

Conservation Contracts for Exhaustible Resources

Nils Chr. Framstad Bård Harstad

University of Oslo

In progress - December 2016

For the most recent version, check:

<http://www.sv.uio.no/econ/personer/vit/bardh/dokumenter/cct.pdf>

Abstract

This paper studies how to best incentivize owners to conserve rather than deplete exhaustible resources. This is an important issue when it comes to forest conservation agreements, but it may also become important for other environmental problems, such as climate change. We present a dynamic model where each resource owner benefits from extracting and selling the resource over time. A third party, or principal, is harmed by the extracted amount or she benefits from conservation. The principal can set up payment schedules that incentivize the owners to conserve. We show that the best contract induces the smallest resource stocks to be depleted first, while the largest stock will be extracted from later. To little is conserved permanently and the speed of extraction is too high. These three results are reversed if and only if it is very costly to protect the resource. By comparison, the first best would require that more is conserved, and that extraction begins where the extraction cost is lowest. The difference to the first best is magnified if some buyers boycott the products.

Keywords: Supply-side environmental policies, exhaustible resources, externalities, climate change, deforestation, REDD, PES, contract theory, principal-agent problems, dynamic games.

1 Introduction

This paper derives the optimal contract when multiple agents benefit from depleting and selling their resources on a common market, but extraction harms a third party; the principal.

This problem arises in many situations. For example, the United Nation's REDD program (for Reducing Emissions from Deforestation and forest Degradation) aims at financially compensating developing countries if they are successful in reducing deforestation in the tropics. Unfortunately, there is little theory guiding us when setting up such contracts. This is a serious problem, since deforestation contributes to 10–20 percent of the global carbon dioxide emissions, which cause global warming.¹ *The Economist* claim that negative externalities from forest loss and degradation cost between \$2 trillion and \$4.5 trillion a year according.² Nevertheless, deforestation continues at a rate of 13 million hectares a year (FAO, 2010).

Similarly, when a climate coalition would like to reduce global emissions, it could consider to pay or incentivize the owners of coal or oil reserves to reduce extraction rates. Such a policy implements the first-best in the model analyzed by Harstad (2012). That model is static, however, and it is not clear how contracts should be designed in a dynamic context.

There are also examples beyond environmental economics: For example, the head of an oil cartel (such as OPEC) benefits from a high price, and thus less global production, and it faces the problem of convincing the other cartel members to achieve this goal.

In all these examples, it is important to study the optimal contract design, the amount that will be extracted, and where extraction will be located. To study these issues, we analyze a simple dynamic model where there are multiple owners of resources. Each owner, or "agent," has a private exhaustible resource stock; the stocks are in general of different sizes and they can be associated with different extraction costs and externalities. In every period, an agent can

¹IPCC (2007 and 2013).

²*The Economist*, September 23, 2010.

extract from his stock and sell the extracted amount on the market. The market price is a decreasing function of the aggregate quantity that is extracted in this period.

In addition to these agents, there is a donor or a principal who faces harm when the resource is extracted. The harm could reflect the present-discounted cost of having more greenhouse gases in the atmosphere, or the present-discounted loss of losing tropical forests and biodiversity for all future. As mentioned, it may also be that the principal prefers less extraction for other reasons, such as the high price this would lead to.

The problem for the principal is to offer a schedule of non-negative side transfers that depends on the extraction rates. The principal can make such offers at the start of a period, and we assume that the principal is committed to honor these contracts. In reality, it is reasonable that there is a limit to how much such a principal can commit. Thus, we do not assume that the principal can commit to schedules or contracts in future periods. In this simple game, we restrict attention to a Markov-perfect equilibria (MPEs) since they are simple, robust, and renegotiation-proof. Furthermore, there is a unique such equilibrium in our model.

As a useful benchmark, note that the first best dictates that the a stock should be completely conserved if its extraction cost, plus the externality, is larger than the marginal benefit of consumption. Furthermore, the first best requires that the least expensive resources should be extracted from first, and that the marginal social value of an extracted amount increases with the discount factor.

The equilibrium we derive is very different from both these benchmarks. First, the *smallest* stock will be exhausted first, and completely, before extraction begins with the stock that is somewhat larger. Second, even if the costs are larger than the benefits, the resource may not be depleted but it will be extracted from (unless the aggregate stock is already small). The principal benefits from permitting extraction since this reduces the market price and thus the outside option for all the agents; side payments can then be lowered ac-

cordingly. Third, each stock is either conserved in full or completely exhausted. Thus, although there may be an interior solution for resource extraction at the aggregate level, there is a corner solution in steady state for every *individual* stock: the smallest stocks will be depleted while the largest stocks will remain untouched. We also describe how the equilibrium extraction level declines over time and, thus, the price increases.

Our results point out that the optimal contracts should vary across agents and over time. If the principal instead offered the same conservation incentive to every agent, then the various stocks would be depleted over time in parallel. The principal would then have to pay more, or get less conservation, relative to the optimal contract which treats the agents asymmetrically. Since existing REDD contracts are remarkably similar across countries, our results describe how these can be improved upon.

We can make further predictions regarding the effect of regulating demand. If consumers are taxed when purchasing the extracted amount, then extraction becomes more valuable and more of the resource is conserved. If instead some buyers boycott the good, then the demand curve becomes steeper and the price more sensitive to changes in supply. This motivates the contract-provider to allow for more extraction in order to reduce side payments to the other agents. Consequently, less is conserved if some buyers boycott the harvested good.

The paper contributes to several strands of literature. The classic theory on how to optimally extract from an exhaustible resource is described by Hotelling (1931): the resource price should increase exponentially over time since later consumption has a lower value when the discount rate is positive. When a principal attempts to slow extraction, as in our paper, then we find that the price increases over time also for two very different reasons. First, the principal finds it very profitable to move extraction forward since this reduces today's price (and increases tomorrow's price), which reduces the agents' temptation to extract today. The reduced temptation means that side transfers can be reduced. Second, once the other agents' aggregate stock becomes smaller, the temptation to permit extraction in order to lower side payments to them weakens, and

extraction decreases every time a new stock has been depleted.

