University of Pennsylvania
ScholarlyCommons

Business Economics and Public Policy Papers

Wharton Faculty Research

11-2013

Scale and Transfers in International Emissions Offset Programs

Arthur van Benthem
University of Pennsylvania

Suzi Kerr

Follow this and additional works at: https://repository.upenn.edu/bepp_papers

Part of the [Public Affairs, Public Policy and Public Administration Commons](#), and the [Public Economics Commons](#)

Recommended Citation

van Benthem, A., & Kerr, S. (2013). Scale and Transfers in International Emissions Offset Programs. *Journal of Public Economics*, 107 31-46. <http://dx.doi.org/10.1016/j.jpubeco.2013.08.004>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/bepp_papers/123
For more information, please contact repository@pobox.upenn.edu.

Scale and Transfers in International Emissions Offset Programs

Abstract

Voluntary emissions offset programs between developing and industrialized countries suffer from adverse selection, because participants will self-select into the program. In contrast, pure subsidies for mitigation lead to full participation and hence efficiency, but require large financial transfers which make them unattractive to industrialized countries. We present a simple model to demonstrate the impact of three policy options on the performance of offset programs: (1) baseline scale increases, (2) offset discounting and (3) setting stringent baselines. With baseline scale increases, entire political jurisdictions such as regions or nations are assigned a single, aggregate baseline and must choose whether to participate as one entity. We find that increasing scale both improves efficiency and reduces transfers from offset buyers to sellers. Offset discounting means paying less than the value of abatement and can be paired with trading ratios between offsets and allowances in a cap-and-trade system. We show that discounting is inefficient, but can make offsets more attractive to industrialized countries. Setting stringent baselines also involves a tradeoff between efficiency and transfers. We finally show that Pareto efficient policies that are individually rational for buyers and sellers entail some combination of discounting and/or stringent baselines: offset policies are never first-best, but can be efficiency improving, especially with increased scale.

Keywords

offsets, deforestation, baselines, adverse selection, climate change policy, opt-in programs

Disciplines

Public Affairs, Public Policy and Public Administration | Public Economics

Scale and Transfers in International Emissions Offset Programs

Arthur van Benthem[†]

The Wharton School
University of Pennsylvania

Suzi Kerr[‡]

Motu Economic and Public Policy Research

Abstract

Voluntary emissions offset programs between developing and industrialized countries suffer from adverse selection, because participants will self-select into the program. In contrast, pure subsidies for mitigation lead to full participation and hence efficiency, but require large financial transfers which make them unattractive to industrialized countries. We present a simple model to demonstrate the impact of three policy options on the performance of offset programs: (1) baseline scale increases, (2) offset discounting and (3) setting stringent baselines. With baseline scale increases, entire political jurisdictions such as regions or nations are assigned a single, aggregate baseline and must choose whether to participate as one entity. We find that scale *both* improves efficiency *and* reduces transfers from offset buyers to sellers. Offset discounting means paying less than the value of abatement, or using trading ratios between offsets and allowances in a cap-and-trade system. We show that discounting is inefficient. While the conventional wisdom – that discounting can be used to reduce the fraction of offsets that are spurious – is invalid, discounting can make offsets more attractive to industrialized countries. Setting stringent baselines also involves a tradeoff between efficiency and transfers. We finally show that Pareto efficient policies that are individually rational for buyers and sellers entail some combination of discounting and/or stringent baselines: offset policies are never first-best, but can be efficiency improving, especially with increased scale. This paper frames the issues in terms of avoiding deforestation but the results are applicable to any voluntary offset program.

Keywords: offsets; deforestation; baselines; adverse selection; climate change policy; opt-in programs.

[†]Corresponding author. Department of Business Economics and Public Policy, The Wharton School, University of Pennsylvania, 3010 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, USA. Phone: (+1) 215 898 3013. Fax: (+1) 215 898 7635. Email: arthurv@wharton.upenn.edu.

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

1 Introduction

Deforestation is responsible for 15-25% of total greenhouse gas emissions, and including forests in climate negotiations appears crucial for the success of any potential climate deal. In fact, many leading academics and policy makers assert that avoiding deforestation is a key short-run, low-cost climate mitigation option (Stern, 2008; Kindermann et al., 2008).¹ A large literature emphasizes that it is critically important to price carbon in forests (Melillo et al., 2009; Wise et al., 2009), but does not take into account the difficulty of designing effective policies to address deforestation in developing countries, where most deforestation occurs (Andam et al., 2008; Pfaff et al., 2007). Many proposals assume the application of efficient price-based policies, yet these are hard to achieve in much of the developing world.

Offset programs, a key current instrument for industrialized countries to transfer resources in order to affect deforestation in developing countries, typically give credit for forest remaining above an estimated and assigned forest baseline.² Such programs suffer from serious problems such as adverse selection and spurious offsets. Adverse selection is caused by a combination of two factors: a *voluntary* element (i.e., forest owners can choose whether or not to opt in to the program) and *asymmetric information* about the baseline (i.e., the forest owners know more about their true baseline than the regulator).³ Spurious offsets occur if the regulator overestimates the baseline deforestation rate.

Are such offset programs doomed to be economically inefficient and, perhaps, undesirable, or can they be designed in ways to make them more efficient? This paper aims to answer this important policy question from a theoretical perspective.⁴ We formally model a voluntary price-based offset program to avoid deforestation, and examine the impact of three key policy levers on the program's economic performance, attractiveness to industrialized and developing countries as a mitigation option, and environmental outcome. The first policy is to increase the scale of offset programs so that entire political jurisdictions such as regions or nations get assigned a *single, aggregate* baseline and must choose whether to voluntarily participate in an international agreement as one entity with *all* its forested land. This is in sharp contrast to small scale, plot-specific baselines and opt-in rules for small local agents (e.g., individual landowners). The second policy is to discount offsets (pay less than the value of abatement), or use trading ratios between offsets and allowances in a

¹Current estimates of the forest carbon supply curve are based on land use responses to commodity prices (Sathaye et al., 2006; Melillo et al., 2009) or on estimates of the opportunity cost of land (e.g. Kindermann et al., 2008).

²In this paper, we focus on offset programs in which buyers and sellers reside in different countries. Buyer countries have committed to reduce emissions, and are currently mostly industrialized countries. In principle, developing countries could become offset buyers if they commit to emissions targets in an international climate agreement.

³Montero (2000), Fischer (2005), Arguedas and Van Soest (2009) and Bushnell (2011) establish theoretical results for the effect of adverse selection on offset programs. Montero (1999) presents empirical evidence for the US acid rain program. Zhang and Wang (2011) use an econometric approach to indirectly assess the additionality and co-benefits of sulfur dioxide emissions reductions in China under the Clean Development Mechanism, and do not find any significant impacts. Millard-Ball (2012) explores similar issues for sectoral emissions baselines. Adverse selection in voluntary environmental programs has also been demonstrated in energy conservation programs (Hartman, 1988).

⁴We first explored these issues in an earlier working paper version (Van Benthem and Kerr, 2010). The current paper is a substantially updated version.

cap-and-trade system.⁵ The third policy is to set stringent baselines.

Specifically, we assess performance using three inter-related criteria. *Efficiency* is determined by whether land goes to its optimal use – land that yields agricultural or timber returns exceeding the positive environmental externalities from the forest should be cleared; land with lower returns should not.⁶ The *average cost* per unit of emissions abatement through avoided deforestation is an indicator of the offset buyers’ value for money. *Quality* or *environmental integrity* of offsets is measured as the percentage of offsets that are not spurious. Spurious offsets lead to a global environmental loss if their presence is not factored in through a more ambitious cap.

We use a combination of analytical results and numerical simulations from a microeconomic model of land use. This model shows that asymmetric information about the forest baseline in voluntary programs leads to trade-offs between efficiency, average cost and offset quality. We then introduce a framework to make explicit the benefits and costs of transfers between industrialized and developing countries, and present the Pareto set of individually rational policies.

