, L. 2009b. "A Clean Development Mechanism with global atmospheric benefits for a post-2012 climate regime." I*nternational Environmental Agreements: Politics, Law and Economics* 9(2): 95-111.

- Stern, N. 2006. *The economics of climate change: The Stern Review.* Cambridge University Press: Cambridge, MA.

- Strand, J. 2010. "Carbon Offsets with Endogenous Environmental Policy". World Bank Policy Research Working Paper 5296.

- Wara, M. and D. Victor. 2008. "A realistic policy on international carbon offsets." Stanford University, Program on Energy and Sustainable Development Working Paper 74.

- Wise, M., K. Calvin, A. Thomson, L. Clarke, B. Bond-Lamberty, R. Sands, S.J. Smith, A. Janetos, and J. Edmonds. 2009. "Implications of Limiting CO2 Concentrations for Land Use and Energy." *Science* 324: 1183-1186.