We also contribute to the literature on contracts in the presence of externalities. Segal (1999) showed that with positive externalities, the principal's contract will be too weak. Since an agent benefits when the other agents conserve (because the price increases when supply is small), the contract generates a positive externality which the principal cannot exploit. It follows that too little is conserved, in equilibrium, relative to the efficient level. This reasoning is present also in the analysis of REDD contracts by Harstad and Mideksa (2016). However, the models of Segal (1999) and Harstad and Mideksa (2016) are static, and they thus fail to characterize how the stocks are, or should be, gradually and sequentially extracted from over time.

Finally, we contribute to the literature on how to design conservation contracts. There are many theories explaining why there is inefficient deforestation in the first place. See, for example, Alston and Andersson (2011) and Angelsen (2010), and the references therein. For empirical studies of the determinants of deforestation, see Burgess et al. (2011), Damette and Delacote (2012), or, for an earlier overview, Angelsen and Kaimowitz (1999). The literature on conservation contracts is small, but growing. For overviews, see Engel et al. (2008) or Kerr (2013), which reveals that the standard approach is to focus on textbook contract-theoretic problems (as surveyed by Bolton and Dewatripont, 2005) such as moral hazard (Gjertsen et al., 2010), private information (Chiroleu-Assouline et al., 2012; Mason, 2013; Mason and Plantinga, 2013), or observability (Delacote and Simonet, 2013). Our approach is more related to the problems that arise when contracts lead to externalities. Relative to that literature, which is discussed above, our model is dynamic and thus captures effects that so far has not been discussed in the literature.

The next section presents our dynamic model and derives some benchmark results. Section 3 derives the equilibrium and our main results, while Section 4 shows when the results are reversed if it is costly to protect the resource (from illegal extraction). Section 5 concludes.

2 A Model of Conservation Contracts

The resource owners: There are n agents and an infinite number of periods. At the beginning of period $t \geq 0$, agent $i \in N \equiv \{1, \dots, n\}$ has a stock of exhaustible resource measured by $y_i^t \geq 0$ and he can extract $x_i^t \in [0, y_i^t]$, leaving $y_i^{t+1} = y_i^t - x_i^t$ for the next period.

The extraction levels are decided on simultaneously and the price depends on the aggregate level of extraction. We will below focus on the design of contracts, rather than the market per se, so we here simplify by assuming that the demand for $x^t = \sum_{i \in N} x_i^t$ is linear and the price p^t is given by:

$$p^t = b - ax^t.$$

So, $b > 0$ is the benefit of the first marginal unit of x^t , while $a > 0$ is the slope of the demand curve.

If we let $c_i \geq 0$ be the unit cost of depleting i 's stock, i 's payoff in period t is given by

$$u_i^t = (p^t - c_i) x_i^t + s_i^t,$$

where $s_i^t \geq 0$ is a possible side payment. Every agent seeks to maximize the present discounted sum $\sum_{t' \geq t} \delta^{t'-t} u_i^{t'}$, where the common discount factor is $\delta \in (0, 1)$.

The principal: In addition to the agents, there is a principal who is harmed by the extraction. We let $e_i > 0$ measure the externality of x_i^t on the principal. The subscript indicates that we allow the externality to vary across the stocks. Some stocks may be more valuable to the principal, or they may generate more emission when extracted. The total environmental harm at time t is then $\mathbf{e} \cdot \mathbf{x}^t$, where $\mathbf{e} \equiv (e_1, \dots, e_n)$ and $\mathbf{x}^t \equiv (x_1^t, \dots, x_n^t)$.

For simplicity, we assume that the agents are exporters and do not internalize the consumer surplus. The principal may or may not internalize the consumer surplus: we let the parameter $w \in [0, 1]$ measure the fraction of the consumer surplus (which is $a(x^t)^2/2$) that is internalized by the principal. Since the side transfers to the agents, $s^t \equiv \sum_{i \in N} s_i^t$, are paid by the principal, the payoff to the

principal in period t is

$$u_0^t = -\mathbf{e} \cdot \mathbf{x}^t - s^t + w \frac{a}{2} (x^t)^2.$$

Thus, we assume that the principal is harmed by the extraction levels. If the principal benefitted from the stock of the resource at any point in time, by receiving the benefit $v_i y_i^t$, then we can set $e_i \equiv v_i \delta / (1 - \delta)$ and interpret e_i as the present discounted loss when another unit of the resource is extraction. Similarly, we can set $e_i \equiv v_i / (1 - \delta)$ also if the principal's concern is the cost of accumulated extraction levels, when this cost can be written as $w \sum_{\tau=0}^{t-1} x^\tau$.

In other words, our model captures well the principal's concern, whether this concern relates to the remaining forest stock, the accumulated greenhouse gases, or whether the principal simply likes the high price p^t associated with a small x^t .

The contracts: At the start of each period, the principal offers each agent a payment schedule $s_i^t(\mathbf{x}^t) \geq 0$, $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$, where the principal commits to pay i the amount $s_i^t(\mathbf{x}^t)$ if the vector of extraction levels is given by \mathbf{x}^t . We can without loss of generality assume that the principal simply suggests a particular vector, $\widehat{\mathbf{x}}^t$, as well as a set of side payments (scalars s_i^t) to be paid if the outcome is exactly this vector $\widehat{\mathbf{x}}^t$, and otherwise promises zero side transfers.

The principal cannot commit to future contracts, and we restrict attention to Markov-perfect equilibria. When deriving the optimal contract in period t , the agents' best outside option will be important. If an agent contemplates to select any other $x_i^t \neq \widehat{x}_i^t$, he will receive no side payment and he can thus ignore the entire contract. In order to simplify, we will *start out* assuming that the best outside option for every agent is simply to deplete the stock and extract it all once and for all. This assumption could be justified if any deviation from the agreement will induce the other agents to produce a lot in the following periods (perhaps because the principal will no longer pay for reduced extraction levels) since then an agent can expect that the price will be even smaller in the future than in the period in which it contemplates to deviate. Alternatively, we can check whether it will be the case that, in equilibrium, the best outside

option is to extract it all, once and for all. It turns out that this is indeed the best outside option if the stocks are small enough. Then, the marginal revenue when extracting the last unit today is larger than the highest possible present-discounted value of extracting a marginal unit in the future (which is δb).