We draw the following conclusions. First, baseline scale increases improve efficiency and quality and reduce infra-marginal transfers from buyers to sellers, leading to lower average cost. Second, discounting and trading ratios are inefficient (since they make participation unattractive to certain sellers of non-spurious offsets), but also reduce transfers. In addition, trading ratios between offsets and allowances have ambiguous environmental effects if the cap is not properly adjusted. While the conventional wisdom – that discounting can be used to reduce the fraction of offsets that are spurious – is invalid⁷, discounting can make offsets more attractive to industrialized countries.⁸ Third, more stringent baselines also reduce efficiency and reduce average cost for industrialized countries but, in contrast to discounting, generally improve the quality of offsets.⁹ Finally, the Pareto set only contains policies that involve some combination of discounting and/or stringent baselines, to guarantee that it is individually rational for industrialized countries to participate. In an international context, offset programs are therefore never first-best, but can be efficiency improving, especially with increased baseline scale.

This research is important since a lot of funds are currently being invested in offset programs with baselines in developing countries. Examples include the Clean Development Mechanism (CDM), Verified Carbon Standards, and, more recently, an international program to reduce deforestation referred to as REDD (Reducing Emissions from Deforestation and Degradation), and at a national level, the payments for ecosystem services program in Costa Rica.¹⁰ Some are beginning

⁵In our model, lowering the price is equivalent to requiring a trading ratio for offsets 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.

⁶Environmental externalities include non-carbon forest benefits such as flood protection, water quality and biodiversity, but these will be ignored if the offsets are designed for carbon markets only.

⁷Chung (2007), Environmental Defense Fund (2007), Greenpeace and Papua New Guinea (2008) and Schneider (2009) all propose discounting as a way to reduce “non-additional” or spurious offsets.

⁸Another related policy, limiting the number of offsets that can be one-for-one exchanged with allowances in a cap-and-trade system, will lower the offset price and hence make offsets more attractive to buyers, but also reduce efficiency and quality. See Bento et al. (2011) for more details.

⁹With stringent baselines, projects can be awarded *fewer* offsets than their actual emissions reductions (Bento et al., 2012).

¹⁰Several studies provide evidence on adverse selection in the context of Costa Rican deforestation (Kerr et al.,

to be implemented on a wider scale and provide rewards based on a regionally or nationally set baseline – notably Norway’s innovative contracts with Guyana, Brazil and Indonesia.¹¹ Our findings suggest that such programs have the potential to be more efficient and attractive as an alternative mitigation option for industrialized countries, provided that developing country governments can effectively respond to the financial incentives and pass them on to local forest owners.

Our paper can be interpreted as an analysis of either adding avoided deforestation offsets 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 provide rewards relative to a forest baseline. In a cap-and-trade market these rewards would be offsets that can be converted to emissions allowances, whereas in the fund these rewards would be dollars. In both programs, buyers in industrialized countries pay for emissions reductions in developing countries. These two approaches are equivalent under the following assumptions. First, the rewards must be the same per unit of avoided deforestation. Second, the emissions cap and the money in the fund can be adjusted so that regardless of which approach is used, global emissions are identical.¹²

Our presentation focuses on deforestation but the results are equally applicable to a variety of voluntary offset programs, national and international, particularly those that involve international public goods and financial transfers between countries, and to climate agreements between industrialized and developing countries more generally.

The remainder of this paper is organized as follows. Section 2 presents a simple model of a voluntary avoided deforestation program that demonstrates the trade-off between efficiency losses from adverse selection and the level of transfers. Section 3 discusses how the offset program’s performance is affected by baseline scale increases, offset discounting and changing the generosity of the assigned baseline. Section 4 presents the set of individually rational, Pareto efficient policies across international borders. Section 5 concludes and summarizes the main policy implications.

2004; Sanchez-Azofeifa et al., 2007; Robalino et al., 2008). Strand (2011) shows that the CDM may serve as an obstacle to implementing more stringent environmental policies, in an effort to influence baseline assignment. A plethora of reports explore the issues associated with the design of REDD. See Angelsen (2008, 2010), Plantinga and Richards (2008) and Busch et al. (2009) for discussions of the challenges.

¹¹For example, see Government of Kingdom of Norway and Government of the Republic of Indonesia (2010).

¹²Suppose industrialized countries (ICs) have a joint emissions cap that requires them to undertake abatement A . The allowance price equals p^* . ICs could use a fund to achieve n further units of abatement (and pay for m infra-marginal, or “spurious”, units), at a price per unit p_c which may be lower than p^* . Global abatement would be $A + n$, where n is a function of p_c . Instead of a fund, ICs could purchase $n + m$ offsets from developing countries (DCs) at price p_c and trade them in their cap-and-trade system. This, however, would not be a fair comparison. Under the fund, global abatement equals $A + n$. Using offsets that can be traded one-for-one with allowances, global abatement will be $A - m$. The environmental outcome is worse than without offsets (and p^* would decrease). To correct this, ICs must increase their abatement target to $A + n + m$. This ensures that, after $n + m$ offsets are purchased from DCs, IC mitigation effort is back at A and the allowance price at p^* . Global abatement is now also $A + n$. If a trading ratio $t : 1$ is applied, global abatement with offsets 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 abatement target would be needed to make the fund and the offset program equivalent. As long as the introduction of offset trading does not cause DCs to include non-forest sectors in an international cap-and-trade program, there is no effect from strategic bargaining about baselines, in which environmentally less (more) concerned countries have an incentive to demand more (fewer) allowances than when they commit to emissions reductions without trading possibilities (Helm, 2003).

2 A Simple Model of Voluntary Opt-In

2.1 A trade-off between efficient subsidies and baselines with adverse selection

Consider a continuum of small plots of forested land, indexed by i . Landowners decide on a plot-by-plot basis whether to clear fully or keep the forest. For each plot, its owner will clear the 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 the forest. Landowner i knows r_i with certainty. The marginal environmental externality from deforestation is defined as δ .¹³ Returns r_i are distributed across i with density f_r . Note that in reality, landowners may own several plots, and do therefore not face an “all-or-nothing” deforestation decision for their total forest area. We will return to this issue in Section 3. The intuition for adverse selection can be illustrated at the single plot level, however, and the single plot model forms a useful building block for further analysis. Therefore, in the remainder this section, we assume that landowners own exactly one plot.

The simplest policy would be to offer a subsidy equal to $p_c = \delta$ per plot that remains forested. 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) deforest.¹⁴ The change in economic surplus ΔS_{eff} from this efficient policy relative to no policy equals:

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

This achieves efficient deforestation but requires a large transfer of resources:

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

The total amount of avoided deforestation is:

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

The average cost ($AC = TT/AD$) to industrialized countries of climate mitigation through avoided deforestation summarizes the program’s value for money for offset buyers. Under the

¹³We implicitly assume that the amount of carbon per hectare of forest is constant. This could be relaxed with little loss of generality. Converting from forest cover to carbon still mostly relies on carbon tables (derived from field work supplemented with LIDAR, an optical remote sensing technology) for different ecological conditions that can be identified with Geographic Information Systems data. This measurement method could induce a bias in the estimates of carbon saved. If the threat to forests is positively correlated with the unobservable errors in carbon measurements, incentives will be poorly targeted and forests with higher than average carbon stocks will be protected at sub-optimally low levels (Kerr et al., 2004). If carbon is not determined ex-ante using carbon tables, rewarding carbon ex-post using LIDAR or on-the-ground measurements will expose the seller to extra risk, in addition to other forest risks such as burning. Both risks can be addressed through insurance type mechanisms.

¹⁴The r_i may be interdependent. In general equilibrium, one landowner’s deforestation decision will alter returns for others. This could operate through leakage: e.g., a landowner who does not deforest reduces the supply of timber, which affects local timber prices. 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. f_r could be thought of as an ex-post distribution of returns when a new set of equilibrium land uses is reached.

subsidy, the average cost is high if many plots have negative returns and would not have been cleared even without the subsidy.

To avoid large transfers to forest-rich developing countries, a second policy option is a voluntary offset program that will pay participants an amount p_c for each hectare of forest exceeding an assigned baseline.¹⁵ Landowners know their true plot level forest baselines BL_i :

$$BL_i = \begin{cases} 1 & \text{if } r_i \leq 0 \\ 0 & \text{if } r_i > 0 \end{cases} \quad (4)$$

We model an “all-or-nothing” decision at the plot level, but note that if a landowner owns multiple plots, this policy assigns a separate baseline for each of his plots. This implies that deforestation does not have to be a binary decision at the landowner level.