3 Analysis and Results

As a benchmark, consider the first best, maximizing aggregate utility. If $b < c_i + e_i \forall i$, the first best is straightforward: Nothing at all should be extracted, since the highest possible value is smaller than the principal's loss. If instead $b > c_i + e_i$ for some i , it is socially optimal to extract at least a marginal amount of x_i^t . Thus, eventually, and in finite time, we reach a stage when $y_i^t = 0$, so that nothing is left. Naturally, it is socially optimal to first extract the resources that can be extracted at the smallest social costs ($c_i + e_i$).

Proposition 0 *The first-best outcomes have the following properties.*

- (i) *The stocks with the smallest social costs $c_i + e_i$ are extracted from first.*
- (ii) *The steady state is reached in finite time T and, then, $y_i^T = 0$ for every stock for which $c_i + e_i < b$, while $y_i^T = y_i^0$ for every stock for which $e_i + c_i \geq b$.*
- (iii) *Before the steady state is reached, the marginal social surplus increases exponentially over time according to:*

$$b - c_i - e_i - a(2 - w)x^t \geq \delta [b - c_i - e_i - a(2 - w)x^{t+1}], \quad (1)$$

if $x_i^t > 0$. The equation (1) holds with equality if also $x_i^{t+1} > 0$, and otherwise $x_i^t = y_i^t$.

Proof. To be added. ■

Note that part (1) simplifies to the standard Hotelling rule

$$p^t = \delta p^{t+1}$$

in the special case in which $w = 1$ and $c_i = e_i = 0$. If instead $w = 0$, and $e_i = 0$, then (1) simplifies to the monopolists' rule of letting marginal revenues grow with exponentially (with the interest rate $1/\delta$).

Another interesting benchmark is the non-cooperative free-market equilibrium if everyone took prices as given. It is straightforward to show that the market equilibrium will then be given by Proposition 0 if just every e_i is replaced by zero. That is, the externalities will not be taken into account, but otherwise the outcome is efficient in that the least costly stocks will be depleted first, for example.

When the agents are not taking the prices as given in the non-cooperative equilibrium, then it is intuitive that several stocks may be extracted from simultaneously.³

3.1 The sequence of extraction

We are now ready to describe the equilibrium when the principal offers conservation contracts that maximize her own utility. Following the sequence of the three parts in Proposition 0, we will (1) describe the equilibrium ordering of how the stocks are being depleted, (2) the steady state, and (3) the extraction path over time before the steady state is reached.

Proposition 0 indicates that the first best cannot pin down the allocation of the x_i^t 's, but only the aggregate x^t , if c_i and e_i are the same for every country. In this case, it is irrelevant, as far as efficiency is concerned, whether one extracts from this or that stock first—or all stocks at the same time. In equilibrium, however, the principal is unlikely to be indifferent. Instead, the

³Deriving the non-cooperative outcome is not the purpose of the present paper, but it is simple to see that if $\delta = c_i = 0$ (i.e., there is only one relevant period), the non-cooperative outcome would be given by the following:

$$\begin{aligned} x_i &= \max \left\{ \frac{b - a \sum_{j \neq i} x_j}{2a}, 0 \right\} \\ &= \frac{b}{a(n+1)} \text{ if } \frac{b}{a(n+1)} \leq y_i^0 \forall i \in N. \end{aligned}$$

In other words, the solution is interior and every agent will extract the exact same amount (unless some stocks are too small).

principal will prefer to allow for extraction from the smallest stock only, while paying every other stock to be conserved until the smallest stock is completely depleted. Thus, the smallest stock will be depleted first.

An intuition for this result is that it is very costly for the principal to deal with multiple districts. Conserving in one district reduces aggregate supply and the other agents become more tempted to extract. This temptation is lowered when the principal lets i extract a little bit more, since the market price decreases accordingly. The reduced price means that the principal can reduce every other agent's side payment, and the total reduction in side payment is proportional to the aggregate stock—minus the stock of district i (i.e., $y_{-1}^t = y^t - y_i^t$). Conditional on y^t being the same, the stock y_{-i}^t is largest when i has the smallest stock.

This reasoning continues to hold even when the stocks c_i and e_i differ. As long as these differences are relatively small, the principal still prefers to extract from the smallest stock first.

Proposition 1 *The stocks with the smallest sizes are depleted first: Stock i is depleted before stock j is extracted from if*

$$y_j^0 - y_i^0 \geq \left(\frac{1 - \delta}{a} \right) (c_i - c_j + e_i - e_j). \quad (2)$$

Proof. The side transfers s_i^t must be set such as to compensate i for delaying exhausting i 's resource. Thus, at any point in time, we have $s_i^t = (b - c_i - ax_{-i}^t - ay_i^t) y_i^t - (b - c_i - ax_{-i}^t - ax_i^t) x_i^t - \delta (b - c_i - ax_{-i}^{t+1} - ay_i^{t+1}) y_i^{t+1}$, so

$$\begin{aligned} s_i^t + \delta s_i^{t+1} &= (b - c_i - ax_{-i}^t - ay_i^t) y_i^t - (b - c_i - ax_{-i}^t - ax_i^t) x_i^t \\ &\quad - \delta (b - c_i - ax_{-i}^{t+1} - ay_i^{t+1}) y_i^{t+1} - \delta^2 (b - c_i - ax_{-i}^{t+2} - ay_i^{t+2}) y_i^{t+2}. \end{aligned} \quad (3)$$