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

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

In practice, however, the regulator cannot observe r_i , but instead observes $\hat{r}_i = r_i + \varepsilon_i$.¹⁶ The observation error ε_i has density $f_\varepsilon \sim (0, \sigma_\varepsilon)$ and is assumed to be symmetric around 0 and independent of f_r . The *predicted* baselines are:

$$\widehat{BL}_i = \begin{cases} 1 & \text{if } \hat{r}_i \leq 0 \\ 0 & \text{if } \hat{r}_i > 0 \end{cases} \quad (6)$$

What happens if the government assigns baseline $\widehat{BL}_i$? When $(r_i > 0, \hat{r}_i > 0)$ or $(r_i \leq 0, \hat{r}_i \leq 0)$, the assigned baseline coincides with the true baseline. The landowner will make the socially efficient decision. However, if $(r_i > 0, \hat{r}_i \leq 0)$, 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

¹⁵If it were practically feasible, a policy that sets $p_c = r_i$ would reduce transfers even further. Mason and Plantinga (2010) describe a model in which the regulator can 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” payment 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 funders in an environmental transfer program was developed in Kerr (1995). Our model does not consider this option, since such sophisticated contracts are not likely in most developing countries.

¹⁶This assumption is similar to Montero (2002)’s approach for modeling asymmetric information between a regulator and firms about *aggregate* abatement costs in a setting of incomplete enforcement of a regulation.

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:

$$\int_0^{p_c} (p_c - r) \left(\int_{-\infty}^{-r} f_\varepsilon(\varepsilon) d\varepsilon \right) f_r(r) dr \quad (7)$$

The amount of avoided deforestation falls by:

$$\int_0^{p_c} \left(\int_{-\infty}^{-r} f_\varepsilon(\varepsilon) d\varepsilon \right) f_r(r) dr \quad (8)$$

Finally, consider the case where $(r_i \leq 0, \hat{r}_i > 0)$. 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 cases described above are summarized in Figure 1.

Figure 1: Adverse selection causes an efficiency loss in the range $0 \leq r_i \leq p_c = \delta$ and increases average cost in the range $r_i \leq 0$ if landowners in these ranges get assigned an incorrect baseline.

The total transfer (TT) in an offset program is lower than under the pure subsidy (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.¹⁷

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

Deforestation is higher 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).¹⁸ Adverse selection increases the average cost relative to a full information voluntary program, but the comparison with average cost under the subsidy is theoretically ambiguous.

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

$$AC = p_c \left(1 + \frac{\int_{-\infty}^0 \left(\int_{-r}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon \right) f_r(r) dr}{\int_0^{p_c} \left(\int_{-r}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon \right) f_r(r) dr} \right) = p_c \left(1 + \frac{OS}{AD} \right) \tag{10}$$

where OS denotes the amount of infra-marginal forest credited, or spurious offsets. 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 ($FOS = \frac{OS}{OS+AD}$) as a measure of the offset quality.

This section has highlighted that adverse selection causes a trade-off between efficiency and average cost for offset buyers. Section 3 presents numerical simulations that indicate that adverse selection can have a dramatic negative impact on the performance of voluntary offset programs. We also show that increasing the baseline scale from the individual plot's forest to, say, a whole country's forested land reduces this trade-off significantly.

2.2 The impact of the observation error distribution on performance measures

The trade-off between efficiency and average cost depends on the distribution of observation errors. In this section, we analyze the impact of the observation error variance on our three performance measures: economic efficiency, average cost (AC) and offset quality (FOS). Appendix A discusses the impact of the marginal costs of avoiding deforestation on these criteria.

Equation (7) shows that any change in $f_{\varepsilon}(\varepsilon)$ that increases the probability mass in the range $[-\infty, -r_i]$, where $0 \leq r_i \leq p_c$, increases the efficiency loss from adverse selection (assuming $p_c = \delta$) and decreases avoided deforestation (AD). A mean preserving spread such that $F'_{\varepsilon}(x) \geq F_{\varepsilon}(x)$

¹⁷In a cap-and-trade system, infra-marginal transfers would be analogous to spurious or non-additional offsets.

¹⁸From (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, $TT(baseline) < TT(full\ information)$.

$\forall x < 0$ is sufficient. If the distribution of errors is normal, an increase in variance will generate such a mean preserving spread.

Under the same assumptions and for $f_\varepsilon(\varepsilon)$ symmetric around 0, AC will increase.¹⁹ 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 observation error $f_\varepsilon(\varepsilon)$. In the remainder of this paper, we will focus mostly on return distributions $f_r(r)$ for which (1) $F_r(0) > 0.5$ and (2) 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, at least in the short term. Landowners have previously chosen not to clear the remaining forest so only land on which relative returns have recently risen will be at risk of clearing. The second assumption implies that there is a higher probability mass for returns just below zero than for returns just above zero, which intensifies the trade-off between efficiency and infra-marginal transfers. These assumptions are only meant to guide intuition but are not necessary for our main results to hold, and we will show robustness to a range of distributional assumptions in Appendix B (normal, exponential and uniform $f_r(r)$).

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 assumptions trivially hold. Additional deforestation occurs because the returns distribution shifts over time. If this shift has both a common (e.g., local technology and infrastructure changes) and an idiosyncratic (e.g., a random plot-specific return shock) element, we would still expect the assumptions to hold.²⁰

We consider $f_r(r) \sim N(-1, 1)$, $f_\varepsilon(\varepsilon) \sim N(0, \sigma_\varepsilon^2)$ and $p_c = \delta = 0.5$ as our central case. Figure 2 plots the various performance criteria 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

¹⁹ OS , defined in (10), will increase because $OS(f'_\varepsilon) = \int_{-\infty}^0 \left(\int_{-r}^{\infty} f'_\varepsilon(\varepsilon) d\varepsilon \right) f_r(r) dr = \int_{-\infty}^0 (1 - F'_\varepsilon(-r)) f_r(r) dr =$ (by symmetry of f_ε) $\int_{-\infty}^0 F'_\varepsilon(r) f_r(r) dr \geq$ (by the mean preserving spread) $\int_{-\infty}^0 F_\varepsilon(r) f_r(r) dr = OS(f_\varepsilon)$. Since AD decreases and OS increases, AC increases.

²⁰The common shock will generate a probability mass of forested land above zero return up to the size of the shock. The idiosyncratic shock will move some land to higher returns (and some to lower), leading to a lower probability mass in the upper tail of the returns distribution. Positive shocks, such as discovery of oil under remote non-productive land, encourage deforestation; negative shocks, such as the creation of a national park, discourage forest clearing.

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

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 spurious offsets are given out only to those with $r_i < 0$. As σ_ε increases, 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.

This section has shown that a mean preserving spread that widens the tails of the observation error distribution unambiguously has negative effects on all three performance measures. Any improvement in our ability to observe returns, or equivalently predict deforestation, would reduce the trade-off between efficiency and transfers.

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 baseline scale, offset price discounting and changing the generosity of the assigned forest baseline.

3.1 Policy 1: increasing the baseline scale

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 many real-world forest carbon programs are indeed small scale (though not literally at the single plot level), other proposals feature baselines for large areas (e.g., a region or a country - if the funder is international). The Costa Rican payments for ecosystems services program is an example of a relatively small scale system where individual landowners choose whether to opt in. Brazil's Amazon Fund, to which Norway has agreed to contribute up to US\$1 billion, sets up a large scale system. A baseline is assigned based on a 10-year average deforestation rate for the entire Amazon, and donor payments are based on the difference between the previous year's deforestation and the baseline.²¹ Empirical evidence suggests that before 2000 the Costa Rican landowner-scale program was ineffective (Sanchez-Azofeifa et al., 2007): it was costly but had no significant impact on deforestation. In contrast, efforts associated with the Amazon Fund have appeared to be effective at reducing deforestation so far (Nepstad et al., 2009 and Soares-Filho et al., 2010).

Larger programs devolve responsibility to large local entities for passing on incentives to individual landowners. Such entities may have more authority and better information to influence local deforestation than an international regulator. We now show that setting baselines for larger areas (consisting of many plots) instead of baselines at the level of an individual landowner's forest (consisting of just a few plots) also mitigates the severity of the trade-off between efficiency, average cost and offset quality.