Departing from any contract x , consider a slight perturbation \tilde{x} given by $\tilde{x}_i^t = x_i^t + \epsilon$, $\tilde{x}_i^{t+1} = x_i^{t+1} - \epsilon$, $\tilde{x}_j^t = x_j^t - \epsilon$, and $\tilde{x}_j^{t+1} = x_j^{t+1} + \epsilon$, while every other extraction level stays unchanged. If $\epsilon \rightarrow 0$, the marginal increase in total side

payments $\sum_{t' \geq t} \sum_{j \in \{1, \dots, n\}} \delta^{t'-t} s_j^t$ is given by the following:

$$\begin{aligned} \frac{\partial \sum_j s_j^t + \delta s_j^{t+1}}{\partial \epsilon} &= a y_i^t - a y_j^t - (b - c_i - a x^t) + (b - c_j - a x^t) \\ &\quad + \delta (b - c_i - a x^{t+1}) - \delta (b - c_j - a x^{t+1}) \\ &= a (y_i^t - y_j^t) + (1 - \delta) (c_i - c_j) \end{aligned}$$

In addition, the principal receives the environmental benefit $(1 - \delta) (e_j - e_i)$, so her total the benefit is

$$a (y_j^t - y_i^t) + (1 - \delta) (e_j + c_j - e_i - c_i).$$

It follows that when this expression is positive, the principal benefits from extracting from i first: Suppose for contradiction the contrary, i.e., $x_i^\tau \cdot x_j^t > 0$ for $\tau > t$. Make the above perturbation at $t, t+1, \dots, \tau-1$ (cancelling out for times $t+1, \dots, \tau-1$). Then the principal has a strict improvement, and the strategy could not be optimal. Finally, we can let $\tau \rightarrow +\infty$ to improve on the case where $x_j^t > 0$ while $x_i^\tau = 0$ for all $\tau > t$ (provided that $y_\tau > 0$). ■

3.2 The steady state

Based on the previous subsection, it is natural to order all the stocks according to decreasing values of size or, more precisely, $y_j^0 + \left(\frac{1-\delta}{a}\right) (c_j + e_j)$, so that $y_j^0 + \left(\frac{1-\delta}{a}\right) (c_j + e_j) > y_{j+1}^0 + \left(\frac{1-\delta}{a}\right) (c_{j+1} + e_{j+1})$ for every $j \in N \setminus n$. Thus, stock n is depleted first, and then stock $n-1$ is depleted before stock $n-2$, etc. Some stocks may be conserved forever. Our next result characterizes the steady state, or the long-run outcome, of the game, corresponding to the second part of Proposition 0. It is convenient to define $Y_i^t \equiv \sum_{j \leq i} y_j^t$ as the sum of the stocks larger than stock i .

Proposition 2 *A unique steady state is reached in finite time. In the steady state, the \hat{i} stocks with the largest $y_j^0 + \left(\frac{1-\delta}{a}\right) (e_j + c_j)$ are conserved while the stocks with smaller $y_j^0 + \left(\frac{1-\delta}{a}\right) (e_j + c_j)$ are exhausted. A necessary condition characterizing the threshold is*

$$y_{\hat{i}}^0 + \frac{e_{\hat{i}} + c_{\hat{i}} - b}{a} \geq Y_{\hat{i}}^0 \geq \frac{e_{\hat{i}+1} + c_{\hat{i}+1} - b}{a}, \quad (4)$$

if $\hat{i} \in \{1, n-1\}$, while the second inequality is skipped if $\hat{i} = n$, and $\hat{i} = 0$ if $e_1 + c_1 - b < 0$. Condition (4) is also sufficient if the heterogeneity in $e_i + c_i$ is relatively small, i.e., if for every i ,

$$(e_i - e_{i-1}) + (c_i - c_{i-1}) \notin \left(ay_{i-1}^0, \frac{a}{1-\delta} (y_{i-1}^0 - y_i^0) \right).$$

Proof. The claim that steady-state is reached in a finite number of steps, will follow from Proposition 3.

Define $\pi_i(x) = x(b - ax - c_i)$. Conserving every stock forever requires that, in every period, $s_i^t = (1 - \delta)\pi_i(y_i^t)$ for every i . By Proposition 1, consider the remaining stock with the largest number i . Suppose the principal considers to extract from i where the aggregate size of the larger/remaining stocks is $Y_{i-1}^t = \sum_{j < i} y_j^t$. Extracting a marginal unit of size $x_i^t = \epsilon$ reduces s_j^t by ϵay_j^t units in this period, while the present-discounted value of the future s_i 's will be reduced from $\pi_i(y_i^t)$ to $\pi_i(y_i^t) - \pi_i(\epsilon) - \delta\pi_i(y_i^t - \epsilon) + \delta\pi_i(y_i^t - \epsilon)$. This marginal extraction is not strictly beneficial to the principal if

$$\epsilon e_i \geq \pi_i(\epsilon) + \sum_{j \neq i} \epsilon ay_j^t,$$

and, when $\epsilon \rightarrow 0$, this condition becomes:

$$e_i \geq b - c_i + \sum_{j \neq i} ay_j^t \Leftrightarrow \sum_{j \neq i} y_j^t = Y_i^t - y_i^t \leq \frac{e_i + c_i - b}{a}.$$

When is this condition also sufficient? Suppose it is not: That is, even if the condition holds for some i ,

$$a(Y_i^t - y_i^t) \leq e_i + c_i - b,$$

it does not hold for $i - 1$, because

$$a(Y_i^t - y_{i-1}^t - y_i^t) \geq e_{i-1} + c_{i-1} - b.$$




Figure 1: *The largest stocks are conserved, while the smallest stocks are depleted.*

Combined, this means,

$$\begin{aligned}
e_i + c_i - b &\geq a(Y_i^t - y_i^t) \geq ay_{i-1}^t + e_{i-1} + c_{i-1} - b \Rightarrow \\
e_i + c_i &\geq ay_{i-1}^t + e_{i-1} + c_{i-1} \Rightarrow \\
e_i + c_i - (e_{i-1} + c_{i-1}) &\geq ay_{i-1}^t \Rightarrow \\
\frac{a}{1-\delta}(y_{i-1}^t - y_i^t) &\geq e_i + c_i - (e_{i-1} + c_{i-1}) \geq ay_{i-1}^t,
\end{aligned}$$

where we have added the condition from Proposition 1. For both equalities to hold, we must have $\delta y_{i-1}^t \geq y_i^t$, which is violated for δ small. ■

If $c_i + e_i < b\forall i$, everything will eventually be extracted, just as in the first best. If $c_i + e_i > b\forall i$, however, the first best dictates that everything should be conserved, but this will be true in equilibrium only if $c_n + e_n - b > aY_{n-1}^0$. Note that the conservation condition coincides with the first-best condition if and only if there is only a single agent, since then $Y_{n-1}^0 = Y_0^0 = 0$.