3.1.1 A multiple plot model

We now consider a single entity (a regional or national jurisdiction or a large indigenous entity) 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:

$$BL_N = \sum_{j=1}^N BL_j \text{ where } BL_j = \begin{cases} 1 & \text{if } r_j \leq 0 \\ 0 & \text{if } r_j > 0 \end{cases} \quad (11)$$

The government observes each r_j with error ε_j : $\hat{r}_j = r_j + \varepsilon_j$. Assume that ε_j is i.i.d. across j . This means that $\hat{r}_j$ has a distribution with mean μ_r and variance $\sigma_r^2 + \sigma_\varepsilon^2$. The distribution of $\hat{r}_j$ is more dispersed than $f_r(r)$. The government could compute an unbiased prediction of the baseline $\widehat{BL}_{N,unbiased}$ as the sum of the expectations of the random variables for the plot-specific baselines. From its point of view, the true baseline for a specific plot BL_j is a Bernoulli random variable with mean p_j and variance $p_j(1 - p_j)$, where $p_j = Pr(r_j < 0|\hat{r}_j) = Pr(BL_j = 1|\hat{r}_j)$.²² Since these

²¹<http://www.regjeringen.no/en/dep/md/Selected-topics/climate/the-government-of-norways-international-norway-amazon-fund.html>

²²Note that $p_j \neq Pr(\hat{r}_j < 0)$, except if $f_r(r)$ is symmetric around zero. If the government naively assumed that r and $\hat{r}$ have the same distribution, it would calculate $p_j = F_\varepsilon(-\hat{r}_j) =$ (if f_ε is symmetric) $1 - F_\varepsilon(\hat{r}_j)$. This would lead

are non-identically but independently distributed across j , the central limit theorem yields that for $N \rightarrow \infty$:

$$BL_N = \sum_{j=1}^N BL_j \xrightarrow{d} N \left(\sum_{j=1}^N p_j, \sum_{j=1}^N p_j (1 - p_j) \right) \quad (12)$$

The unbiased assigned baseline $\widehat{BL}_{N,unbiased} = \sum_{j=1}^N p_j$, is a single, cumulative baseline for all N plots (as opposed to N plot-specific baselines with N separate opt-in decisions).

3.1.2 The effect of increased baseline scale on efficiency

With the N -plot aggregate 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 at all. This does not mean that the entity must protect the entire forest area to get any rewards, but rather that they accept a baseline (and agree to being monitored) for the entire area and are rewarded based on the difference between baseline and monitored forest levels. The difference with the single plot model is illustrated by Figure 3.

Figure 3: Single versus multiple plot policy.

Figure 3 contrasts the single plot with the multiple plot case. In the single plot case, an inefficiency occurs when the true baseline is clearing the forest (0), but the government assigns a no-deforestation baseline (1). In the multiple plot case, assigning a more favorable baseline ($\widehat{BL}_N <$ 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, $f_{\hat{r}}(\hat{r}) \sim N(-1, 2)$. The probability that $\hat{r} > 0$ exceeds the probability that the true return $r > 0$. Therefore, if the government used a bottom-up plot-level approach to estimate $\hat{r}$ and assigned a zero baseline for all plots with positive $\hat{r}$, the baseline would be biased downwards.

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:

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

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 no-deforestation baseline, and opt out. However, in some cases the efficiency of the new system will be lower than with plot-specific baselines. This happens when the baseline is so unfavorable that the entity opts out *with all N plots*. In contrast, in the single plot program, some critical plots with returns between 0 and p_c will get assigned a correct baseline and some deforestation will be efficiently avoided.

By the law of large numbers, as $N \rightarrow \infty$, $\frac{BL_N}{N} \rightarrow \frac{\widehat{BL}_{N,unbiased}}{N}$: the standard error of the average baseline *per plot* goes to zero. However, the standard error of BL_N does not converge to zero. Therefore, it is possible that the entity gets assigned a baseline that is 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 $N \rightarrow \infty$ and the efficient solution will be obtained.

In the limit, a larger baseline scale will lead to the same efficient outcome as under the full information voluntary program.²³ However, real-world programs can only be scaled up to a finite number of plots.²⁴ We therefore turn to numerically exploring the effects of moderate increases in baseline scale.

3.1.3 Numerical simulations of increased baseline scale

We numerically simulate the differences between single plot versus multiple plot baselines for a “country” with a total forest area of N_{total} plots. We compare offset programs that assign N_{total}/N separate baselines covering N plots each for different values of $N \in \{1, 2, \dots\}$. We show results for certain example error and returns distributions, but the results hold quite generally for many distribution (we discuss an exception below). Throughout this section, we assume $f_\varepsilon(\varepsilon) \sim N(0, 0.5^2)$, unbiased assigned baselines $\widehat{BL}_{N,unbiased}$, and $p_c = \delta = 0.5$. The central case returns distribution

²³Another benefit of scaling up is that it reduces leakage: reduced deforestation within the project increases pressure to clear in neighboring regions and gains are offset by losses elsewhere. As project scale rises, these leakage effects diminish.

²⁴The next section illustrates that, for *finite* N and certain unusual land return distributions, it is possible that efficiency initially *decreases* with N .

is $f_r(r) \sim N(-1, 1)$, but we also consider alternative distributions below and, in more detail, in Appendix B. Table 1 demonstrates what happens to the performance criteria as N increases.

Table 1: The impact of increasing the baseline 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	9.76%	18.30%	58.42%	97.47%	100%

Notes: Each simulation contains $N_{total} = 10,000$ plots, with baselines assigned for sets of $N \in \{1, 2, 10, 100\}$ plots. Simulations are repeated 100 times. $\delta = p_c = 0.5$. $f_\varepsilon(\varepsilon) \sim N(0, 0.5^2)$. Efficiency gains, AD and TT are all normalized per 100 plots. Assigned baselines are unbiased.

Figure 4: A graphical representation of the impact of baseline scale on the performance criteria.

Table 1 shows that increasing the baseline scale has dramatic consequences for its performance. For the central case, N -plot aggregate baselines increase efficiency and AD and reduce AC as N increases. In this example, $N = 100$ is enough to approach the efficient solution. Hence, scale mitigates adverse selection. As explained above, the observation error *normalized per plot* decreases as N grows, and the probability of opt-in becomes very high (97.47%), although the program will only induce a change in deforestation decisions for a limited number of the N plots (those for which $0 \leq r \leq p_c$). Figure 4 represents the same results graphically. The interval between the error bars contains 95% of all realized values.

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 results by analyzing the effects of baseline scale increases 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 more exotic distribution: a symmetric bimodal normal $BMN(0.5, 0.1^2)$ distribution with modes at -0.5 and 0.5 and standard deviation $\sigma_r = 0.1$ around each mode. The latter distribution is unlikely to represent reality, but illustrates that – for finite N – efficiency does not necessarily monotonically increase in N .

Figure 5 summarizes the effects on efficiency and FOS of increasing baseline scale and compares it to the efficient solution. For the $N(0, 1)$ distribution, the effects are similar to the central case distribution. Both efficiency and FOS improve with scale. The $BMN(0.5, 0.1^2)$ 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 $BMN(0.5, 0.1^2)$ 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* in N between $N = 1$ and 10. However, such distributions are highly stylized and unlikely to represent true returns distributions.²⁵ AC and FOS decrease with scale for the three example distributions, although the effect is also theoretically ambiguous for finite N .²⁶ The simulation results are consistent with the empirical evidence discussed above: the small-scale baselines in Costa Rica have performed poorly (Sanchez-Azofeifa et al., 2007) relative to nationwide baselines for Brazil (Nepstad et al., 2009 and Soares-Filho et al., 2010).

Figure 5: The impact on efficiency (left panel) and FOS (right panel) of increasing the baseline scale for alternative return distributions, for $N = 1, 2, 10$ and 100.

We now test the robustness of our conclusions in another way. An important assumption has

²⁵We generated results with many alternative probability distributions. The only counterexamples we found for the “scale improves efficiency” result are distributions with a probability mass that is heavily concentrated around δ .

²⁶Consider 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: baseline scale adversely impacts both FOS and AC .

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_{ur}^2 and $\sigma_{u\varepsilon}^2$, such that $\sigma_r^2 = \sigma_{ur}^2/(1 - \rho_r^2)$ and $\sigma_\varepsilon^2 = \sigma_{u\varepsilon}^2/(1 - \rho_\varepsilon^2)$. Table 2 summarizes the main findings for the central case.²⁷

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	100-plot	100-plot	Maximum efficiency:
		$\rho_r = \rho_\varepsilon = 0$	$\rho_r = \rho_\varepsilon = 0.5$	$\rho_r = \rho_\varepsilon = 0.9$	$\rho_r = \rho_\varepsilon = 0.99$	$\sigma_\varepsilon = 0$
	(1)	(2)	(3)	(4)	(5)	(6)
Efficiency gain	1.33	2.47	2.41	2.02	1.59	2.53
<i>AC</i>	0.91	0.51	0.52	0.61	0.80	0.50
<i>FOS</i>	45.29%	1.86%	3.54%	17.76%	37.69%	0%
<i>AD</i>	5.34	8.95	8.72	7.27	5.72	9.18
<i>TT</i>	4.88	4.56	4.52	4.42	4.59	4.59
Opt-in	9.76%	97.47%	94.96%	78.40%	37.28%	100%

Notes: Each simulation contains $N_{total} = 10,000$ plots, with baselines assigned for sets of $N = 1$ or $N = 100$ plots and various degrees of spatial correlation. Simulations are repeated 100 times. $\delta = p_c = 0.5$. $f_\varepsilon(\varepsilon) \sim N(0, 0.5^2)$. Efficiency gains, *AD* and *TT* are all normalized per 100 plots. Assigned baselines are unbiased.

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 quickly cancel out across plots, but are persistent. High spatial correlation reduces the probability of participation for a given N and therefore adversely impacts the performance criteria. A larger baseline scale will counteract the effects of spatial correlation: setting baselines at the landowner level (one or several plots) will be much less efficient than setting baselines for large regions ($N \gg 100$).²⁸

3.2 Policy 2: discounting the payment per hectare

In the analysis above, the payment per hectare p_c equals the marginal externality from deforestation δ . A second potential policy is to vary p_c . This section analyzes what happens to the three performance measures. Reducing p_c is equivalent to the practice of “offset discounting” sometimes observed in real-world offset programs (Kollmuss and Lazarus, 2011). 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, or a trading ratio is applied between offsets and allowances

²⁷Results for the $N(0, 1)$ distribution, and the various other distributions that we simulated, are qualitatively similar, but the relative efficiency loss from high spatial correlation varies across distributions.

²⁸If u_ε is i.i.d., a large enough N will eliminate the effect of spatial correlation. However, if there are unobservable factors that affect large numbers of plots similarly, adverse selection can remain even for large-scale baselines.

in a cap-and-trade system. As discussed in Section 1, this is often promoted as a way to correct for spurious offsets. 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 ($(\bar{r}|0 \leq r \leq \delta) > p_c$) would opt out. Paying more than δ (a “premium” rather than a “discount”) reduces efficiency because some entities will opt in even though their private gains from deforestation exceed the environmental value of the forest. In the single plot model with symmetric information, the change in efficiency relative to a no policy case was given by (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:

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

Proposition 1. 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 performance criteria as the payment p_c varies.

Proposition 2. 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 2 demonstrates that, contrary to the intended effect, the fraction of offsets that is spurious, FOS , *increases* when the offset price is discounted: as p_c falls the share of offsets that is spurious rises toward 1. The intuition is that discounting never changes the decision for spurious sellers, but sellers of “real” offsets will opt out if their return is between p_c and δ . Discounting does reduce transfers, which can make the offset program more attractive to industrialized countries. The effect on AC is theoretically ambiguous for low values of p_c , but eventually AC must increase.²⁹

We conclude that, in the single plot model, efficiency is maximized by paying $p_c = \delta$. Discounting increases FOS , leads to less AD , but requires fewer transfers.

3.2.2 Discounting in the multiple plot model

In the multiple 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 positive 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 is 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 6 illustrates that raising p_c above δ can increase efficiency ($\delta = 0.5$ and $N = 10$ or 100), and also shows the impact of discounting on other criteria. For $N = 10$, raising p_c above δ (to $p_c = 0.6$) slightly increases efficiency. This increased efficiency coincides with higher AC , 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%. The efficient p_c is close to 0.5.

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 $\delta < p_c < p'$, the trade-off between efficiency and average cost remains.

Analogous to Proposition 2 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

²⁹Since 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 .

Figure 6: The impact of changing p_c and baseline 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 .

opt in. At a higher price $p_c + \varepsilon$, 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 improving the quality of offsets is *not* a valid reason to advocate discounting: offset quality always falls if p_c decreases. 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 global mitigation will fall. If, however, discounting is achieved by trading ratios (offsets and allowances trade in $(t : 1, t > 1)$ ratios) this could increase global mitigation for high enough t , even if the cap is not adjusted.³⁰ Discounting also, almost always, reduces the average cost to industrialized countries of real mitigation through offsets. Thus discounting, while inefficient, makes industrialized countries more willing to fund mitigation in developing countries. Section 4 explains why discounting may be *necessary* for a self-interested, rational international buyer to participate in an offset program.

3.3 Policy 3: changing the generosity of the assigned baseline

A third 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 baseline for plot i :

$$\widehat{BL}_i = \begin{cases} 1 & \text{if } \hat{r}_i \leq r^* \\ 0 & \text{if } \hat{r}_i > r^* \end{cases} \quad (16)$$

where r^* is a specified return set by the government. The government, aware of adverse selection,

³⁰Note 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 .

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 performance 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 3. More generous baselines (weakly) increase efficiency and AD , but require a (weakly) higher TT .

Proof. The efficiency change relative to no policy equals $\Delta S(r^*) = \int_0^{p_c} (p_c - r) \left(\int_{r^*-r}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon \right) f_r(r) dr$. By Leibniz' Rule, this expression is globally weakly decreasing in r^* . Hence, efficiency is maximized if $r^* \rightarrow -\infty$. AD is given by $AD(r^*) = Pr(0 \leq r \leq p_c, \hat{r} > r^*) = \int_0^{p_c} \left(\int_{r^*-r}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon \right) f_r(r) dr$, which is globally weakly decreasing in r^* . TT is given by $TT(r^*) = p_c \int_{-\infty}^{p_c} \left(\int_{r^*-r}^{\infty} f_{\varepsilon}(\varepsilon) d\varepsilon \right) f_r(r) dr$, which is also globally weakly decreasing in r^* . ■

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 an unconditional 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. The shape of AC and FOS is dependent on the return distribution $f_r(r)$, but for reasonable returns distributions like our central case, AC and FOS are decreasing in r^* (as baselines become more stringent).

The ambiguity of the effect on FOS can be demonstrated as follows. Using (9) and (10), $OS/AD = IT/TT$ 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^* , OS/AD and hence FOS is decreasing in baseline stringency. This condition will certainly hold if the baseline error is bounded from above. 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 .

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 7 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) reduces

efficiency and AD , but also reduces AC and FOS . For very stringent baselines FOS becomes negative: the environmental gains are greater than the number of traded offsets. Again, efficiency, average cost and offset quality are conflicting policy aims: efficiency requires setting generous baselines, while minimizing average cost and maximizing offset quality requires setting stringent baselines.

Figure 7: The impact of changing baseline generosity on the performance criteria, for $f_r(r) \sim N(-1, 1)$ and $N = 100$. 84 is an unbiased baseline.

Section 3 has shown that only increasing baseline scale improves all performance measures simultaneously (for “typical” returns distributions). Offset discounting and changing baseline generosity lead to a trade off between performance criteria. Discounting reduces efficiency, AD and offset quality, but reduces the average cost for funders. Making assigned baselines more stringent reduces efficiency and AD , but improves offset quality and reduces average cost. This 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 performance criteria discussed in this paper.

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

	Policy option		Performance criteria		
	Required scale of baseline (1)	Maximize efficiency (2)	Maximize offset quality (minimize <i>FOS</i>) (3)	Maximize value for money (minimize <i>AC</i>) (4)	Maximize avoided deforestation (5)
Increase scale	Finite N $N \rightarrow \infty$	(+) +	(+) +	(+) +	(+) +
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 Constrained Pareto Efficient Policies

The previous sections have illustrated that avoided deforestation policies involve trade-offs between environmental outcomes, financial transfers and efficiency. This section shows how the efficient policy differs from the Pareto set of individually rational policies.