If $n > 1$, however, the principal will in general permit extraction even if this is socially suboptimal. The reason is that by allowing for a little bit of extraction, say $x_i^t = \epsilon > 0$, the equilibrium price falls and this reduces every agent's temptation to opt for the outside option. Starting from the steady state, the net side payment to i himself can be reduced by the profit he makes, $\epsilon(b - a\epsilon)$. In addition, $j \neq i$ realizes that the price is reduced by $a\epsilon$ units, and thus j 's reduced temptation to extract implies that the principal can reduce the

side payment s_j^t without violating j 's incentive to conserve: the sum $\sum_{j \neq i} s_j^t$ can be reduced by $a\epsilon \sum_{j \neq i} y_j^t$. Thus, the larger is $\sum_{j \neq i} y_j^t$, the larger is the benefit of allowing $x_i^t = \epsilon > 0$. Since $\sum_{j \neq i} y_j^t$ is largest for the smallest possible y_i^t , we confirm the intuition that it is optimal to first let the smallest agent extract. Furthermore, since $\sum_{j \neq i} y_j^t$ stays unchanged as y_i^t is depleted, it is optimal to completely exhaust i 's stock if it is optimal to start extracting from him in the first place. This explains why the steady state is always a corner solution where every stock is either left untouched or completely exhausted, as illustrated in Figure 1. However, once there is a sufficiently small number of stocks left, the cost-savings $a\epsilon \sum_{j \neq i} y_j^t$ of allowing for $x_i^t > 0$ is reduced and it may eventually be optimal for the principal to pay for conservation of every remaining stock.

Note that the condition for sufficiency always hold when $c_i + e_i$ is the same for every country. In this case, the proposition simplifies.

Corollary 1 *Suppose $c_i = c$ and $e_i = e$.*

- (i) *The first best is that nothing is extracted if $b < c + e$.*
- (ii) *In equilibrium, every stock is depleted except for the $\hat{\imath}$ largest stocks, where $\hat{\imath}$ is given by*

$$y_{\hat{\imath}}^0 + \frac{c + e - b}{a} \geq Y_{\hat{\imath}}^0 \geq \frac{c + e - b}{a}$$

- (ii) *With a large number of small stocks ($y_{\hat{\imath}} \rightarrow 0$), the steady-state conservation level is*

$$y^T = \max \left\{ 0, \frac{c + e - b}{a} \right\}.$$

The last equation emphasizes (as should also be clear from formula (4)) that more is conserved if e is large, if b is small, and if a is small. As noted already, the large a means that the s_j^t 's can be reduced by quite a lot if $x_i^t > 0$, and it then becomes beneficial for the principal to allow for more extraction.

3.3 The speed of extraction

We are now ready to characterize the extraction path over time. As hinted to already, the principal prefers to allow for more extraction than what is socially optimal since this reduces every agent's temptation to go for the outside option and extract more. This reduced temptation means that the principal can lower the side payments to these agents, without fearing that they will extract more. In line with the intuition above, the equilibrium path coincides with the first-best path if there is only one stock left in the game, but otherwise the speed of extraction will be too high.

Proposition 3 *Let l^t denote the largest i such that $x_i^t > 0$.*

- (i) *Extraction x^t is decreasing in t .*
- (ii) *The marginal social surplus increases more than exponentially in time. For any two consecutive periods, we have:*

$$[b - e_l - c_l - a(2 - w)x^t] + ay_{-l}^t \geq \delta [b - e_l - c_l - a(2 - w)x^{t+1}] \quad (5)$$

with equality if $x_1^t < y_1^t$.

Proof. Similarly to the proof of Proposition 1, consider a small perturbation, now referred to as \hat{x} , where $\hat{x}_i^t = x_i^t + \epsilon$ and $\hat{x}_i^{t+1} = x_i^{t+1} - \epsilon$, while every other extraction level stays unchanged. With (3), the marginal increase in total side payments $\sum_{t' \geq t} \sum_{j \in \{1, \dots, n\}} \delta^{t'-t} s_j^t$ (when $\epsilon \rightarrow 0$) is given by the following:

$$\begin{aligned} \frac{\partial \sum_j s_j^t + \delta s_j^{t+1}}{\partial \epsilon} &= -ay_{-i}^t - (b - c_i - ax_{-i}^t - 2ax_i^t) + ax_{-i}^{t+1} \\ &\quad + \delta (b - c_i - ax_{-i}^{t+1} - 2ax_i^{t+1}) - \delta ax_{-i}^{t+1} \\ &= -ay_{-i}^t - (b - c_i - 2ax^t) + \delta (b - c_i - 2ax^{t+1}) \end{aligned}$$

Moving extraction forward in this way is also implying an environmental loss to the principal given by $(1 - \delta)e_i$, and a consumer surplus gain given by $wa(x^t - \delta x^{t+1})$. Thus, the total benefit of moving forward a marginal amount is

$$\begin{aligned} wa(x^t - \delta x^{t+1}) - (1 - \delta)e_i + ay_{-i}^t + (b - c_i - 2ax^t) - \delta (b - c_i - 2ax^{t+1}) &\geq 0 \Rightarrow \\ b - e_i - c_i - a(2 - w)x^t &\geq -ay_{-i}^t + \delta [b - e_i - c_i - a(2 - w)x^{t+1}]. \end{aligned}$$

Clearly, this equation holds with equality if it is optimal with $x_i^t x_i^{t+1} > 0$, while if the equality is strict, then $x_i^{t+1} = 0$. ■

One should remember here that $[b - e_l - c_l - a(2 - w)x^t]$ is simply the marginal social value when extracting one more unit of stock l . In the special case in which $c_l = 0$ and $w = 1$, this term equals $p^t - e_l$, and the equation then becomes $(p^t - e_l) + aY_{l-1}^t = \delta(p^{t+1} - e_l)$ if $x_i^t x_i^{t+1} > 0$. So, the term $(p^t - e_l)$ increases very fast in t , not only because of the discount factor, but also thanks to the term aY_{l-1}^t : the principal finds it optimal to move extraction forward in time to save on side transfers to the other agents. This is particularly beneficial to the principal when there are many other agents/stocks around. Hence, the speed of extraction is more different from the first best if Y_{l-1}^t is large, i.e., early in the game.

As corollaries, we obtain the following:

Corollary 2

- (i) *The speed of extraction ($x^t - \delta x^{t+1}$) is too large relative to the first best if and only if $n > 1$.*
- (ii) *The extraction level x^t is piecewise concave in t , and the price p^t is piecewise convex in t .*

Another comparison between the first best and the equilibrium path concerns the concavity of x_t , as a function of t . Suppose $c_i = c$ and $b_i = b$. In the first best, x_t is always concave, implying that the equilibrium price is convex, as a function of t . In equilibrium, however, x_t is concave only as long as we keep on extracting from the same agent. The concavity can be seen from Proposition 3 by noting that x^t is a weighted average of x^{t+1} and $\frac{y_{-l}^t}{(1-\delta)(2-w)} - \frac{e_l + c_l - b}{2a}$. As soon as $y_l^t = 0$, however, and extraction begins with $x_{l-1}^t > 0$, then x^t becomes a weighted average of x^{t+1} and $\frac{y_{-(l-1)}^t}{(1-\delta)(2-w)} - \frac{e_l + c_l - b}{2a}$, where the first

term is reduced. Thus, x^t , as a function of t , becomes piecewise—and not globally—concave. As a consequence, the price becomes piecewise—but not globally—convex as a function of t .

3.4 Comparative static and policy implications

There are several interpretations of the model and of who the principal may be. The principal could be the World Bank, the UN, or individual countries such as Norway. However, when it comes to forest and land conservation, also non-governmental organizations (NGOs) are offering conservation contracts. Clearly, the preferences of the principal may be different if the principal is an NGO or whether it should be interpreted as the UN. If the principal is "benevolent" and takes the agents' profit into account, then the outcome is clearly first best, as described by Proposition 0. In other words, suppose the principal has a preference not only for conservation, but also for the resource-owners profit (by maximizing $v_0^t = \sum_{i=0}^n u_0^t$). Then, perhaps paradoxically, the principal decides to conserve *more*—not less. Intuitively, when the principal does not internalize the profit to the agents, then she permits more extraction since this reduces the agent's outside-option payoffs and thus the side transfers which the principal must pay.

Corollary 3. *Suppose the principal internalizes the profits u_i^t 's. Compared to the results above,*

- (i) *extraction may switch from a small stock j to a large stock i if just $c_j + e_j > c_i + e_i$;*
- (ii) *the stock that is conserved forever increases.*

In the model above, the market for the extracted good is parametrized by a and b . If demand decreases, b will decrease or a will increase. Policies that influence the market may influence b or a , depending on the policy instrument. A boycott, where a fraction of the buyers will exit the market, will be similar

to a larger a . The explanation is particularly straightforward when all the m buyers are identical and associated with a demand function $b - a'$. The total demand is then given by $b - a$, where $a \equiv a'/m$. Clearly, a smaller m implies that a will increase.

Note that the larger a is beneficial to the principal. The reason is that when a is large, then when the principal permits one agent to extract a marginal amount, the side payments to the other agents are lowered by quite a lot (namely, by an amount proportional to a). While this effect is beneficial to the principal, it also motivates her to take advantage of this effect by indeed allowing an agent to extract a little bit more. This means that a larger amount is eventually extracted, and the steady state level that is conserved forever is reduced.

Corollary 4. *A boycott increases a and has the following consequences:*

- (i) *extraction may switch from a large stock j to a small stock i , even if $c_j + e_j < c_i + e_i$;*
- (ii) *the stock that is conserved forever declines.*

In contrast, an emission tax, or a consumption tax, will be similar to a reduction in b : that is good for efficiency and conservation as well as for the principal, it turns out. A smaller b will reduce the temptation to extract and the principal will be able to motivate conservation at a smaller price. In equilibrium, a larger fraction of the resource will be conserved forever.

Corollary 5. *A consumer tax decreases b and has the following consequences:*

- (i) *there is no influence on whether stock i or stock j will be depleted first, as long as both will be depleted;*
- (ii) *the stock that is conserved forever becomes larger.*

3.5 Alternative contracts

Above, we allowed the principal to set the side payment schedules as arbitrary functions of the vector of extraction levels. Thus, the principal cannot do better than with these contracts, if the agents are free to opt out and cannot be forced to participate. However, the principal can certainly be much worse off if there are (realistic) constraints to the contracts that are offered.

Real-world REDD contracts have two features. First, they are linear, in that they specify a baseline level \bar{x}_i^t and a rate ϕ_i^t such that the subsidy is linear in the amount that deforestation is reduced relative to this baseline level: $s_i^t = \phi_i^t (\bar{x}_i^t - x_i^t)$. While the incentives will be influenced by the rate ϕ_i^t , the baseline will not affect x_i^t , on the margin, but just ensure that i does not prefer to ignore the contract. In principle, these linear subsidy schemes can achieve the same outcome as the general contracts offered above. However, to ensure that $x_i^t > 0$ for the smallest stock, while $x_j^t = 0$ for the largest stock, it may be necessary to set $\phi_j^t > \phi_i^t$. Furthermore, ϕ_i^t may need to be reduced in time, t , in order to implement the increase in x_i^t which we have proved to be optimal above.

The second feature of real-world contracts is that they are remarkably similar across countries. In particular, the subsidy rate ϕ_i^t is fixed over i . For the REDD contracts offered by the Norwegian government, for example, ϕ_i^t is for every country set equal to the amount of avoided carbon dioxide emissions, multiplied by five dollar a ton. When ϕ_i^t is the same for every i , then it is reasonable that when $x_i^t > 0$, we also have $x_j^t > 0$. Thus, several resource stocks will be depleted in parallel, even though the optimal contract requires that the smallest stock is depleted before one begins extraction from a larger stock. In other words, requiring $\phi_i^t = \phi_j^t$ leads to too little conservation conditional on the money that is spent, or too much money being spent conditional on the amount of conservation that is achieved. Our results illustrate how real-world contracts can be improved upon by letting the contracts vary over time as well as over districts.