4.1 Globally efficient vs. constrained Pareto efficient policies

A global social planner, who can force participation and faces no costs of transfers, could maximize efficiency and then redistribute the surplus to meet distributional (and political) objectives. This is a useful benchmark, which we refer to as the "globally efficient policy". A subsidy per hectare with ex-post international redistribution of funds would achieve the efficient allocation.³¹ At the national level, it is reasonable to assume that the government can transfer funds across industries, potentially to compensate losers and/or to make all parties better off (Bovenberg et al., 2008). However, in an international context, unconstrained money transfers across international borders are not realistic. Instead, buyers (industrialized countries, or ICs) and sellers (developing countries, or DCs) can only bargain over prices and baselines, but not commit to ex-post financial transfers.

³¹Various global climate change models (e.g., DICE/RICE) calculate "cooperative efficiency gains" in this way or assume that countries set binding national targets while ignoring beneficial international environmental spillovers (Nordhaus, 1996).

We therefore focus on policies in which both parties will participate voluntarily ex-ante. Our basic policy design ensures that the DC participation constraint is always met, since they can opt out. We now define the extremes of the Pareto set of policies that make both parties at least as well off compared to a no policy case:

$$\max_{\widehat{BL}, p_c} S_{IC} \text{ s.t. } \Delta S_{DC} \geq 0 \quad \& \quad \max_{\widehat{BL}, p_c} S_{DC} \text{ s.t. } \Delta S_{IC} \geq 0 \quad (17)$$

where S denotes surplus and ΔS surplus change. Stakeholders should in theory be able to negotiate a policy in the Pareto set. However, we will show below that a globally efficient policy can no longer be negotiated.

To define surplus, consider a fund that ICs use to pay DC governments that opt in to the program.³² DC forest owners receive p_c per hectare, which leads to AD hectares of avoided deforestation and OS units of infra-marginal forest receiving payments (offsets that are spurious). They value each dollar received at face value. DCs forego returns r_j on each of the AD hectares of avoided deforestation. In this simple framework, we make the strong assumption that the global environmental gain δ is fully valued by ICs, and is not valued at all by DCs.³³ ICs transfer resources to the fund at face value plus an additional cost or benefit γ per dollar. If $\gamma = 0$, a \$1 dollar transfer costs the ICs exactly \$1. Alternatively, $\gamma > 0$ if distortionary taxes need to raise money for the fund. $\gamma < 0$ if rich ICs derive satisfaction from donating funds to poor DCs.

Under these assumptions, the costs and benefits to the different countries are:

$$\Delta S_{DC} = AD \left(\frac{\int_0^{p_c} (p_c - r) f_r(r) dr}{\int_0^{p_c} f_r(r) dr} \right) + OS p_c = AD (p_c - E[r|0 \leq r \leq p_c]) + OS p_c \quad (18)$$

$$\Delta S_{IC} = \delta AD - p_c(AD + OS)(1 + \gamma) \quad (19)$$

$$\Delta S_{global} = AD (\delta - E[r|0 \leq r \leq p_c]) - \gamma p_c(AD + OS) \quad (20)$$

where ΔS_{global} assumes equal marginal utility of money for ICs and DCs. This specification of surplus combines the performance criteria discussed in previous sections: efficiency (captured by the first term of (20)), avoided deforestation, total transfers (included in the second term in (19)), and average cost (defined by p_c , AD and OS). Hence, this represents a framework to trade off these performance criteria subject to participation constraints.

We can now compute the Pareto set with boundaries defined by (17) above. Our model does not predict which point in the Pareto set will be chosen, but rules out feasible but Pareto-dominated

³²As discussed in Section 1, under assumptions we maintain in this paper, including avoided deforestation in a cap-and-trade system and rewarding it through an international fund are equivalent.

³³In reality, some of the environmental benefits accrue to DCs. If DC governments value them, they should intervene even without an IC-induced policy. This shifts the baseline upwards, but would not affect our qualitative results.

policies. The set of *feasible* policies neither contains the globally efficient policy nor any policy in which $p_c = \delta$ if there is remaining adverse selection and $\gamma \geq 0$: (19) shows that if $OS > 0, \gamma \geq 0$ and $p_c = \delta$, ICs will not participate.

As suggested in Section 3, two options could induce IC participation. First, offsets could be discounted ($p_c < \delta$). Second, the baseline could be set stringent enough to achieve $OS < 0$. Both will reduce efficiency.

4.2 Numerical simulations of optimal policy choices

We now illustrate the Pareto set with numerical simulations. Initially we assume $\gamma = 0$. The IC participation constraint becomes $\delta AD - TT \geq 0$. We simulate a multiple plot program ($N = 100$), which mitigates but does not eliminate adverse selection. Using the same parameters as in Section 3, the true baseline is 84 (out of 100 hectares) and the environmental gain is $\delta = 0.5$.

Figure 8 shows contour plots of the IC, DC and global surplus for policies that involve a combination of discounting (p_c varies between 0 and 0.65) and changing the generosity of the assigned baseline ($\widehat{BL}$ varies between 82 and 92).

Figure 8 indicates that DC surplus is decreasing in $\widehat{BL}$ and increasing in p_c (upper right panel): more generous policies are preferred. However, the surplus maximizing choice for ICs is an interior solution: $(p_c, \widehat{BL}) = (0.35, 86)$ (upper left panel). This involves a combination of discounting and a more stringent than unbiased baseline. The intuition is that a small change from $p_c = \delta$ to $p_c = \delta - \varepsilon$ increases IC surplus by $AD\varepsilon$, while AD decreases only slightly as landowners with returns between p_c and δ opt out. Similarly, the gains from reduced transfers from a slight increase in the stringency of the assigned baseline exceed the efficiency losses. Global surplus is decreasing in $\widehat{BL}$ and increasing as $p_c \rightarrow \delta$ (lower left panel).

The bottom right panel shows the Pareto set. The gray dotted lines indicate the environmental externality ($\delta = 0.5$) and the unbiased baseline ($BL = 84$). Point A maximizes the surplus that ICs can derive from an offset program; point B makes ICs indifferent about participating in the policy. Because of adverse selection, the Pareto set does (just) not include the unbiased baseline and efficient price policy $(p_{c,eff}, \widehat{BL}_{unbiased}) = (0.50, 84)$, illustrating an economic rationale for the simultaneous use of discounting and stringent baselines. Which policy in the Pareto set gets implemented ultimately depends on the relative bargaining positions of ICs and DCs. The global social planner's constrained optimal choice in the Pareto set would be point B, which assigns as much surplus as possible to DCs.

These results assume that $\gamma = 0$: the net effect of deadweight losses from raising revenues cancels out against any utility that ICs may receive from transferring funds to poor countries. Changing γ will affect the location of Pareto set. If $\gamma > 0$, transfers are more costly for ICs and the Pareto set moves further away from the globally efficient solution.³⁴ For large values of γ , the

³⁴For example, when $\gamma = 0.25$, the ICs' optimal policy choice would be $(p_c, \widehat{BL}) = (0.25, 85)$, while, within the Pareto set, DCs are best off with $(p_c, \widehat{BL}) = (0.35, 84)$. These policies result in a smaller global efficiency gain compared to when $\gamma = 0$.

Figure 8: Contour plots of IC, DC and global surplus for $\gamma = 0$, $f_r(r) \sim N(-1, 1)$, $N = 100$, and various combinations of $(p_c, \widehat{BL})$.

Pareto set covers a region with very low p_c and very high $\widehat{BL}$. In that case, no globally meaningful avoided deforestation policy is possible. In contrast, if $\gamma < 0$, the Pareto set will move southeast and possibly include $(p_{c,eff}, \widehat{BL}_{unbiased}) = (0.50, 84)$.³⁵ This is not an unrealistic assumption for certain ICs such as Norway, which has shown a willingness to finance avoided deforestation programs that are hard to justify if no donation motive were present.