Corollary 6

- (i) *The optimal contracts vary across countries and change when resources are being depleted.*
- (ii) *Harmonized contracts achieve too little conservation and at too high costs to the principal.*

4 Alternative Outside Options and Models of Extraction

This section extends the model above by (i) allowing resource extraction to be illegal or protection costly and (ii) letting the outside option be different than simply depleting the entire resource once and for all. We study these two generalizations together because, it turns out, none of them will make much of a difference in isolation. Combined, however, they shed light on the limitations of the above results. To concentrate on the new results, we here assume agents are identical except for size, so that the principal's direct utility depends only on aggregates ($c_i = c$ and $e_i = e$). (We remark though, that it is possible to generalize the result to heterogeneous agents by modifying the size ranking criterion in a similar way to the previous section.)

Costly protection. For some resources, such as tropical forests, extraction may not be easily regulated by the local governments. Instead, there may be a lot of pressure from illegal loggers on the forest that is to be conserved. The profit from illegally cutting a unit of the forest is the equilibrium price minus the extraction cost, $p^t - c$, and this must be compared to the expected penalty which the loggers are facing. The expected penalty can be raised by raising the fine or the number of years in jail, but there is a limit to how much the penalty can be increased in this way. For a given penalty level, the expected penalty is instead raised by increasing the monitoring probability. Since this is costly, we can let $\alpha \geq 0$ measure the cost of raising the expected penalty. To be effective, the expected penalty must be at least as large as the profit when cutting, $p^t - c_i$,

and in equilibrium it will be exactly so large. Thus, if agent i wants to protect the remaining amount $y_i^t - x_i^t$, i 's payoff in a period is given by the following:

$$u_i^t = \beta (p^t - c) x_i^t - \alpha (p^t - c) (y_i^t - x_i^t) + s_i^t,$$

where parameter $\beta \geq 0$ measures the benefit of sales revenues, while α measures the cost of monitoring the remaining forest. This utility function is better defended in Harstad and Mideksa (2016) who study contracts in a static model.

Outside option. An important simplifying assumption above was to suppose that if an agent deviates from the principal's contract, then he sells the entire stock, once and for all. We presented a condition suggesting that this assumption could be reasonable when the stocks were small, and argued that it could also hold if one could expect "everything to break down" if an agent cheats and breaks the contract. More generally, however, one should allow an agent to consider any extraction level $x_i^t \leq y_i^t$ at any point in time. At the moment, we are working on such a model, where an agent can deviate by any amount he wants before entering a new period where the principal re-optimizes by offering new contracts. In this section, however, we will grossly simplify by assuming that an agent's outside option is to extract x_i^t and leave unextracted $y_i^{t+1} = y_i^t - x_i^t = F(y_i^t)$, where $F(\cdot)$ is an exogenous increasing function which takes values on $[0, y_i^t]$. This function could measure some capacity constraints, and we thus assume that also the principal must take the constraint $x_i^t \leq y_i^t - F(y_i^t)$ into account. Although assuming an exogenous F is a severe shortcut, it is still a generalization relative to the analysis above, were $F(y_i^t) = 0$ for any y_i^t . If an agent has deviated in a particular period, then the principal will re-optimize by suggesting new contracts in the following period. After all, we consider Markov-perfect strategies and deviations from the contract in the past is not payoff relevant when going forward

None of these extensions would make much of a difference in isolation. With the outside option given by (a sufficiently well-behaved) $F(\cdot)$, it would still be the case that extraction will start with the smallest stock if $\alpha = 0$. The same would be true if we allowed for $\alpha > 0$ but had $F(\cdot) = 0$. Together, however, it

turns out that the principal may want to extract from the largest stock first if $F'(\cdot)$ and α are sufficiently large.

The intuition for why Proposition 1 may be overturned is the following. Remember that the principal receives the sum of payoffs minus every agent's outside option. If the protection cost α is large, then it is costly for agent i to protect the remaining amount $F(y_i^t)$, but this cost is reduced if the price is smaller, e.g., if another agent j is allowed to extract. Thus, there is a positive (pecuniary) externality from x_j^t to i 's outside option. This externality is larger if $F(y_i^t)$ is large, so, if y_i^t is large. Since improving i 's outside option is costly to the principal, it is better for her to instead let the large district extract. In short, the (pecuniary) externality from j 's extraction on i 's outside option is positive when α is large, but it is negative when α is small, as in Sections 2-3. This difference overturns the result from Proposition 1.

Proposition 4 *Let $Z^t \equiv [0, \max y_i^t]$ and assume that $F'(z) \in [0, 1/\delta) \forall z \in Z^t$. Then, the smallest stocks are extracted first at t if*

$$\alpha < \underline{\alpha} \equiv \beta \cdot \min_{z \in Z^t} \left\{ \frac{1 - \delta F'(z)}{[1 - \delta F'(F(z))] \cdot F'(z)} - 1 \right\},$$

while the largest stocks are extracted first if

$$\alpha > \bar{\alpha} \equiv \beta \cdot \max_{z \in Z^t} \left\{ \frac{1 - \delta F'(z)}{[1 - \delta F'(F(z))] \cdot F'(z)} - 1 \right\},$$

where both ratios are interpreted as $+\infty$ if $F'(z) = 0$.

Proof. Since we consider homogeneous agents, we can normalize c_i to zero without loss of generality. Consider the following perturbation of an arbitrary strategy: at time t , increase the contractual extraction x_i^t from a single agent i , and, provided the agent takes the contract, decrease extraction in the next period – and for some agent $j \neq i$, perform the opposite perturbation. That is, we consider the marginal effect of a partial decrease in y_i^{t+1} and a partial increase (of same size) in y_j^{t+1} . By the envelope theorem we can disregard everything from period $t+2$ on, and because aggregates are kept constant, neither prices nor the principal's direct utility change, nor the sum of the agents' direct utilities

in the contract. Neither do the outside values change for any other agents than i and j .