We have shown that the location of the Pareto set depends on the ICs' cost (or benefit) of transferring funds. Another determinant of the location is the degree of adverse selection. The results in Figure 8 assume $N = 100$, representing a large forest holding entity such as a country. However, a country may not simply be large collection of individual plots with asymmetric information on the plot level. In addition, there may be asymmetric information at the government

³⁵For example, when $\gamma = -0.25$, the ICs' optimal policy choice would be $(p_c, \widehat{BL}) = (0.50, 87)$, while DCs are best off with $(p_c, \widehat{BL}) = (0.60, 83)$. The Pareto set now includes policies that are more efficient, such as $(p_{c,eff}, \widehat{BL}_{unbiased}) = (0.50, 84)$.

level. For instance, a DC government may know more about its own intentions to protect forests (e.g., through establishing national parks, or better enforcement) than a potential IC offsets buyer. If so, adverse selection persists even when negotiating baselines and prices with large DCs.

Qualitatively, one way to model the effect of country level adverse selection on the location of the Pareto set is by varying N . Figure 9 demonstrates the effect of reducing N from 100 to 10.

Figure 9: The Pareto set of a low degree ($N = 100$) versus a high degree ($N = 10$) of country level adverse selection, for $f_r(r) \sim N(-1, 1)$ and $\gamma = 0$.

Figure 9 shows that more adverse selection at the country level (lower N) moves the Pareto set further away from the globally efficient solution. Contracts now require more discounting and more stringent baselines. This reduces efficiency, total surplus and the amount of avoided deforestation.

We conclude that under reasonable assumptions ($\gamma \geq 0$), ICs should find it in their best interest to propose an avoided deforestation policy which features a combination of discounting and/or stringent baselines. Offset policies such as REDD, while never first-best, can be efficiency improving, especially with increased baseline scale. Total surplus gains are larger with increased IC generosity, but smaller with significant remaining adverse selection at the country level.

5 Conclusion and Policy Implications

This paper analyzes the performance of voluntary avoided deforestation programs in an international context. Using a microeconomic model of land use, we demonstrate a trade-off between three key performance criteria when there is asymmetric information: efficiency, minimizing average cost to offset buyers (typically industrialized countries), and quality of offsets. We then analyze how the performance of offset programs is affected by three policy choices: increasing the baseline scale

of projects, offset discounting (reducing the payment per hectare to below the value of the environmental externality) and changing the generosity of the baseline. Finally, we present a framework to make explicit the role of benefits and costs of international transfers, and demonstrate how the Pareto set of individually rational policies differs from the first-best outcome.

We have three main findings. First, under almost all circumstances, offset programs perform better when they are scaled up so that entire political jurisdictions such as regions or nations get assigned a single, aggregate baseline and must choose whether to opt in as one entity. This forces offset sellers to decide on participating with all of their forested land, rather than being offered the flexibility of plot-specific baselines under which they can opt in with several plots and opt out with others. This makes it less easy for the seller to exploit his information advantage, and leads to an efficiency gain. Second, offset discounting and trading ratios between offsets and allowances in a cap-and-trade system reduce efficiency but also reduce transfers to offset-selling developing countries. Contrary to conventional wisdom, offset discounting increases the percentage of offsets that are spurious, and even high trading ratios have ambiguous effects on global carbon emissions. Third, we show that setting more stringent baselines reduces efficiency and transfers to developing countries but, in contrast to discounting, generally improves the quality of offsets. The Pareto sets presented in Section 4 highlight the main rationale of using discounting and/or stringent baselines: they may be necessary to convince self-interested offset buyers to participate in the program. We then show that feasible offset programs are never first-best, but can be designed such that they yield an efficiency gain relative to no policy.

We offer two key messages for policy makers. First, scale up programs and baselines as much as possible. Programs with region-wide baselines will be more efficient than the currently common payments for ecosystem service programs, which deal with baselines for individual landowners. These programs will also offer better value for money to funders. An example of a large scale program is the Norwegian agreement with Brazil, where an Amazon-wide forest baseline is set. In contrast, Costa Rica offered the flexibility of separate baselines for each individual landowner. This flexibility maximized the potential for landowners to exploit information advantages over the regulator, and resulted in a very inefficient outcome. Other advantages of scaling up are that baseline deforestation rates might be somewhat easier to predict at the national or regional scale than at the individual plot or landowner level, and that there is less scope for leakage.

Second, recognize that the primary purpose of offset discounting and setting stringent baselines is to reduce the cost to industrialized countries to convince them to participate in avoided deforestation offset programs, but not too “correct” the problem of spurious offsets. While it is possible to design a program that is individually rational for both industrialized and developing countries and efficiency-improving, voluntary offset programs should be used with caution and – everything else equal – replaced with mandatory emissions reduction regulations whenever politically feasible.

That being said, we conclude that well-designed, large-scale offset programs that provide financial incentives to developing country governments are more likely to lead to significant efficiency gains than existing small-scale programs. This does require that such governments have the ability

to effectively respond to such incentives (e.g., by enforcing illegal logging laws, national park creation, stimulating alternatives for fuel wood, or promoting better agricultural practices) and pass them on to local forest owners. While not first-best policies, large-scale offset programs could meet the expectations of those who promote avoided deforestation as a key climate mitigation option in the short run.

Acknowledgments

The authors would like to thank Antonio Bento and Larry Goulder for valuable comments and suggestions. We also thank participants in seminars at Stanford University, Resources for the Future, the World Bank, NZARE, and the 2010 AIMES open science conference in Edinburgh for comments and suggestions. We also thank Junyang Qian for excellent research assistance. Suzi Kerr would like to thank the Stanford Program on Energy and Sustainable Development for support.

References

- [1] Andam, K., Ferraro, P., Pfaff, A., Robalino, J., Sanchez-Azofeifa, A., 2008. "Measuring the Effectiveness of Protected-Area Networks in Reducing Deforestation." *Proceedings of the National Academy of Sciences* 105(42): 16089-16094.
- [2] Angelsen, A. (ed.), 2008. *Moving Ahead with REDD: Issues, Options and Implications*. CIFOR, Bogor, Indonesia.
- [3] 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*.
- [4] Arguedas, C., Van Soest, D.A., 2009. "On reducing the windfall profits in environmental subsidy programs." *Journal of Environmental Economics and Management* 58(2): 192-205.
- [5] van Benthem, A.A., Kerr, S., 2010. "Optimizing Voluntary Deforestation Policy in the Face of Adverse Selection and Costly Transfers." Motu Working Paper 10-04. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1655568
- [6] Bento, A., Kanbur, R., Leard, B., 2011. "Designing Efficient Markets for Carbon Offsets with Distributional Constraints." Working paper, Cornell University.
- [7] Bento, A., Kanbur, R., Leard, B., 2012. "Super-Additionality: A Neglected Force in Markets for Carbon Offsets." Working paper 2012-09, Charles H. Dyson School of Applied Economics and Management, Cornell University.
- [8] Bovenberg, A.L., Goulder, L.H., Jacobsen, M.R., 2008. "Costs of alternative environmental policy instruments in the presence of industry compensation requirements." *Journal of Public Economics* 92(5-6): 1236-1253.
- [9] Busch J., Strassburg, B., Cattaneo, A., Lubowski, R., Bruner, A., Rice, R., Creed, A., Ashton, R., Boltz, F., 2009. "Comparing climate and cost impacts of reference levels for reducing emissions from deforestation." *Environmental Research Letters* 4(4).

- [10] Bushnell, J.B., 2011. "The Economics of Carbon Offsets." NBER Chapters, in: *The Design and Implementation of US Climate Policy*, National Bureau of Economic Research: 197-209.
- [11] Chung, R.K., 2007. "A CER discounting scheme could save climate change regime after 2012." *Climate Policy* 7(2): 171-176.
- [12] Environmental Defense Fund, 2007. "CDM and the Post-2012 Framework." Discussion Paper, Vienna, August 27-31, AWG/Dialogue.
- [13] Fischer, C., 2005. "Project-based mechanisms for emissions reductions: balancing trade-offs with baselines." *Energy Policy* 33(14): 1807-1823.
- [14] Government of Norway and Government of Indonesia, 2010. Letter of Intent on "Cooperation on reducing greenhouse gas emissions from deforestation and forest degradation."
- [15] Greenpeace and Papua New Guinea, 2008. REDD proposals submitted to the United Nations. Available at <http://unfccc.int>.
- [16] Hartman, R.S., 1988. "Self-Selection Bias in the Evolution of Voluntary Energy Conservation Programs." *Review of Economics and Statistics* 70(3): 448-458.
- [17] Helm, C., 2003. "International emissions trading with endogenous allowance choices." *Journal of Public Economics* 87(12): 2737-2747.
- [18] 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.
- [19] Kerr, S., Hendy, J., Liu, S., Pfaff, A., 2004. "Tropical forest protection, uncertainty and the environmental integrity of carbon mitigation policies." Motu Working Paper 04-03.
- [20] Kindermann, G., Obersteiner, M., Sohngen, B., Sathaye, J., Andrasko, K., Rametsteiner, E., Schlamadinger, B., Wunder, S., Beach, R., 2008. "Global Cost Estimates of Reducing Carbon Emissions Through Avoided Deforestation." *Proceedings of the National Academy of Sciences* 105(30): 10302-10307.
- [21] Kollmuss, A., Lazarus, M., 2011. "Discounting Offsets: Issues and Options." *Carbon Management* 2(5): 539-549.
- [22] Mason, C.F., Plantinga, A.J., 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.
- [23] Melillo, J.M., Reilly, J.M., Kicklighter, D.W., Gurgel, A.C., Cronin, T.W., Paltsev, S., Felzer, B.S., Wang, X., Sokolov, A.P., Schlosser, C.A., 2009. "Indirect Emissions from Biofuels: How Important?" *Science* 326: 1397-1399.
- [24] Millard-Ball, A., 2012. "The trouble with voluntary emissions trading: Uncertainty and adverse selection in sectoral crediting programs." *Journal of Environmental Economics and Management*, forthcoming.
- [25] 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.

- [26] Montero, J.P., 2000. "Optimal design of a phase-in emissions trading program." *Journal of Public Economics* 75(2): 273-291.
- [27] Montero, J.P., 2002. "Prices versus quantities with incomplete enforcement." *Journal of Public Economics* 85(3): 435-454.
- [28] Nepstad, D., Soares-Filho, B.S., Merry, F., Andre, L., Moutinho, P., Carter, J., Bowman, M., Cattaneo, A., Rodrigues, H., Schwartzman, S., McGrath, D., Stickler, C.M., Lubowski, R.N., Piris-Cabezas, P., Rivero, S., Alencar, A., Almeida, O., Stella, O., 2009. "The end of deforestation in the Brazilian Amazon." *Science* 326(5958): 1350-1351.
- [29] Nordhaus, W.D., Yang, Z., 1996. "A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies." *American Economic Review* 86(4): 741-765.
- [30] Pfaff, A., Kerr, S., Lipper, L., Cavatassi, R., Davis, B., Hendy, J., Sanchez-Azofeifa, A., 2007. "Will buying tropical forest carbon benefit the poor? Evidence from Costa Rica." *Land Use Policy* 24(3): 600-610.
- [31] Plantinga, A.J., Richards, K.R., 2010. "International Forest Carbon Sequestration in a Post-Kyoto Agreement." In: Aldy, J.E., and R.N. Stavins, eds. *Post-Kyoto International Climate Policy*. Cambridge University Press.
- [32] Robalino, J., Pfaff, A., Sanchez-Azofeifa, A., Alpizar, F., Leon, C., Rodriguez, C.M., 2008. "Deforestation Impacts of Environmental Services Payments: Costa Rica's PSA Program 2000-2005." Resources for the Future Discussion Paper 08-24, Washington, DC.
- [33] Sanchez-Azofeifa, A., Pfaff, A., Robalino, J.A., Boomhower, J.P., 2007. "Costa Rica's Payment for Environmental Services Program: Intention, Implementation, and Impact." *Conservation Biology* 21(5): 1165-1173.
- [34] Sathaye, J., Makundi, W., Dale, L., Chan, P., Andrasko, K., 2006. "GHG mitigation potential, costs, and benefits in global forests." *Energy Journal* 27(3): 127-162.
- [35] Schneider, L., 2009. "Assessing the additionality of CDM projects: practical experiences and lessons learned." *Climate Policy* 9(3): 242-254.
- [36] Soares-Filho, B.S., Moutinho, P., Nepstad, D., Anderson, A., Rodrigues, H., Garcia, R., Dietzsch, L., Merry, F., Bowman, M., Hissa, L., Silvestrini, R., Maretti, C., 2010. "Role of Brazilian Amazon protected areas in climate mitigation." *Proceedings of the National Academy of Sciences* Early edition. doi:10.1073/pnas.0913048107.
- [37] Stern, N., 2008. "The Economics of Climate Change." *American Economic Review: Papers and Proceedings* 98(2): 1-37.
- [38] Strand, J., 2011. "Carbon offsets with endogenous environmental policy." *Energy Economics* 33(2): 371-378.
- [39] Wise, M., Calvin, K., Thomson, A., Clarke, L., Bond-Lamberty, B., Sands, R., Smith, S.J., Janetos, A., Edmonds, J., 2009. "Implications of Limiting CO₂ Concentrations for Land Use and Energy." *Science* 324: 1183-1186.
- [40] Zhang, J., Wang, C., 2011. "Co-benefits and additionality of the clean development mechanism: An empirical analysis." *Journal of Environmental Economics and Management* 62(2): 140-154.

Appendix A: The Impact of Different Marginal Costs of Avoiding Deforestation on the Performance Measures

Although abatement cost is largely a function of geography and economic factors and not a policy choice, it will influence where avoided deforestation programs will lead to the largest gains. In our model, abatement costs are represented by the foregone net return from deforestation r : the marginal abatement cost curve is determined by the distribution $f_r(r)$.

We first consider which distributions $f_r(r)$ lead to the largest efficiency gain from a 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 channels. 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 $[0, \varepsilon \ll \delta]$ relative to returns between $[\delta - \varepsilon, \delta]$ increases efficiency. The first channel by itself is not sufficient for an overall efficiency gain.

Figure A.1 illustrates this by using three different returns distributions that imply different marginal abatement cost curves. As the distribution changes from case 1 ($N(-1, 0.5^2)$) 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 efficiency gains is ambiguous.

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

Figure A.1: The ambiguous relationship between returns distributions, avoided deforestation and efficiency.

Proposition A.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 A.1).³⁶

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 - $\int_0^{p_c} f'_r(r) dr > \int_0^{p_c} f_r(r) dr \forall r \in [0, p_c = \delta]$.

Proposition A1 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.

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 A.2 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.

Moving from case 1 to case 2 unambiguously raises efficiency and lowers AC . This follows from the statement below Proposition A1, 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. Efficiency and AD fall because the probability mass of returns between $[0, \delta]$ decreases slightly, and the distribution becomes almost flat in the region $[0, \delta]$: 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 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, which reduces the fraction of offsets that are spurious (from 67% to 57%) as well as AC .

³⁶A 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 $f_\varepsilon(\varepsilon) \sim Uniform(-k, k)$ 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.

Figure A.2: Impact of changing $f_r(r)$ on the performance criteria, with $p_c = \delta = 0.5$ and $\sigma_\varepsilon = 0.5$.

Appendix B: Robustness of the Simulation Results to Alternative Returns Distributions

We present figures similar to Figure 4, but for $f_r(r) \sim N(0, 1)$, $f_r(r) \sim \text{Exponential}(1)$ translated to the left by a distance of 1, and $f_r(r) \sim \text{Uniform}(-3, 1)$. We further assume $f_\varepsilon(\varepsilon) \sim N(0, 0.5^2)$, unbiased assigned baselines $\widehat{BL}_{N, \text{unbiased}}$, and $p_c = \delta = 0.5$.

Figure B.1: The impact of baseline scale increases on the performance criteria: $f_r(r) \sim N(0, 1)$.

Figure B.2: The impact of baseline scale increases on the performance criteria: $f_r(r) \sim \text{Exponential}(-1)$ translated to the left by a distance of 1.

The uniform distribution is bounded, which might be more appealing than a distribution that assigns a positive probability to very high returns: such plots should have been cleared already. The qualitative results, however, do not depend on such distributional details:

Figure B.3: The impact of baseline scale increases on the performance criteria: $f_r(r) \sim \text{Uniform}(-3, 1)$.