The principal will therefore minimize the sum of those agents' outside options to minimize side transfers. Agent i 's outside option will benefit from agent j 's initially reduced supply, but take a hit from the subsequent increased supply, a net effect of a times $\beta y_i^t - (\beta + \alpha)F(y_i^t) - \delta[\beta F(y_i^t) - (\beta + \alpha)F(F(y_i^t))]$. The principal's transfer expenditure will thus increase by the difference between this, and the corresponding expression for j . It suffices to differentiate wrt. stock:

$$\begin{aligned} & \frac{d}{dz} [\beta z - (\beta + \alpha)F(z) - \delta[\beta F(z) - (\beta + \alpha)F(F(z))]] \\ &= \beta[1 - \delta F'(z)] - (\beta + \alpha)[1 - \delta F'(F(z))]F'(z) \end{aligned}$$

which is positive iff $\beta + \alpha < \beta[1 - \delta F'(z)]/([1 - \delta F'(F(z))]F'(z))$, by the sign of F' and $1 - \delta F'(F(z))$. The conclusion follows. ■

5 Conclusions

This paper analyses a model in which a principal tries to slow or reduce the agents' extraction levels of their privately owned exhaustible resources. This is relevant for conservation contracts regarding tropical forests, for example, but it may also become an important policy tool when dealing with other problems, such as climate change.

Even when the first best dictates that nothing or everything should be extracted in the long run, the equilibrium outcome is likely to be interior. While the first best requires that one should first extract from the stocks that have the lowest emissions or extraction costs, it is the smallest stock(s) that will be depleted first, in equilibrium. The principal's optimal contract is very asymmetric in that it treats relatively similar resource owners very differently. This suggests a way of improving existing conservation contracts on deforestation, which currently treat every country in the same way. In particular, if the subsidy rate for reducing deforestation is the same for every country, then too little is conserved for every dollar that is spent, relative to the efficient contract. We

have also noted that other policies, such as boycotts, are likely to worsen the problem and encourage a contract-provider to conserve less and to switch extraction from stocks that are inexpensive to stocks that are instead small (even though they may be associated with larger emissions and extraction costs).

Our model is just a first cut when studying optimal conservation contracts in dynamic settings. We make a large number of strong assumptions that ought to be relaxed in future research. In particular, one should allow for alternative outside options than we have done in this analysis, and one may also want to study more carefully the market for the extracted product. In this way, future research will bring us another few steps forward in understanding the best way of conserving natural resources.

References

Alston, Lee J. and Andersson, Krister (2011): "Reducing Greenhouse Gas Emissions by Forest Protection: The Transaction Costs of implementing REDD". *Climate Law* 2(2): 281-89.

Angelsen, Arild (2010): "The 3 REDD 'I's," *Journal of Forest Economics* 16(4): 253-56.

Angelsen, Arild and Kaimowitz, David (1999): "Rethinking the Causes of Deforestation: Lessons from Economic Models," *The World Bank Research Observer* 14(1): 73-98.

Bolton, Patrick and Dewatripont, Mathias (2005): *Contract Theory*, MIT Press.

Burgess, Robin; Hansen, Matthew; Olken, Ben; Potapov, Peter and Sieber, Stefanie (2011): "The Political Economy of Deforestation in the Tropics," *Quarterly Journal of Economics* 127(4): 1707-54.

Chiroleu-Assouline, M., Poudou, J.-C., and Roussel, S. (2012): "North/South Contractual Design through the REDD+ Scheme," FEEM Working Paper, (89):106–112.

Delacote, P. and Simonet, G. (2013): "Readiness and Avoided Deforestation Policies: on the use of the REDD Fund," Working Paper 1312, Chaire Economie du Climat.

Damette, Olivier and Delacote, Philippe (2012): "On the economic factors of deforestation: What can we learn from quantile analysis?" *Economic Modelling* 29 (6): 2427-34.

Engel, Stefanie; Pagiola, Stefano and Wunder, Sven (2008): "Designing payments for environmental services in theory and practice: An overview of the issues," *Ecological Economics* 65(4): 663-74.

Gjertsen, H., Groves, T., Miller, D. A., Niesten, E., Squires, D., and Watson, J. (2016): "The Optimal Structure of Conservation Agreements and Monitoring," Department of Economics, University of California San Diego.

FAO (2010): *State of the World Forests*. FAO, Rome.

Harstad, Bård (2012): "Buy Coal! A Case for Supply-Side Environmental Policy," *Journal of Political Economy* 120(1): 77-115.

Harstad, Bård, and Mideksa, Torben (2016): "Conservation contracts and political regimes," forthcoming, *Review of Economic Studies*.

Hotelling, Harold (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy* 39(2): 137-75.

IPCC (2007): *The Fourth Assessment Report*, Bonn, Germany: IPCC.

IPCC (2013): *The Fifth Assessment Report*, Cambridge University Press.

Maskin, Eric and Tirole, Jean (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory* 100(2): 191-219.

Mason, C. F. (2013): "Optimal Contracts for Discouraging Deforestation with Risk Averse Agents,".Mimeo.

Mason, C. F. and Plantinga, A. J. (2013): "The Additionality Problem with Offsets: Optimal Contracts for Carbon Sequestration in Forests," *Journal of Environmental Economics and Management*, 66(1):1-14.

Segal, I. (1999): "Contracting with Externalities," *Quarterly Journal of Economics*, 114(2):337-88.

Sinn, Hans-Werner (2008): "Public Policies against Global Warming: A Supply Side Approach," *International Tax and Public Finance* 15:360-94.

Sinn, Hans-Werner (2012): *The Green Paradox*, MIT Press.

Stern, Nicholas (2008): "The Economics of Climate Change," *American Economic Review* P&P 98(2): 1-37.

Tirole, Jean (1998): *The Theory of Industrial Organization*, 10th edition. Cambridge, MA: The MIT Press.